A Reputational Theory of Firm Dynamics*

SIMON BOARD† AND MORITZ MEYER-TER-VEHN‡

This Version: June 13, 2014
First Version: April 2010.

Abstract

We propose a firm lifecycle model in which the firm privately invests in its quality and thereby its reputation. Over time, both the firm and the market learn about the firm's evolving quality via infrequent breakthroughs. The firm can also exit if its value becomes negative, giving rise to selection effects. In a pure-strategy equilibrium, incentives are single-peaked: the firm shirks immediately following a breakthrough, works for intermediate levels of reputation and shirks again when it is about to exit. This investment behavior yields predictions for the distribution of firm productivity and the survival rate. Finally, we compare the model to two variants: one in which the firm's investment is publicly observed, and a second in which the firm has private information about its product quality.

1 Introduction

> "New technologies come and go. This is simply the nature of our business. The Samsung brand is the only asset that will live on beyond our products."
>
> Sue Shim, CMO at Samsung.

Models of firm dynamics seek to generate the large variability in productivity and profits seen within industries, whereby some firms invest in their assets and grow, while others disinvest and shrink (e.g. Syverson (2011)). One of a firm's most important assets is its reputation. Philip Kotler writes, "In my field of marketing, brand reputation is everything," with Interbrand valuing Apple's brand at \$98bn (from a market cap of \$475bn). Furthermore, surveys of Boards of Directors reveal that reputation risk, in particular pertaining to product quality and the firm's public perception, is one of their primary concerns.[1]

*We have received helpful comments from Andy Atkeson, Heski Bar-Isaac, V. Bhaskar, Alessandro Bonatti, Andrew Clausen, Christian Hellwig, Hugo Hopenhayn, Johannes Hörner, Yuliy Sannikov and seminar audiences at Bocconi, Essex, LSE, Princeton, PSE, Toronto, TSE, Warwick, the ES Winter Meetings, Gerzensee, Duke IO Theory Conference, Mannheim Reputation Conference, SAET, SED, SWET. We gratefully acknowledge financial support from NSF grant 0922321. Keywords: Reputation, Self-esteem, Exit, Lifecycle, Brands, Firm dynamics, Career concerns. JEL: C73, L14

†Department of Economics, UCLA. http://www.econ.ucla.edu/sboard

‡Department of Economics, UCLA. http://www.econ.ucla.edu/mtv

[1]Sources: The Kotler quote comes from Diermeier (2011). The Shim quote and Interbrand's brand value comes from its 2013 Global Brand Survey; this report seeks to measure a product's demand after controlling for price and product features. The final line comes from EisnerAmper's 2012 Board of Directors' Survey.

This paper proposes a model of lifecycle dynamics in which a firm's quality and its reputation are its most important assets. We suppose a firm privately- chooses its level of investment, while the market and firm learn about the resulting quality. We analyze the firm's optimal investment and exit strategy to obtain predictions concerning the distribution of firm revenue product, and the resulting selection effects. Consequently, this is a natural lens through which to study the evolution of reputation data, such as JD Power scores, Yelp reviews, and Facebook Likes.

Our model draws inspiration from two canonical models of firm dynamics. In Jovanovic (1982) firms, along with consumers, learn about their quality over time; in Ericson and Pakes (1995) firms invest in the quality of their products and are subject to idiosyncratic shocks. We combine learning and investment, thereby putting reputation at center stage. A firm's reputation reflects the market's belief about quality, and is therefore very different from traditional capital assets. First, reputation evolves according to Bayes' rule and therefore may be volatile even though the underlying product quality is fairly stable. For example, movie reviews can dramatically change box office numbers even though the movie's quality is constant. Second, reputation depends on the market's *beliefs* about the firm's investment, rather than actual investment. As the firm cannot control these beliefs, investment incentives are dampened by moral hazard. Moreover, as we explore in this paper, investment incentives critically depend on the structure of information in the market.

In the model, a long-lived firm sells a product of high or low quality to a continuum of identical short-lived consumers, and product quality is a stochastic function of the firm's past investments. Consumers observe neither the firm's investment history nor the resulting quality. Rather, they learn about quality via public breakthroughs that can only be produced by a high-quality product; the market's belief that quality is high, x_t, is called the *reputation* of the firm. Likwise, the firm does not observe quality directly, but learns about it via the same breakthroughs; unlike the market, it can recall its past investments. The firm's belief that quality is high, z_t, is called the *self-esteem* of the firm. In a pure-strategy equilibrium, reputation and self-esteem coincide on path. At each point in time, the consumers' willingness to pay and the firm's revenue equals its reputation. The firm can exit the market at any time, and it does so when its reputation falls below some threshold. We aim to characterize the firm's optimal investment and exit decisions over its lifecycle. Since the market can only observe breakthroughs, we consider recursive equilibria where reputation is solely a function of the time since the last breakthrough.

To understand the firm's optimal investment strategy, note that investment raises the firm's self-esteem but, since investment is not observed by the market, does not affect its reputation. Hence investment incentives are determined by the marginal value of self-esteem. In turn, this marginal value can be characterized as an integral over future reputational dividends that derive from an increased chance of breakthroughs.

One problem with this characterization of investment incentives is that this integral itself depends on the knowledge of optimal future investment, and does not yield a sufficient condition for an optimal investment strategy. Therefore, we prove existence of optimal strategies and equilibrium via a different route, by defining an adequate topology on the strategy space and applying the Kakutani-Fan-Glicksberg theorem.

We then show that in any equilibrium, the firm shirks when its reputation is near the exit threshold. Intuitively, investment only pays off if investment affects quality *and* quality is revealed via a breakthrough. When the firm is ε from exit, the probability of these two events is of order ε^2, and therefore the firm prefers to shirk. The exit threshold itself is determined by an indifference condition where the option value of obtaining a breakthrough offsets the immediate losses in operating profit.

In a pure-strategy equilibrium, investment incentives are single-peaked in the time since a breakthrough. This means that any nontrivial equilibrium is work-shirk or shirk-work-shirk. In the latter, the firm shirks upon witnessing a breakthrough and obtaining a high reputation, works for intermediate levels of reputation, and shirks again when its reputation comes close to the exit threshold. Intuitively, as the firm's reputation falls the benefit of a breakthrough grows, leading to increasing incentives; however, the firm also gets closer to the exit point, leading to decreasing incentives. We show that the sum of these two effects is single-peaked. After proving this qualitative property, we use the baseline model to simulate the lifecycle of a stylized restaurant. This exhibits a shirk-work-shirk equilibrium and demonstrates how reputation quickly declines when the firm is believed to be shirking, and slowly declines when the firm is believed to be working.

In the final part of the paper, we consider two variants of the baseline model (see Figure 1), illustrating how the information structure in the market affects the firm's incentives to invest. In addition, these variants allow us isolate the two different economic forces that lead to single-peaked incentives in the baseline model. In the first variant ("observable investment"), the market observes the firm's investment, while both the firm and market must still learn about the firm's quality over time. Here, investment incentives decrease monotonically as time passes without a breakthrough and the firm approaches the exit threshold. When compared to the baseline model, the elimination of moral hazard means that investment increases both a firm's self-esteem and its reputation. In equilibrium, this increases the amount of investment, slows the decline in reputation, delays the exit time and raises the firm's value. In the second variant ("known quality"), we suppose the firm is privately- informed about its own quality, on which it can condition its exit and investment choices. In equilibrium, the high-quality firm can then signal its type by remaining in the market, and the low-quality firm randomizes over exit times in order to keep reputation constant, as in Bar-Isaac (2003). Investment incentives then increase monotonically over time, as the value of a breakthrough increases. Unlike the baseline model, investment incentives stay large close to the exit threshold because the firm can immediately- see if investment pays off, and will then choose to remain in the market. Combining the increasing incentives in the "known quality" variant and the decreasing incentives in the "observable investment" variant yields the single-peaked incentives in the baseline model.

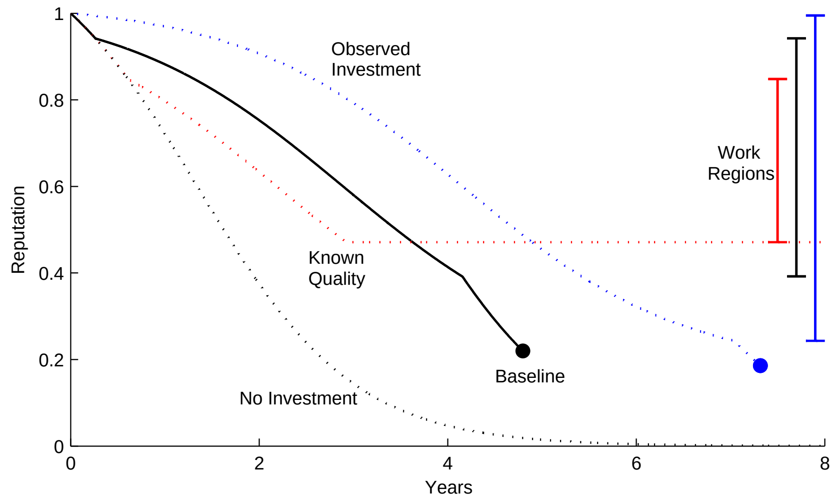

Figure 1: **Reputation Trajectories.** This figure shows the equilibrium trajectory of reputation under the baseline model, the two extensions − where the market observes the firm's investment, and where the firm knows its own quality − and the trajectory of reputation if the firm never invests. The parameters are described in Section 4.

1.1 Literature

The paper embeds the reputation framework of Board and Meyer-ter-Vehn (2013) into a firm lifecycle model. In contrast to that paper we assume that the firm does not observe its quality, introducing a role for the firm's self-esteem. The possibility of exit qualitatively changes the firm's investment incentives. The model variant with known quality also draws heavily on Bar-Isaac (2003), adding investment to that model.

We contribute to the growing literature on learning models with moral hazard; such models have the feature that private and public beliefs differ off-path. Kovrijnykh (2007) introduces exit into a three-period career concerns model. Bonatti and Hörner (2011, 2013) consider effort incentives in a strategic experimentation game. Sannikov (2014) considers a contract design problem in which the agent's effort has long-run effects on her employer's performance. Cisternas (2014) analyzes a general model of two-sided learning with moral hazard; his incentive equation is analogous to our "marginal value of self-esteem."

In addition to Jovanovic (1982) and Ericson and Pakes (1995), there are a variety of other models of firm dynamics. Hopenhayn (1992) assumes firm capabilities change over time according to a Markov process, and looks at the resulting entry and exit patterns. Cabral (2014) and Abito, Besanko, and Diermeier (2012) consider reduced-form models of reputational firm dynamics, whereby reputation is modeled as a state variable akin to capital stock, but is not derived from Bayes' rule. Gale and Rosenthal (1994) and Rob and Fishman (2005) consider the dynamics of repeated games equilibria where incentives arise from punishment strategies. On the empirical side, Foster, Haltiwanger, and Syverson (2013) model the slow demand growth of new entrants by assuming the level of current demand depends on the stock of past demand, which the authors interpret as the "growth of customer base or building a reputation." Bronnenberg, Dubé, and Gentzkow (2012) use a model of brand capital to study the dynamics of brand shares when customers move between cities and their preferences depend on past purchases.

2 The Baseline Model

Players and actions: There is one long-lived firm and a continuum of short-lived consumers, also referred to as *the market*. Time $t \in [0, \infty)$ is continuous. At every time t the firm chooses an *investment* level $A_t \in [0, \overline{a}]$ where $\overline{a} < 1$; it may also choose to exit the market, thereby ending the game. Following Holmström (1999) and Mailath and Samuelson (2001), the firm produces one unit of output at every time, and consumers are assumed to purchase the firm's output at a price equal to their willingness to pay.

At time t the firm's product *quality* is $\theta_t \in \{L, H\}$, where $L = 0$ and $H = 1$. Initial quality θ_0 is exogenous; subsequent quality depends on investment and technology shocks. Specifically, shocks are generated according to a Poisson process with arrival rate $\lambda > 0$. Quality θ_t is constant between shocks, and determined by the firm's investment at the most recent technology shock $s \leq t$; i.e., $\theta_t = \theta_s$ and $\Pr(\theta_s = H) = A_s$. This captures the idea that quality is a lagged function of past investments.

Information: Consumers observe neither quality nor investment, but learn about quality through public *breakthroughs*. Given quality θ, breakthroughs are generated according to a Poisson process with arrival rate $\mu\theta$. We write h^{t-} for histories of breakthrough arrival times before time t, h for infinite histories, and $\varnothing$ for histories with no breakthroughs.

The firm does not observe product quality either, but does recall its past actions. It chooses an investment plan $A = \{A_t\}_{t \geq 0}$ and an exit time $T \in [0, \infty)$ that is predictable with respect to the associated filtration. From the firm's perspective, investment A controls the distribution of quality $\{\theta_t\}_{t \geq 0}$ and thereby the histories of breakthroughs h; we write $\mathbb{E}^A$ for expectations under this measure and call $Z_t = \mathbb{E}^A \left[\theta_t | h^t\right]$ the firm's *self-esteem* at time $t < T$. This reflects the firm's belief of its own quality given its past investment and the history of breakthroughs.

We write pure market beliefs over investment and exit as $\tilde{A} = \{\tilde{A}_t\}_{t \geq 0}$ and $\tilde{T}$, and mixed beliefs as distributions over pure beliefs $F = F(\tilde{A}, \tilde{T})$. If market beliefs are focused on a unique strategy $(\tilde{A}, \tilde{T})$ its belief about quality, the firm's *reputation*, is given by $X_t := \mathbb{E}^{\tilde{A}}[\theta_t | h^t]$ as long as $t < \tilde{T}$. If the market holds mixed beliefs $F(\tilde{A}, \tilde{T})$ it faces an additional layer of uncertainty. Such a distribution induces a joint distribution over $\{\theta_t\}_{t \geq 0}$, h, and exit times $\tilde{T}$; writing $\mathbb{E}^F$ for expectations under this measure, the firm's reputation is given by $X_t = \mathbb{E}^F[\theta_t | h^t, \tilde{T} > t]$ for all $t < T(F)$, where $T(F) := \min\{t : F(\tilde{T} \leq t) = 1\}$ is the first time at which the market expects the firm to exit with certainty. When the firm fails to exit, that is at times $t \geq T(F)$, the market revises its beliefs about the firm's strategy to some arbitrary $F'(\tilde{A}, \tilde{T})$ with $t < T(F')$ and reputation equals $X_t = \mathbb{E}^{F'}[\theta_t | h^t, \tilde{T} > t]$; initially, self-esteem and reputation are exogenous and coincide, $X_0 = Z_0$.

Payoffs: The firm and consumers are risk-neutral and discount future payoffs at rate $r > 0$. At time t, the firm produces one unit with flow value θ_t. Given the public information h^{t-}, consumers' willingness to pay then equals the firm's reputation X_t. We assume that the price equals the willingness to pay, so consumers' expected utility is 0. Investment has a constant marginal flow cost of $c > 0$ and the firm's operating costs equal $k \in (0, 1)$. The firm's flow profits are thus given by $X_t - k - cA_t$.

Given the firm's strategy (A, T) and market beliefs about this strategy $F(\tilde{A}, \tilde{T})$, the firm's expected present value equals

$$\mathbb{E}^A \left[\int_{t=0}^{T} e^{-rt}(X_t - k - cA_t)dt \right]. \tag{2.1}$$

Market beliefs determine the firm's revenue $X_t = \mathbb{E}^F[\theta_t | h^t, \tilde{T} > t]$ for a given history h, actual investment A determines the distribution over histories of breakthroughs h, and the exit time T determines the integration domain.

To ensure the analysis is interesting we assume throughout the paper that

$$z^\dagger := \lambda/\mu < 1 \quad \text{and} \quad z^\dagger - k + \mu z^\dagger(1-k)/r < 0. \tag{2.2}$$

This assumption ensures that, in the absence of a breakthrough, the firm's reputation declines and the firm eventually exits.

2.1 Recursive Strategies

Both reputation and self-esteem are reset to $X = Z = 1$ at a breakthrough; between breakthroughs the market observes no information about the firm's performance. For this reason, we consider recursive strategies which only depend on the time since the last breakthrough. Formally, we call strategy (A, T) *recursive* if there exists a deterministic process $a = \{a_t\}$ with $a_t \in [0, \overline{a}]$ and $\tau \in [0, \infty]$ such that if the last breakthrough before t was at $s < t$, then $A_t = a_{t-s}$ and $T \leq t$ iff $\tau \leq t - s$. We write recursive strategies as (a, τ) and the resulting self-esteem as $z = \{z_t\}$, where $z_0 = 1$. Similarly we call beliefs F recursive if they assign probability one to recursive strategies $(\tilde{a}, \tilde{\tau})$, and denote the induced reputation by $x_t = \mathbb{E}^F[\theta_t | h^t = \varnothing, \tilde{\tau} > t]$, where $x_0 = 1$.

Self-esteem evolves according to Bayes' rule. Thus, at a breakthrough, self-esteem jumps to one. Absent a breakthrough, self-esteem is governed by $\dot{z}_t = g(a_t, z_t)$ where the drift g is given by

$$g(a_t, z_t) = \lambda(a_t - z_t) - \mu z_t(1 - z_t). \tag{2.3}$$

The first term derives from the technology process: with probability λdt a technology shock hits in $[t, t + dt)$, previous quality becomes obsolete, and the current quality is determined by the firm's investment. This term is positive if investment a_t exceeds the firm's self-esteem z_t and negative otherwise. The second term derives from the absence of breakthroughs.

Given recursive reputation $x = \{x_t\}$, we can restrict the firm's problem to recursive strategies. From the firm's perspective breakthroughs arrive at rate μz_t, so we truncate the integral in (2.1) at the first breakthrough to see that the firm's continuation value at time t is given by

$$V(t, z_t) = \sup_{a,\tau} \int_{s=t}^{\tau} e^{-\int_t^s (r+\mu z_u)du}(x_s - k - ca_s + \mu z_s V(0,1))ds. \tag{2.4}$$

We write optimal recursive strategies as (a^*, τ^*) and the associated self-esteem as $z^* = \{z_t^*\}$.

Lemma 1 (Existence of Optimal Strategy) *For any reputation $x = \{x_t\}$, an optimal strategy (a^*, τ^*) maximizing (2.4) exists. Furthermore, there exists $\overline{\tau} < \infty$ such that for all x and optimal strategies (a^*, τ^*), we have $\tau^* \leq \overline{\tau}$.*

Proof. We prove this lemma by defining a topology on the space of strategies (a, τ) such that the space is compact and the objective function in (2.4), is continuous in (a, τ) in this topology.

For the compactness argument, we first argue the second part of the Lemma, that there exists $\overline{\tau} < \infty$ at which the firm exits for any reputational trajectory x induced by beliefs F. Given $\overline{a} < 1$ and assumption (2.2), reputational drift $g(a_t, z_t)$ is negative and bounded away from zero for $z_t \in [z^\dagger, 1]$ and any $a_t \in [0, \overline{a}]$. Write $z_t(a)$ for the solution of $\dot{z}_t = g(a_t, z_t)$ for fixed investment strategy $a = \{a_t\}$ and define $\overline{\tau}$ such that $z_{\overline{\tau}}(\overline{a}) = z^\dagger$. Hence $z_t(a) \leq z^\dagger$ for all $t \geq \bar{\tau}$ and any a. By the law of iterated expectations, reputation $x_t = \mathbb{E}^F[z_t(a)|t < \tilde{\tau}, h^t = \varnothing]$ is a conditional expectation over self-esteem and is thus bounded by $x_t \leq z_t(\bar{a})$ both on-path and off-path. Thus, $x_t \leq z^\dagger$ for all $t \geq \bar{\tau}$. Equation (2.1) implies $V(0,1) < (1-k)/r$, and so the integrand in (2.4) is bounded above by $z^\dagger - k + \mu z^\dagger (1-k)/r$, which is negative by assumption (2.2). Hence the firm exits by time $\bar{\tau}$, as required.

To prove the existence of an optimal strategy, let B be the space of measurable investment functions $\{a_t\}_{t \in [0, \overline{\tau}]}$. B is naturally embedded in the (rescaled) unit ball of $L^2([0, \overline{\tau}], \mathbb{R})$. In the weak topology this unit ball is compact by Alaoglu's theorem, and as a closed subset of this unit ball, B is also compact. In this topology a sequence $\{a_t^n\}$ converges to $\{a_t\}$ if $\int_0^{\overline{\tau}} (a_t^n - a_t)\xi_t dt \to 0$ for all test functions $\xi \in L^2([0, \overline{\tau}], \mathbb{R})$. While this topology is coarse enough to make B compact, it is fine enough for the trajectory $\{z_t\}_{t \in [0, \overline{\tau}]}$ to be continuous (in the sup-norm) in $\{a_t\}_{t \in [0, \overline{\tau}]}$ (see Davis (1993, Theorem 43.5)). Thus, firm value (2.4) is continuous in $(\{a_t\}, \tau)$ and is maximized by some $(\{a_t^*\}, \tau^*)$. $\square$

Remarks: The central ingredients of our model are investment and learning, and one may wonder how our analysis depends on the specification of these two processes. For the investment process, we assume that the firm's previous quality becomes obsolete at a breakthrough. However our analysis extends to models in which high-quality firms, say, are more successful at investing as long as quality at a breakthrough is separable in previous quality and current investment. For example, a model with $\Pr(\theta_s = H) = (A_s + \theta_{s-})/2$ is isomorphic to our model with arrival rate of technology shocks equal to $\lambda/2$. One may also question how our assumption of potentially declining quality fits some applications, e.g. technology firms. To apply our model to such settings, one can think of a firm's quality as its advantage over a competitive fringe that advances one rung on a quality ladder at each technology shock.

On the learning side, we focus on breakthroughs that reveal high quality with certainty, termed "perfect good news learning" in Board and Meyer-ter-Vehn (2013), for two reasons. First, it is tractable because it makes the model recursive in the time since the last breakthrough. Second, it precludes downward jumps in self-esteem and reputation, allowing us to study investment when exit is imminent. The intuition for our qualitative results $-$ that investment incentives are low when reputation is close to perfect or near the exit point $-$ extends to imperfect good news Poisson or Brownian signals; in comparison, under "perfect bad news" exit would only follow a breakdown. While formally proving qualitative properties would be much harder in these models, the firm's optimal investment condition (3.1), and thus our numerical solution strategy, remains valid for these processes.[2]

The baseline model considers the lifecycle of a single firm. Following Atkeson, Hellwig, and Ordonez (2012), one can extend the model to a competitive market with a continuum of firms where consumers care about aggregate quality $\int_i x_{t,i} di$ and the market price is determined by free entry. If one is interested in the restaurant industry, one could model such entry by assuming these new firms pay a fixed cost and start with an exogenous reputation and self-esteem, analogous to Jovanovic (1982). Alternatively, if one is interested in the labor market for professionals, one could assume there is a distribution of admits each year whose reputation and self-esteem must exceed a cutoff.

3 The Firm's Problem

In this section we analyze the firm's optimal strategy, (a^*, τ^*) for arbitrary (recursive) beliefs, inducing reputation $x = \{x_t\}$. First, we characterize the firm's investment incentives as the integral of a series of dividends that result from having higher self-esteem. Second, we explore the qualitative properties of the firm's strategy, showing that the firm shirks when it is close to bankruptcy and, if reputation decreases over time, investment incentives are single-peaked. In Section 4 we close the model in equilibrium by assuming that $\{x_t\}$ is derived from correct market beliefs $F(\tilde{a}, \tilde{\tau})$.

Investment does not directly affect reputation, but raises the firm's self-esteem and thereby raises the probability of breakthroughs, which do boost reputation and revenue. Using (2.3), a small increase in investment raises self-esteem by λ. Assuming that the value function V is differentiable with respect to z, the marginal benefit of investing at time t is $\lambda V_z(t, z_t^*)$; thus optimal investment must satisfy

$$a_t^* = \begin{cases} 0 & \text{if} \quad \lambda V_z(t, z_t^*) < c, \\ \bar{a} & \text{if} \quad \lambda V_z(t, z_t^*) > c. \end{cases} \tag{3.1}$$

Next, observe that the firm's value $V(t, z)$ is convex in z. This follows because z is the firm's private belief about its quality, and the value of information is convex since firms with extreme values of

[2] As there is no "time since the last breakthrough t" for such learning processes, the natural Markovian state variables are reputation x_t and self-esteem z_t. These states capture all payoff-relevant information if the market holds point beliefs about investment, e.g. in a pure strategy equilibrium. Otherwise, if the market has mixed beliefs about investment, x_t is not a sufficient statistic for market beliefs and one would have to keep track of the market's belief about z_t, that is, the market's second-order belief about quality.

z can attain the same average value as a firm with moderate z by mimicking its strategy.[3] This result has two implications. First, even where it is not differentiable, the value function admits directional derivatives $V_{z^-}(t,z), V_{z^+}(t,z)$. Second, investment today and investment in the future are dynamic complements. Intuitively, investment today raises the firm's self-esteem z_{t+dt} and its remaining time in the industry; this raises the marginal benefit of self-esteem and investment tomorrow since the firm has a greater chance of benefiting from the resulting breakthroughs. This strategic complementarity stands in contrast to the strategic substitutability in Bonatti and Hörner (2011). There, a player who exerts more effort today is more pessimistic about the state of the project tomorrow when his effort fails to result in the desired breakthrough. Here, to the contrary, investment today makes the firm more optimistic about its prospects tomorrow.

We next express the marginal value of self-esteem in terms of future reputational dividends. This expression is the work horse of our paper, helping us study optimal investment.

Lemma 2 (Marginal Value of Self-Esteem) *If $V_z(t,z_t^*)$ exists, it equals*

$$\Gamma(t) := \int_t^{\tau^*} e^{-\int_t^s r+\lambda+\mu(1-z_u^*)du} \mu \left[V(0,1) - V(s,z_s^*)\right] ds. \tag{3.2}$$

More generally $V_{z^-}(t,z_t^*) \leq \Gamma(t) \leq V_{z^+}(t,z_t^*)$.

Proof. This follows by applying the envelope theorem to a variant of (2.4). See Appendix A.1. □

Equation (3.2) is an integral version of the adjoint equation for the firm's control problem. Intuitively, self-esteem raises the probability of a breakthrough and, since it is persistent, pays off dividends over time. That is, incremental self-esteem dz raises the probability of a breakthrough in $[t, t+dt]$ by $\mu dz dt$; the value of a breakthrough equals $V(0,1) - V(t,z_t^*)$. We thus call the integrand $\mu(V(0,1) - V(s,z_s^*))$ the *reputational dividend of self-esteem.* The dividend stream from the increment dz depreciates for three reasons. First, time is discounted at rate r; second, at rate μz_t^* a breakthrough arrives, self-esteem jumps to one, and the increment disappears; third, reputational drift (2.3) is not constant in z, and its derivative equals $g_z(a_t^*, z_t^*) = -(\lambda + \mu(1-2z_t^*))$. Summing these three components yields the discounting term seen in (3.2).

If $V(t,z_t^*)$ is differentiable in z, then its derivative coincides with $\Gamma(t)$. If there are multiple optimal strategies (a^*, τ^*) for which $\Gamma(t)$ does not coincide, then V is not differentiable at (t, z_t^*). However, for any (a^*, τ^*) (3.2) is well-defined and bounded by the directional derivatives of V. Hence if $\lambda\Gamma(t) > c$ then $\lambda V_{z^+}(t,z_t^*) > c$ and the firm finds it profitable to work, whereas if $\lambda\Gamma(t) < c$ then $\lambda V_{z^-}(t,z_t^*) < c$ and the firm finds it profitable to shirk. This implies:

Lemma 3 (Optimal Investment) *Given $x = \{x_t\}$, any optimal strategy (a, τ) satisfies*

$$a_t^* = \begin{cases} 0 & \text{if } \lambda\Gamma(t) < c \\ \overline{a} & \text{if } \lambda\Gamma(t) > c \end{cases} \tag{3.3}$$

[3] More formally, let (a^*, τ^*) be an optimal strategy for a firm with self-esteem z at time t. If two neighboring firms, high and low, with initial states $(x, z+\varepsilon)$ and $(x, z-\varepsilon)$ both mimic strategy (a^*, τ^*), then their average expected payoff equals $V(t,z)$. Since these firms can weakly raise their payoff by reoptimizing, V is weakly convex in z.

for almost all t.[4]

Fixing a candidate strategy (a, τ), equation (3.3) gives a necessary condition for a best response. However, this is not sufficient since this approach is the continuous-time analogue to checking only "one-step deviations on path." Since actions are dynamic complements, the possibility of multi-stage deviations must be taken seriously.

With these preliminary results, we can characterize the qualitative properties of equilibrium.

Theorem 1 (Shirk at End) *Given* $x = \{x_t\}$ *and any optimal strategy* (a^*, τ^*), *there exists* $\varepsilon > 0$ *such that* $a_t^* = 0$ *for almost all* $t \in [\tau^* - \varepsilon, \tau^*]$.

Proof. Lemma 2 implies that $\Gamma(t)$ is of order $O(\tau^* - t)$, and so $\lambda\Gamma(t) < c$ for $t \in [\tau^* - \varepsilon, \tau^*]$. The theorem thus follows from Lemma 3. $\square$

Intuitively, investment incentives vanish at the exit time τ^* because there is no time left for the investment to pay off. More formally, the benefit of investment is of second order because both a technology shock and a breakthrough must arrive in the remaining time interval for the investment to avert exit.

When the firm is close to exit it will therefore cease to invest, accelerating its demise. For example, in the beer industry, Goldfarb (2007) argues that Schlitz realized that the rise of Miller would have a large impact on its future profitability. The firm chose therefore to change the preservatives, switch to lower quality accelerated batch fermentation, disinvest in the brand, and fire much of its marketing team. More generally, this implies that the death of firms is a quick process, with reputation quickly declining.[5]

Next, we assume that reputation $x = \{x_t\}$ strictly decreases in time. This assumption means the longer the firm fails to prove itself by generating a breakthrough, the more pessimistic the market becomes about its product quality. This is satisfied if market beliefs $F(\tilde{a}, \tilde{\tau})$ are focused on a single strategy $(\tilde{a}, \tilde{\tau})$ and the market draws no positive inference about $\{\tilde{a}_t\}_{t \leq \tilde{\tau}}$ in the off-path event that the firm fails to exit at time $\tilde{\tau}$.[6]

When $\{x_t\}$ strictly decreases, the firm's value $V(t, z)$ strictly decreases in t and strictly increases in z. Intuitively, value decreases in t because a low-t firm can mimic the strategy of a high-t firm, yielding the same probability of a breakthrough, but higher revenue prior to a breakthrough. Similarly, value increases in self-esteem because a firm with high self-esteem can mimic a firm with low self-esteem, yielding the same revenue prior to a breakthrough, but a higher probability of a breakthrough.[7]

[4] The firm can change its strategy at a measure zero set of times without affecting payoffs, so any statements about optimal investment hold only almost always; we included this technical qualification in the lemma and theorem statements, but omit it in the text. It can be eliminated by restricting the firm to forward-continuous strategies (see Board and Meyer-ter-Vehn (2013)).

[5] This contrasts with the model variant with known quality analyzed in Section 5.2.

[6] This assumption is not satisfied in a mixed strategy equilibrium, where reputation rises when the low-investment firm exits (see Section 4.2).

[7] See Appendix A.2 for a formal argument.

The next lemma computes the partial derivative $V_t(t, z_t^*)$. We then use this in Theorem 2 to show that investment incentives are single-peaked.

Lemma 4 (Marginal Value of Reputation) *Assume that $\{x_t\}$ strictly decreases. Whenever the partial derivative $V_t(t, z_t^*)$ exists, it is equal to*

$$\Psi(t) := \int_t^{\tau^*} e^{-\int_t^s r+\mu z_u^* du} dx_s. \tag{3.4}$$

Moreover, $\Psi(t) < 0$ for $t < \tau^*$.

Proof. Rewrite the firm's continuation value (2.4) by writing $\sigma = s - t$ for the time since t and $(\{a_\sigma^*\}, \zeta^*)$ for the optimal strategy starting at t. Then,

$$V(t, z_t^*) = \int_{\sigma=0}^{\zeta^*} e^{-\int_0^\sigma (r+\mu z_{t+u}^*)du}(x_{t+\sigma} - k - ca_\sigma^* + \mu z_{t+\sigma}^* V(0, 1))d\sigma.$$

As $z_{t+\sigma}^*$ is determined by initial self-esteem z_t^* and $\{a_\sigma^*\}$, it is independent of t. The envelope theorem thus yields (3.4). As $\{x_t\}$ is assumed to strictly decrease, $dx_s < 0$ and hence $\Psi(t)$ must be negative. $\square$

Theorem 2 (Single-Peaked Incentives) *If $\{x_t\}$ strictly decreases, then investment incentives $\Gamma(t)$ are single-peaked, with boundary conditions $\Gamma(0) > 0$, $\dot{\Gamma}(0) > 0$ and $\Gamma(\tau^*) = 0$.*

Proof. Taking the derivative of investment incentives (3.2) and setting $\rho(t) := r + \lambda + \mu(1 - z_t^*)$ yields the adjoint equation

$$\dot{\Gamma}(t) = \rho(t)\Gamma(t) - \mu(V(0, 1) - V(t, z_t^*)). \tag{3.5}$$

Now assume that $\rho(t)$ and $V(t, z_t^*)$ are differentiable. Then $\dot{\rho}(t) = -\mu \dot{z}_t^*$ and $\frac{d}{dt}V(t, z_t^*) = \dot{z}_t^* \Gamma(t) + \Psi(t)$; in Appendix A.3 we show that these functions are indeed absolutely continuous and extend our arguments to that case. The derivative of the adjoint equation equals

$$\begin{aligned} \ddot{\Gamma}(t) &= \rho(t)\dot{\Gamma}(t) + \dot{\rho}(t)\Gamma(t) - (-\mu \frac{d}{dt} V(t, z_t^*)) \\ &= \rho(t)\dot{\Gamma}(t) - \mu \dot{z}_t^* \Gamma(t) + \mu \dot{z}_t^* \Gamma(t) + \Psi(t) \\ &= \rho(t)\dot{\Gamma}(t) + \Psi(t) \end{aligned} \tag{3.6}$$

Since $\Psi(t) < 0$, $\dot{\Gamma}(t) = 0$ implies $\ddot{\Gamma}(t) < 0$, hence $\Gamma(t)$ is single-peaked.

At $t = \tau^*$, equation (3.2) immediately- implies that $\Gamma(\tau^*) = 0$. At $t = 0$, equation (3.2) implies $\Gamma(0) > 0$ because the integrand $\mu(V(0, 1) - V(s, z_s^*))$ is strictly positive for $s \in [0, \tau^*]$. Equation (3.5) then implies $\dot{\Gamma}(0) = \rho(0)\Gamma(0) > 0$. $\square$

Theorem 2 implies that the optimal strategy takes one of three forms:

1. *Full-shirk.* The firm shirks for almost all t.

2. *Shirk-work-shirk.* The firm shirks immediately- following a breakthrough, then works for intermediate reputations, and shirks near the exit time.

3. *Work-shirk.* The firm works after a breakthrough, but shirks when they are close to exit.

Intuitively, the evolution of investment incentives is shaped by two countervailing forces. On the downside, as t increases the firm forgoes the reputational dividends over $[t, t + dt]$, as captured by the second term in (3.5). This negative effect becomes more important over time as the reputational dividend increases, as captured by the positive term $-\mu \frac{d}{dt} V(t, z_t^*)$ in (3.6). On the upside, an increase in t brings future and larger dividends closer, as captured by the first term in (3.5). Ignoring the time dependence of $\rho(t)$, this positive effect becomes less important over time once incentives start decreasing. Thus once incentives decrease, the negative effect keeps growing while the positive effect decreases, and so incentives decrease until exit.

Theorem 2 is a surprisingly robust result. First, much of the analysis would be identical for convex cost function $c(a)$; then single-peaked incentives would translate into single-peaked investment. Second, the result only requires that payoffs $\{x_t\}$ decrease over time; thus it holds if, say, high-reputation firms make more sales, implying revenue is convex in reputation. Third, the result applies to the firm's optimal strategy – treating payoffs $\{x_t\}$ as an exogenous process – and therefore does not require that market beliefs be correct.

Turning to the firm's exit behavior, we next assume that reputation $\{x_t\}$ is continuous. In equilibrium, reputation $\{x_t\}$ must be continuous on-path. Off-path reputation is continuous if the market does not draw inferences about previous investment from a failure to exit.

Theorem 3 (Exit Time) *If $\{x_t\}$ is continuous, the optimal exit time τ^* satisfies*

$$x_{\tau^*} - k + \mu z_{\tau^*} V(0, 1) = 0. \tag{3.7}$$

Proof. Recall that the firm's value is given by (2.4). When the firm shirks, its flow payoff is $x_t - k$, and its option value of staying in the market has a flow value of $\mu z_t^* V(0, 1)$. Thus, if $x_t - k + \mu z_t^* V(0, 1) > 0$, then the firm can secure itself strictly positive payoffs by shirking and staying in the market until (3.7) holds. Conversely, if $x_{\tau^*} - k + \mu z_{\tau^*}^* V(0, 1) < 0$ then the continuity of $\{x_t\}$ implies this inequality also holds for t just before τ^*, and the firm would have been better off exiting a little earlier. □

At the end of its life the firm's flow profits $x_t - k$ are negative but it remains in the market for the option value of a last-minute breakthrough that boosts its reputation and self-esteem to one. Over time, losses grow and the option value diminishes. The firm exits when they exactly offset each other.

The exit condition (3.7) implies that $V(t,z)$ is strictly convex in z on $\{(t,z) : V(t,z) > 0\}$. This follows because, in the argument in footnote 3, a high/low self-esteem firm mimicking a firm with intermediate self-esteem can strictly increase its profits by exiting later/earlier. Strict convexity and the investment condition (3.1) imply the existence of a threshold $z(t)$ with the property that the firm invests when $z_t^* > z(t)$, and disinvests when $z_t^* < z(t)$. Strict convexity also implies that best responses are strictly ordered in the following sense: Let (a^+, τ^+) and (a^-, τ^-) be optimal strategies and z^+, z^- the associated trajectories and assume that $z_t^+ > z_t^-$ for some $t < \min\{\tau^+, \tau^-\}$. Then $a_t^+ \geq a_t^-$ for almost all $s > t$, $z_s^+ > z_s^-$ for all $s > t$, and $\tau^+ > \tau^-$.

4 Equilibrium Analysis

So far we have studied the firm's optimal strategy for arbitrary beliefs $F = F(\tilde{a}, \tilde{\tau})$ and associated revenue trajectories $\{x_t\}$. In this section, we close the model by assuming that market beliefs are correct. To analyze the properties of correct beliefs we derive a more explicit expression for reputation, x_t. Breakthroughs arrive with intensity $\mu z_s(\tilde{a})$, so the probability of no breakthrough before time t equals $w_t(\tilde{a}) := \exp(-\mu \int_0^t z_s(\tilde{a}) ds)$. Bayes' rule then implies

$$x_t = \frac{\mathbb{E}^F \left[z_t(\tilde{a}) w_t(\tilde{a}) \mathbb{I}_{\{t<\tilde{\tau}\}} \right]}{\mathbb{E}^F \left[w_t(\tilde{a}) \mathbb{I}_{\{t<\tilde{\tau}\}} \right]}, \qquad t < \tau(F). \tag{4.1}$$

Definition: An *equilibrium* consists of a distribution over recursive investment and exit strategies $F = F(a, \tau) \in \Delta(B \times [0, \bar{\tau}])$ and a recursive revenue trajectory $x = \{x_t\}_{t \in [0,\bar{\tau}]} \in B$ such that:
(a) Given $\{x_t\}$, any strategy (a, τ) in the support of F solves the firm's problem (2.4).
(b) Reputation $\{x_t\}$ is derived from F by Bayes' rule via (4.1) for $t < \tau(F)$.

Note that this definition does not impose sequential optimality of strategy (a, τ) and thus corresponds to Nash equilibrium rather than sequential equilibrium. However, this is merely for notational convenience: as the firm's investment is unobservable, deviations do not affect beliefs and revenue. Thus, any equilibrium is outcome-equivalent to a sequential equilibrium. In fact, all of the analysis in the last section starting at states t, z_t^* extends immediately- to optimal strategies starting at any state t, z.

Theorem 4 (Existence) *An equilibrium exists.*

Proof. See Appendix A.4. □

The proof of Theorem 4 applies the Kakutani-Fan-Glicksberg Theorem to the best-response correspondence, mapping revenue $\{x_t\}$ to optimal strategies (a^*, τ^*), and the Bayesian updating correspondence (4.1), mapping mixed strategies $F(\tilde{a}, \tilde{\tau})$ to revenue $\{x_t\}$.[8] The key step in the proof is to define the appropriate weak topology that renders the strategy space compact (as in Lemma 1)

[8]The latter correspondence is multi-valued because x_t can take any value in $[z_t(0), z_t(\bar{a})]$ for $t \geq \tau(F)$.

and the two correspondences continuous.[9]

4.1 Pure-Strategy Equilibria

In a pure-strategy equilibrium the market is certain of the firm's strategy (a, τ) and so $x_t = z_t^*$ for all $t < \tau^*$. Thus, $\{x_t\}$ decreases and Theorem 2 implies that investment incentives $\Gamma(t)$ are single-peaked.

Equilibrium investment behavior depends on the investment cost c. If c is high, the firm always shirks. If c is intermediate, initial incentives $\lambda\Gamma(0)$ are insufficient to motivate effort, and any equilibrium is shirk-work-shirk. After a breakthrough, such a firm rests on its laurels because it has little to gain from an additional breakthrough; as its reputation and self-esteem drop, it starts investing and works hard for its survival, but eventually gives up and shirks before exiting the market. Finally, if c is small, then any equilibrium is work-shirk. This has the flavor of a probationary equilibrium where the market assumes a firm invests for a fixed period of time after each breakthrough, but then grows suspicious.[10]

The incentives at high reputations also depend critically on the level of $\overline{a}$. As $\overline{a} \to 1$, investment at $t \approx 0$ is impossible to sustain in equilibrium because with such market beliefs, reputation would remain close to 1 and dividends would remain small forever, undermining investment incentives. This same force is seen in Mailath and Samuelson (2001).

Figures 2-4 illustrate a pure-strategy equilibrium. This simulation considers a restaurant that has revenues of \$$x$ million a year, capital cost of $k = \$500,000$, investment cost of $c = \$125,000$ and an interest rate of $r = 20\%$ (incorporating a risk premium). Good news arrives when the restaurant is written up in the local paper; we set $\mu = 1$, so that on average a good restaurant is reviewed positively once a year. Finally, we set $\lambda = 0.2$, so that on average a technology shock arrives every 5 years. In these figures we replace the firm's state variable t with its time-t reputation x_t to aid comparison with the models in Section 5.[11]

This pure-strategy equilibrium is shirk-work-shirk, exhibiting work on $x \in [0.39, 0.94]$ with an exit threshold of $x^e = 0.22$. In Figure 2, the left panel shows the work region and the value function. The value function exhibits kinks at the edges of the work region, but smooth pasting at x^e. The right panel shows the distribution of surviving firms' reputation after 10 years, if all firms start at $x = z = 1$. This shows how firms tend to bunch in the work region, where drift is relatively slow. In

[9] We have not proved the existence of a *pure strategy* equilibrium. Given Theorem 2, we know that for small costs any equilibrium must be work-shirk. Hence one natural approach is to pick a cutoff t for believed investment, map it into the optimal cutoff $t^*(t)$ for actual investment, and find a fixed point. The problem with this approach is that this mapping may be discontinuous because strategies are dynamic complements.

[10] Formally, if $\overline{a}$ is large and a pure strategy equilibrium exists for all parameter values of $c > 0$, then there exist parameter values $c < c' < c''$ such that equilibrium investment a^* must be work-shirk for low costs c, shirk-work-shirk for intermediate costs c', and full-shirk for high costs c''.

[11] To solve for equilibirum numerically, we fix a candidate investment strategy $\{a_t\}$, calculate the resulting payoffs, verify that investment a_t satisfies the first-order condition (3.1), and finally determine the exit time τ by setting the firm's continuation value equal to zero. This approach is valid according to standard verification arguments (e.g. Davis (1993, Theorem 45.16)). We calculate one equilibrium but, since we have not shown uniqueness, there may be others.

Figure 3, the left panel shows the investment incentives $V_z(x, z)$ on the entire state-space (x, z). One can see that V_z increases in z, illustrating the convexity of the value function; it is also single-peaked in x, illustrating how incentives are low when reputation is high and when near the exit point. The right panel shows the incentives along the equilibrium path where $x = z$; this coincides with the 45° line on the horizontal plane of the left panel. These are clearly single-peaked, as shown in Theorem 2. Finally, Figure 4 shows three typical 10-year lifecycles for firms starting at $x = z = 1$. The left and center firms survive the 10-year period, experiencing 11 and 10 breakthroughs respectively; the right firm exits after 6 years after only a single breakthrough.

For empirical implementation, one may be concerned about the realism of having a large number of firms close to the upper end of the distribution, as in the right panel of Figure 2. With restaurants, one can view this as representing the degree of occupancy, say, on a weekend evening. If one is interested in explaining the size distribution of firms, there are a number of natural model extensions. First, as the mass of firms at the top results from our stylized breakthrough learning process, one could consider other learning processes as discussed at the end of Section 2. Second, one could relax the linear relationship between revenue and reputation. In competitive industries this relationship is convex, and so revenue has a thinner right tail than reputation. Third, one could endow the firm with regular capital in addition to reputational capital, so a restaurant with great reviews would need time to expand the franchise.

4.2 Mixed-Strategy Equilibria

The model may potentially exhibit mixed-strategy equilibria, as illustrated in Figure 5. In this picture, the work region is the area above some function $z(t)$. There are then two optimal trajectories of self-esteem for the firm: the 'low' path is full-shirk, while the 'high' path is shirk-work-shirk. The dynamic complementarity of investment means that the firm that works when the paths divide then *strictly* prefers to continue working, while the firm that shirks at the dividing line then *strictly* prefers to continue shirking. The firm's reputation, which equals the market's belief about the firm's self-esteem, is sandwiched between these two self-esteem paths.

When both firms remain in the market, reputation and self-esteem decline as normal. After sufficient time without a breakthrough, the 'low' firm will wish to exit. If the firm was believed to exit deterministically at some time τ, reputation would jump up at τ, undermining incentives so actually exit. Thus, equilibrium requires that the 'low' firm randomizes between exiting and remaining in the market over some period $[\underline{\tau}, \overline{\tau}]$. During this exit period, self-esteem declines, and so reputation has to rise so as to satisfy the firm's indifference condition (3.7). Once the 'low' firm has exited with probability one, reputation coincides with the self-esteem of the 'high' firm. Such a firm exits deterministically when its self-esteem has fallen sufficiently. Since reputation increases at some times prior to exit, we can no longer conclude that investment incentives are single-peaked. However, Theorem 1 holds, and the 'high' firm shirks near the exit time.[12]

[12] A mixed strategy equilibrium has the interesting technical property that smooth-pasting fails at the exit time. In particular, for the 'low' firm, the investment incentives depend on the choice of the exit time. For example, $\Gamma(\underline{\tau}) = 0$ if $\tau^* = \underline{\tau}$ but $\Gamma(\underline{\tau}) > 0$ if $\tau^* > \underline{\tau}$. As $V_z(\underline{\tau}, z_{\underline{\tau}}-) \leq \Gamma(\underline{\tau}) \leq V_z(\underline{\tau}, z_{\underline{\tau}}+)$ for any exit time, firm value $V(t, z)$ cannot be

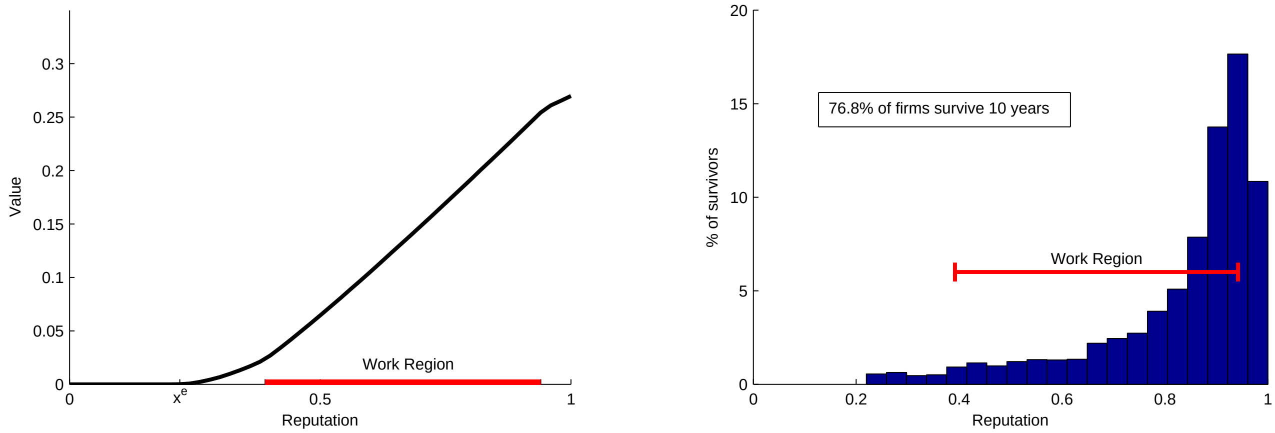

Figure 2: **Value Function and Distribution of Firms.** The **left** panel shows the firm's value as a function of its reputation. The **right** panel shows the resulting distribution of firms after 10 years, assuming all firms start at $x = z = 1$. This figure assumes capital cost $k = 0.5$, interest rate $r = 0.2$, maximum effort $\overline{a} = 0.9$, investment cost $c = 0.125$, technology shock rate $\lambda = 0.2$, and breakthrough rate $\mu = 1$.

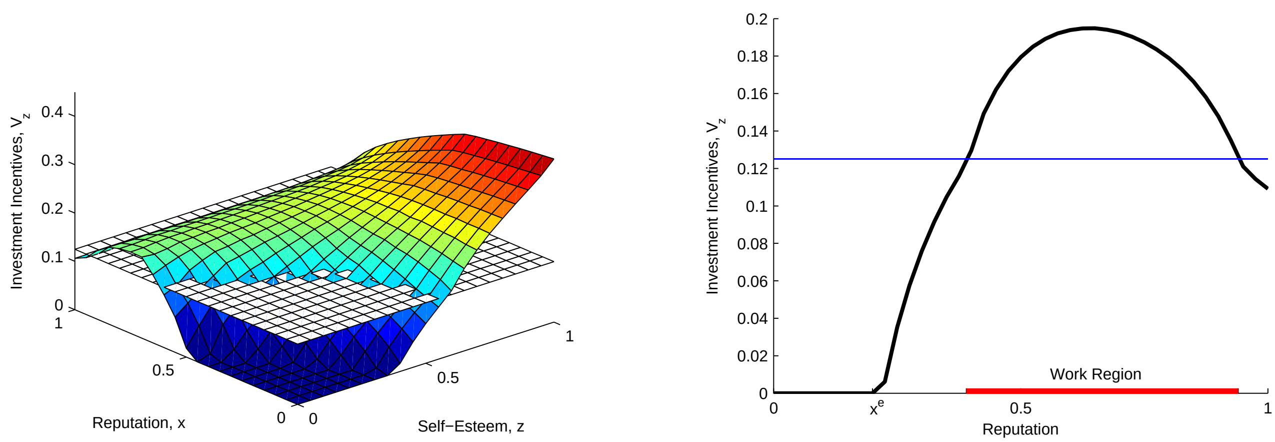

Figure 3: **Investment Incentives.** The **left** panel shows the firm's investment incentives $V_z(x_t, z_t)$ as a function of reputation x_t and self-esteem z_t. The **right** panel shows the investment incentives along the equilibrium path, where $x_t = z_t$. The parameters are the same as in Figure 2.

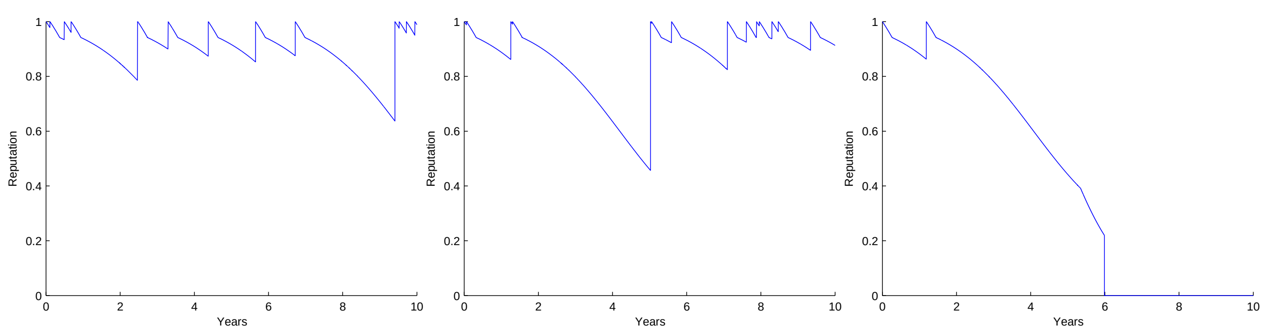

Figure 4: **Firm Lifecycles.** This figure shows three sample paths starting at $x = z = 1$. The **left** and **center** firms survive 10 years, the **right** firm exits after 6 years. The parameters are the same as in Figure 2.

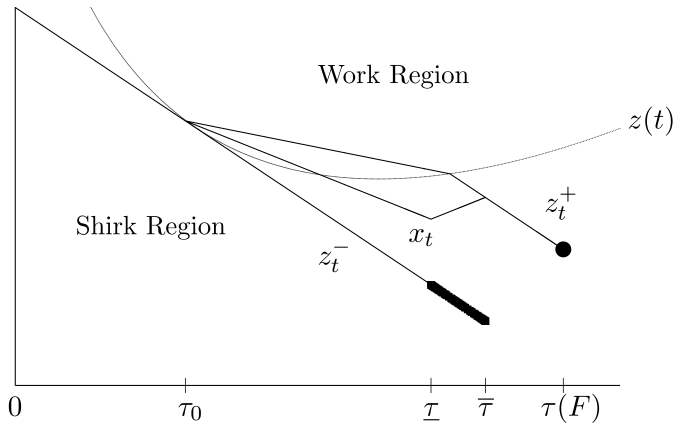

Figure 5: **Mixed Strategy Equilibrium.** This picture illustrates the qualitative features of a mixed strategy equilibrium. At time τ_0 the firm mixes between working and shirking resulting in self-esteem paths $\{z_t^+, z_t^-\}$ and reputation x_t. The z_t^- firms exit over $[\underline{\tau}, \overline{\tau}]$, while the z_t^+ firms exit at $\tau(F)$.

More generally, a mixed-strategy equilibrium could have more than two optimal investment plans a^*, but the qualitative features of the above example, such as gradual exit during which reputation recovers, carry over from the above discussion.

5 Model Variants

In this section, we consider two natural variants of our baseline model. In Section 5.1 we assume the market observes the firm's investment; this maintains the assumption that both the firm and the market learn about the firm's quality. In Section 5.2 we analyze a model in which the firm knows its own quality, while the market learns it via breakthroughs as before; this maintains the moral hazard assumption. Additionally, Section 5.3 considers the model in between the baseline model and the version with known quality, assuming the firm does not know its quality, but instead observes private signals in addition to the public signals.

5.1 Observable Investment

In the baseline model, the firm's investment was unobserved by the market; in this section we eliminate moral hazard, and assume that the market can directly observe the firm's investment. Theorem 5 shows that investment incentives decrease over time. Comparing this model without moral hazard to our baseline model, Theorem 6 shows that moral hazard decreases investment.

Suppose that the market observes the history of signals h^t and the firm's past investment $\{a_s\}_{s\leq t}$. Since the market has the same information as the firm, reputation and self-esteem coincide $x_t = z_t$; we can thus write firm value as a function of self-esteem alone. Analogous to (2.4), we truncate the

differentiable at $(t, z) = (\underline{\tau}, z_{\underline{\tau}})$. This failure of smooth-pasting is analogous to that in Keller and Rady (2014).

integral at a breakthrough, yielding

$$\hat{V}(z_t) = \sup_{a,\tau} \int_t^\tau e^{-\int_t^s r+\mu z_u du} \left[z_s - a_s c - k + \mu z_s \hat{V}(1)\right] ds. \tag{5.1}$$

Denote $(\hat{a}, \hat{\tau})$ as the optimal strategy and $\hat{z} = \{\hat{z}_t\}$ the associated self-esteem. Since the firm controls both self-esteem and reputation, the analysis reduces to a decision problem (rather than finding an equilibrium). Existence of an optimal strategy then follows as in Lemma 1.

The rest of the analysis follows the pattern of Section 3. First, the value function $\hat{V}(z)$ is strictly convex. The firm's expected profit is linear in its belief, so optimizing over its investment and exit strategies yields a convex function. Economically, a firm that invests today expects to live longer and so gains more from a marginal gain in its self-esteem. Next, by (2.3), investment raises self-esteem (and reputation) at rate λ. Optimal investment is thus characterized by the derivative of $\hat{V}$, which is given by

$$\hat{\Gamma}(t) := \int_t^{\hat{\tau}} e^{-\int_t^s r+\lambda+\mu(1-\hat{z}_u)du} \left[1 + \mu(\hat{V}(1) - \hat{V}(\hat{z}_s))\right] ds, \tag{5.2}$$

whenever it exists.[13] More generally, $\hat{\Gamma}(t)$ is sandwiched between the upper and lower derivatives of $\hat{V}(\hat{z}_t)$, which exist by convexity. When compared to the investment incentives with moral hazard (3.2), investment now affects reputation and revenue directly. This accounts for the "1" term in the integrand. In any optimal strategy, investment therefore satisfies

$$\hat{a}_t = \begin{cases} 0 & \text{if } \lambda\hat{\Gamma}(t) < c \\ 1 & \text{if } \lambda\hat{\Gamma}(t) > c. \end{cases}$$

Finally, the optimal exit time $\hat{\tau}$ satisfies $\hat{z}_{\hat{\tau}} - k + \mu \hat{z}_{\hat{\tau}} \hat{V}(1) = 0$.

Theorem 5 (Observable Investment Characterization) *Assume that the market observes the firm's investment. Then investment incentives $\hat{\Gamma}(t)$ decrease in t, with $\hat{\Gamma}(\hat{\tau}) = 0$.*

Proof. Given assumption (2.2), drift $g(a_t, z_t)$ is boundedly negative on $[z^\dagger, 1]$ and the firm exits before its reputation hits $z^\dagger$. Since z_t decreases and the value function is strictly convex, $\hat{V}'(z_t)$ strictly decreases in t (whenever this derivative exists). Then, $\hat{\Gamma}(t)$ also strictly decreases because it is continuous and coincides with $\hat{V}'(z_t)$ whenever this derivative exists. □

By Theorem 5, optimal investment is either full-shirk or work-shirk, depending on the cost of investment. That is, there exists $\hat{t} \in [0, \hat{\tau})$ such that $a_t = 1$ for almost all $t < \hat{t}$ and $a_t = 0$ for almost all $t > \hat{t}$. As time progresses without a breakthrough, the firm gets closer to its exit point and any investment pays off over a shorter horizon, reducing incentives.

When comparing this result to the baseline model, note that with moral hazard the firm benefits from its investment solely through the reputational dividends obtained when the market learns about the quality via a breakthrough. Immediately- following a breakthrough, these dividends are zero and

[13] See Appendix A.5 for a proof.

so investment incentives (3.2) must increase. In contrast, with observable investment the dividends are initially 1, and as time progresses the loss of immediate dividends outweighs the benefit of getting future dividends sooner. Hence the investment incentives (5.2) decrease over time.

We next wish to show that investment incentives with observable investment are larger than those in the baseline model. Intuitively, the additional 1 term indicates that investment incentives are higher; however many other terms in the integral are endogenous. To state the result formally, for any equilibrium of the baseline model, define $t^* := \sup\{t : \Pr(\lambda\Gamma(t) \geq c) > 0\}$ as the last time at which investment is optimal.

Theorem 6 (Impact of Moral Hazard) *With observable investment, the firm works strictly longer than in any baseline equilibrium, $t^* < \hat{t}$.*[14]

Proof. See Appendix A.6. □

Intuitively, in the baseline case, the firm raises its self-esteem when it invests. With observable investment, the firm raises both its self-esteem and its reputation. Since a higher reputation is good for the firm, there is a higher marginal benefit to investment and, in equilibrium, the firm invests more.

Figure 6 simulates the firm's optimal strategy in this model variant for the same parameters as in Figures 1-4. In this example, the firm works for $z \in [0.24, 1]$ with an exit cutoff $\hat{z}^e = 0.19$. The left panel shows that the additional investment at high reputations is very valuable for the firm, raising the value of a firm with perfect reputation from 0.27 to 0.34. Intuitively, the investment at the top means that reputation initially falls slowly, and a breakthrough is likely to occur before the reputation has fallen significantly (see Figure 1). As a result, most firms remain close to $z = 1$ after 10 years, with 96.8% surviving, as shown in the right panel.

5.2 Privately-Known Quality

In the baseline model, we assume that the firm learns about its quality through the same public breakthroughs as the market and acquires private information only if the market's beliefs about investment are mixed or incorrect. In this section we suppose the firm knows its own quality, but the market still does not. This model is reasonable if the firm has much better information than the market; for example, a restaurant owner may receive direct, non-public feedback from his patrons. This introduces the possibility that the firm signals its quality simply by remaining in business.

As before, we focus on strategies and beliefs that are recursive in the time since the last breakthrough, t. Additionally, we assume that the firm conditions its strategy on current quality but not past quality, which is payoff-irrelevant. Thus, a recursive strategy consists of an investment plan $a = \{a_t\}$ and an exit time τ for the firm with either quality level.

[14] Theorem 6 holds for both pure and mixed equilibria of the baseline model. In a mixed-strategy equilibrium, firm investment $\{a_t\}$ and thus investment incentives $\Gamma(t)$ are random variables. In a pure-strategy equilibrium one can also state the result in terms of reputation space. That is, if the firm with moral hazard works at some x, then the firm with observable investment also works at this x.

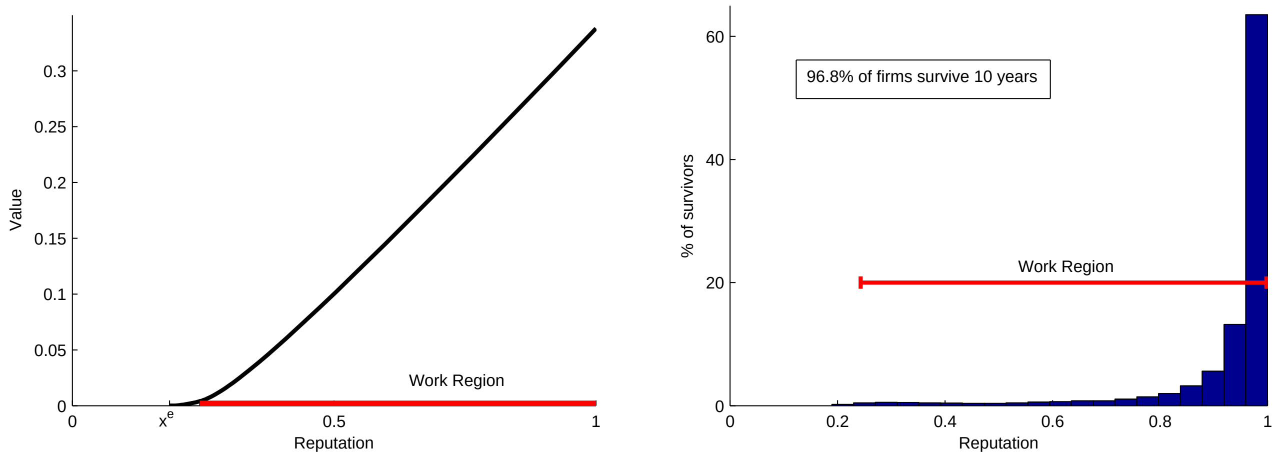

Figure 6: **Market Observes Firm's Investment.** The **left panel** shows the firm's value as a function of its reputation. The **right panel** shows the resulting distribution of firms after 10 years, assuming all firms start at $x = z = 1$. The parameters are the same as in Figure 2.

To analyze the value of a low-quality firm at time t, we truncate its cash flow expansion at the first technology shock, obtaining

$$V(t,0) = \sup_{a,\tau} \int_t^\tau e^{-(r+\lambda)(s-t)} \left[x_s - ca_s - k + \lambda(a_s V(s,1) + (1-a_s)V(s,0)) \right] ds. \tag{5.3}$$

Compared to a low-quality firm, a high-quality firm additionally enjoys breakthroughs with present value $V(0,1) - V(s,1)$ at arrival rate μ, giving

$$V(t,1) = \sup_{a,\tau} \int_t^\tau e^{-(r+\lambda)(s-t)} \left[x_s - ca_s - k + \lambda(a_s V(s,1) + (1-a_s)V(s,0)) + \mu(V(0,1) - V(s,1)) \right] ds. \tag{5.4}$$

Investment raises the quality of the firm at rate λ. Writing $\Delta(s) = V(s,1) - V(s,0)$ for the *value of quality*, optimal investment is thus characterized by the bang-bang condition

$$a_s = \begin{cases} 0 & \text{if } \lambda\Delta(s) < c, \\ 1 & \text{if } \lambda\Delta(s) > c. \end{cases} \tag{5.5}$$

Importantly, optimal investment is independent of the firm's quality, allowing us to write it as a deterministic function of the time since the last breakthrough $a^* = \{a_t^*\}$. Intuitively, investment only pays off if there is a technology shock, in which case the firm's current quality is irrelevant. For the firm's exit decision, write τ^θ for the optimal exit time(s) of a firm with current quality θ.

An equilibrium in this model variant consists of a distribution over strategies, $F^\theta(a,\tau)$ for $\theta \in \{L, H\}$, and a reputation trajectory $\{x_t\}$ such that: (1) all equilibrium strategies are optimal, and (2) reputation x_t is derived from the distributions via Bayes' rule whenever possible.

We restrict attention to equilibria where reputation $\{x_t\}$ is continuous; this assumption holds on the equilibrium path, but off the equilibrium path it implies the firm cannot be punished for failing to exit by a downward jump in beliefs. We ignore such equilibria because it is implausible for the market to interpret failure to exit as a signal of low quality. We also restrict attention to equilibria with weakly-decreasing reputation; this is satisfied in any markovian equilibrium.

The following result is a version of Bar-Isaac (2003, Proposition 2).

Lemma 5 (Exit) *In an equilibrium with continuous, weakly-decreasing reputation $\{x_t\}$, there exists a time $\underline{\tau} < \infty$ such that the exit time of the low-quality firm τ^L has support $[\underline{\tau}, \infty)$, and revenue and firm value are constant for $t \in [\underline{\tau}, \infty)$ and satisfy*

$$x_t - k + \overline{a} \max\{\lambda V(t,1) - c, 0\} = 0. \tag{5.6}$$

The high-quality firm never exits, i.e. $\tau^H = \infty$.

Proof. Since reputation $\{x_t\}$ continuously decreases, firm value $V(t,\theta)$ continuously decreases in t. For an optimal investment strategy a^*, the flow payoff of the low- and high-quality firm $-$ the integrands in (5.3) and (5.4) $-$ continuously decreases in t as well. Thus, exiting is optimal exactly when flow-payoffs are zero; since firm value is zero at an exit-time, flow payoffs of the low-quality firm are given by (5.6). As flow payoffs of the high-quality firm exceed those of the low-quality firm by the last, positive term in (5.4), the latest possible exit time of the low-quality firm must strictly precede the earliest possible exit time of the high-quality firm.

To see that the low-quality firm starts exiting at some finite $\underline{\tau}$, note first that (2.2) implies the low-quality firm exits with certainty before its reputation falls to $x^\dagger := \lambda/\mu$; at this reputation the negative flow profits $x^\dagger - k$ from staying in the market exceed the option value of staying in the market $\lambda V(t,1)$, which is bounded above by $\lambda(1-k)/r$. Moreover, unless the market expects the low-quality firm to start exiting and draws a positive inference from its failure to exit, reputational drift $g(\tilde{a}, x)$ is strictly negative on $[x^\dagger, 1]$ and takes reputation below $x^\dagger$ in finite time. Thus, in equilibrium the low-quality firm must eventually exit, and we define $\underline{\tau}$ as the earliest time at which it does so.

After $\underline{\tau}$ reputation must be constant. Otherwise, if it started to decrease at some time $t > \underline{\tau}$, the flow payoffs of the low-quality firm turn strictly negative and the low-quality firm would exit with certainty; thus reputation would jump to one, undermining incentives to exit. Therefore, the firm's problem becomes stationary after $\underline{\tau}$, all exit times $\tau^L \in [\underline{\tau}, \infty)$ are optimal and the high-quality firm never exits. Finally, in order to keep reputation constant at $x_{\underline{\tau}}$, the low-quality firm must exit at constant rate $-g(a_{\underline{\tau}}, x_{\underline{\tau}})/x_{\underline{\tau}}(1 - x_{\underline{\tau}})$ to offset the negative reputational drift due to learning. $\square$

Given the low-quality firm's indifference at the exit threshold, we can assume that it always remains in the market, and therefore follows the same strategy as the high-quality firm. Subtracting (5.3) from (5.4), we obtain the following expression for the equilibrium value of quality

$$\Delta(t) = \int_t^\infty e^{-(r+\lambda)(s-t)} \mu(V(0,1) - V(s,1)) ds. \tag{5.7}$$

The integrand in (5.7) represents the *reputational dividend of quality*: high quality does not affect the firm's reputation and revenue immediately but gives rise to future breakthroughs that arrive at rate μ and boost the firm's reputation to one. These dividends depreciate at both the time-discount rate r and the quality obsolescence rate λ. This is analogous to the investment incentives (3.2) in the baseline model.

Theorem 7 (Known Quality Characterization) *In an equilibrium with continuous, weakly decreasing reputation $\{x_t\}$, investment incentives $\Delta(t)$ increase over time.*

Proof. As s rises, the firm's value $V(s,1)$ falls and reputational dividends $V(0,1) - V(s,1)$ grow. Hence an increase in t leads to an increase in the value of quality (5.7), and in investment via the optimality equation (5.5). $\square$

Intuitively, breakthroughs are most valuable to a firm with low reputation since a breakthrough takes the firm from its current reputation to $x = 1$. Thus, the optimal investment strategy is either full shirk, shirk-work or full work, depending on the investment cost c.

In contrast to the baseline model, all states $(t, \theta) \in [0, \infty) \times \{0, 1\}$ are on-path in this model variant, and so conditions (5.5) and (5.6) are sufficient as well as necessary for an equilibrium strategy. Thus, disinvestment at times $t \in [0, \underline{t})$, investment thereafter and exit of low-quality firms after time $\underline{\tau}$ constitutes an equilibrium if (5.5) and (5.6) are satisfied. Equilibrium existence can thus be established by Brouwer's fixed-point theorem applied to $(\underline{t}, \underline{\tau}) \in [0, \underline{\tau}] \times [0, \overline{\tau}]$, as in Board and Meyer-ter-Vehn (2013, Theorem 2).

The increasing investment incentives in Theorem 7 are in sharp contrast to the single-peaked, eventually-vanishing investment incentives in Theorem 2. In the baseline model, the firm gives up near the exit threshold and coasts into liquidation; with privately-known quality, the firm fights until the bitter end. As discussed above, with unknown quality, the firm's investment at times $t \in [\tau^* - \varepsilon, \tau^*]$ pays off only if a technology shock arrives and a breakthrough arrives that averts exit. The probability of this joint event is of order ε^2, hence the expected gain eventually falls short of the investment costs. With known quality, only a technology shock is required for investment to pay off, because a boost in quality is immediately witnessed by the firm, which then averts exit. Thus, investment incentives are of order ε at all times, and are actually maximized when the firm is about to exit by Theorem 7.

With observed investments, incentives decrease as the firm gets closer to exit (Theorem 5). With known quality, incentives increase as the breakthroughs become more valuable to the firm (Theorem 7). The single-peaked incentives in the baseline model (Theorem 2) can be viewed as a combination of these two effects.

Figures 7-8 illustrate an equilibrium for the same parameters used in Figures 1-4. In Figure 7, the left panel shows the value function of high- and low-quality firms with different reputations. One can see the kink at the start of the work region. In this example, the value of a firm with perfect reputation is 0.22, compared with 0.27 under unknown quality. The right panel plots the distribution of firms after 10 years, starting at $x = 1$. There is a bulge of firms in the work region,

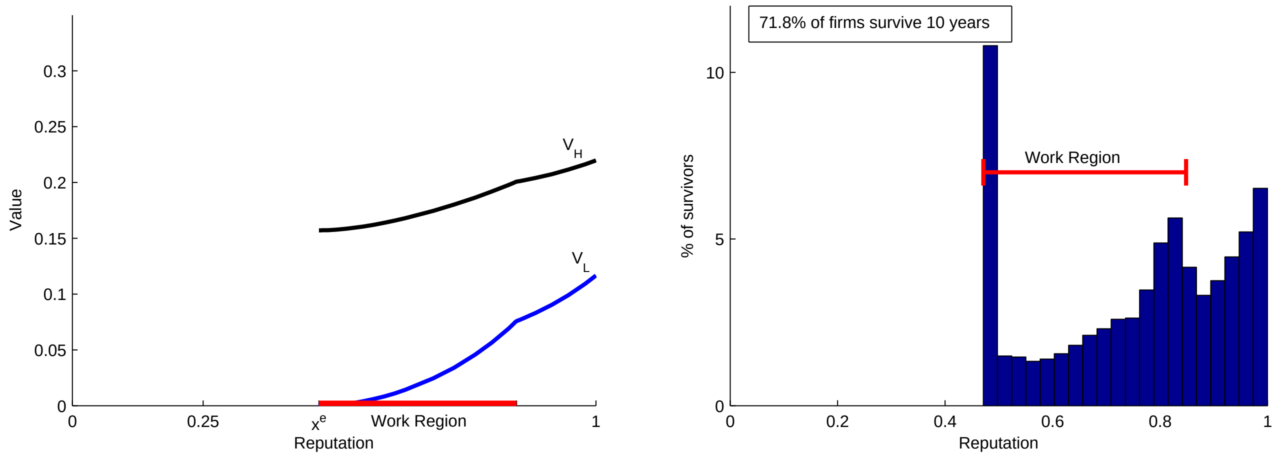

Figure 7: **Privately Known Quality.** The **left** panel shows the firm's value as a function of its reputation. The **right** panel shows the resulting distribution of firms after 10 years, assuming all firms start at $x = z = 1$. The parameters are the same as Figure 2.

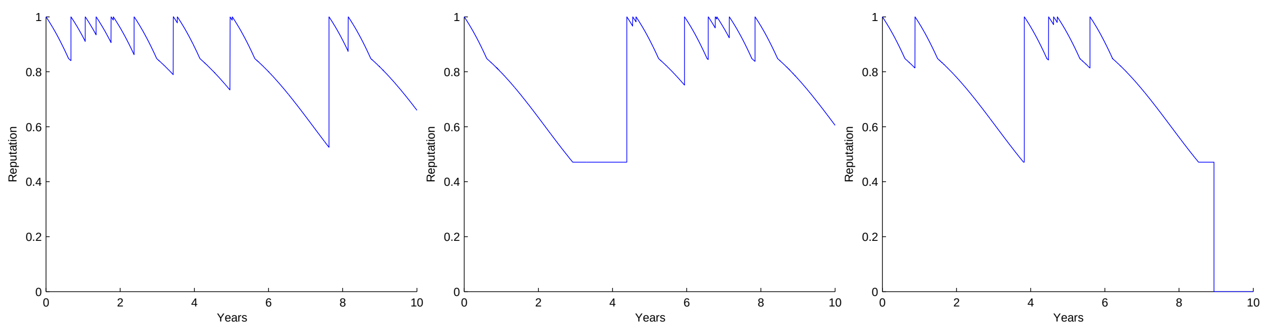

Figure 8: **Firm Lifecycles with Privately-Known Quality** This figure shows three sample paths starting at $x = z = 1$. The **left** and **center** firms survive 10 years, although the center firm hits the exit point and survives the randomization. The **right** firm exits after 9 years. The parameters are the same as in Figure 2.

where the firm's reputation begins to decline more slowly; relative to the baseline distribution, this distribution appears censored. The exit threshold is much higher, with a large number of firms massing at this threshold, where low-quality firms randomize between exiting and remaining in the market. Overall 71.8% firms survive 10 years, compared to 76.8% in the baseline case. Figure 8 shows three typical firms' life-cycles. The first firm never comes near the exit threshold; the second is temporarily indifferent between exiting and not, and ultimately survives; the third exits after 9 years.

5.3 Imperfect Private Information

The model variant with privately-known quality predicts investment behavior opposite that of the baseline model when the firm is about to exit. To further illuminate this contrast, we now bridge the two extremes by nesting them in a class of models with imperfect private information.

We model the firm's new information via private breakthroughs that arrive at rate ν and reveal high quality with certainty. Thus, at a public breakthrough reputation and self-esteem jump to one; at a private breakthrough reputation is continuous while self-esteem jumps to one; absent either breakthrough, self-esteem is governed by $\dot{z} = g(a_t, z_t)$ with

$$g(a_t, z_t) = \lambda(a_t - z_t) - (\mu + \nu)z_t(1 - z_t).$$

When $\nu = 0$ we recover the unknown quality case. As $\nu \to \infty$, this model approximates the known quality case in the sense that self-esteem z_t converges to 0 or 1 in distribution for any time $t > 0$ and any investment strategy a.[15]

This model is recursive in the time since the last public breakthrough, t. A recursive strategy for the firm then specifies investment a_t and exit time τ as a function of the history of private breakthroughs. Writing optimal strategies as (a, τ) and the resulting process of self-esteem as z^*, we truncate the firm's cash flow expansion at either kind of breakthrough to obtain

$$V(t, z_t^*) = \int_t^{\tau^*} e^{-\int_t^s (r+(\mu+\nu)z_u^*)du} \left[x_s - ca_s^* - k + \mu z_s^* V(0,1) + \nu z_s^* V(s,1)\right] ds.$$

The additional term $\nu z_s^* V(s,1)$ captures the firm's continuation value after a private breakthrough. As in Lemmas 2-3, investment incentives are given by

$$\Gamma(t) = \int_t^{\tau^*} e^{-\int_t^s (r+\lambda+(\mu+\nu)(1-z_u^*))du} \left[\mu(V(0,1) - V(s, z_s^*)) + \nu(V(s,1) - V(s, z_s^*))\right] ds. \tag{5.8}$$

These disappear at the exit time τ^* as in Theorem 1, so the firm will shirk close to the exit threshold. Thus, even as $\nu \to \infty$ and the model approaches the known-quality case, the firm shirks before exiting, in contrast to Theorem 7. However, this does not imply a discontinuity: The integrand in (5.8) increases in ν, so while the firm shirks for some time before exit in equilibrium, the length of this shirking time may converge to zero as ν grows large.

6 Conclusion

This paper models the lifecycle of a firm whose primary assets are its quality and its reputation. In the baseline model, the firm privately invests in its quality, while both the market and firm learn about the success of past investments. We characterize investment incentives and show they are single-peaked in the time since a breakthrough. This yields predictions about the distribution of firm revenue product and the industry turnover rate. Finally, we investigate two variants of our model: one where investment is publicly observed, and a second where the firm privately knows its quality.

[15]To see this, note that for any $\varepsilon > 0$ there exists $\delta > 0$ and $\nu^* > 0$ such that for all $\nu > \nu^*$ we have either $z_{t+\delta} < \varepsilon$ if no breakthrough arrived in $[t, t+\delta]$, or $\Pr(z_{t+\delta} > 1 - \varepsilon) > 1 - \varepsilon$ if a breakthrough arrived at $t' \in [t, t+\delta]$ and thus $\Pr(\theta_{t+\delta} = H | \theta_{t'} = H) \geq e^{-\lambda\delta}$.

We believe this model has a wide variety of applications that lend themselves to empirical investigation For example, a film studio invests in its personnel and its products; the studio and Hollywood then learn about its production quality via its hit movies. In academic labor markets, researchers invest in their skills; the profession then learns about the success through publications. At the international level, countries makes policy choices concerning government spending and privatizations; both the country and their sovereign debt holders then learn about the country's solvency via public statistics.

A Appendix

A.1 Proof of Lemma 2

Fix time t, self-esteem z_t, firm strategy (a, τ) (not necessarily optimal), write $z = \{z_s\}_{s \geq t}$ for future self-esteem, and let

$$\Pi(t, z_t) = \int_{s=t}^{\tau} e^{-\int_t^s r+\mu z_u du}(x_s - ca_s - k + \mu z_s \Pi(0,1))ds \tag{A.1}$$

be the firm's continuation value, where the integral of the cash-flows is truncated at the first breakthrough as in (2.4). We will show that $\Pi(t, z)$ is differentiable in z with derivative

$$\Pi_z(t, z_t) = \int_{s=t}^{\tau} e^{-\int_t^s r+\lambda+\mu(1-z_u)du}\mu(\Pi(0,1) - \Pi(s, z_s))ds \tag{A.2}$$

Equation (3.2) then follows by the envelope theorem, Milgrom and Segal (2002).

To show (A.2) we first state two claims, both of which follow immediately from Board and Meyer-ter-Vehn (2013).

Claim 1: For any bounded, measurable functions $\phi, \rho : [0, \tau] \to \mathbb{R}$, the function

$$\psi(t) = \int_t^{\tau} e^{-\int_t^s \rho(u)du}\phi(s)ds \tag{A.3}$$

is the unique solution to the integral equation

$$f(t) = \int_{s=t}^{\tau} (\phi(s) - \rho(s)f(s))ds. \tag{A.4}$$

This is proved for $\tau = \infty$ and constant ρ in Board and Meyer-ter-Vehn (2013, Lemma 5). The proof generalizes immediately to finite τ and measurable functions $\rho(t)$.

Claim 2: For any times $s > t$ and fixed investment a, time-s self-esteem z_s is differentiable in time-t self-esteem z_t. The derivative is

$$\frac{dz_s}{dz_t} = \exp\left(-\int_{u=t}^{s} (\lambda + \mu(1 - 2z_u))du\right).$$

This follows by the same arguments as in Board and Meyer-ter-Vehn (2013, Lemma 8B).

Setting $\psi(s) = e^{-r(s-t)}\Pi(s, z_s)$, $\rho(s) = \mu z_s$ and $\phi(s) = e^{-r(s-t)}(x_s - ca_s - k + \mu z_s \Pi(0,1))$, equation (A.1) becomes (A.3). Applying Claim 1, we get (A.4) which becomes

$$\Pi(t, z_t) = \int_{s=t}^{\tau} e^{-r(s-t)}(x_s - ca_s - k + \mu z_s(\Pi(0,1) - \Pi(s, z_s)))ds.$$

Taking the derivative with respect to z at $z = z_t$ and applying Claim 2, we get

$$\begin{aligned}\Pi_z(t, z_t) &= \int_{s=t}^{\tau} e^{-r(s-t)} \frac{dz_s}{dz_t}(\mu(\Pi(0,1) - \Pi(s, z_s)) - \mu z_s \Pi_z(s, z_s))ds \\ &= \int_{s=t}^{\tau} e^{-\int_t^s r+\lambda+\mu(1-2z_u)du}(\mu(\Pi(0,1) - \Pi(s, z_s)) - \mu z_s \Pi_z(s, z_s))ds.\end{aligned}$$

Setting $\rho(s) = \mu z_s$, $\phi(s) = e^{-\int_t^s r+\lambda+\mu(1-2z_u)du}\mu(\Pi(0,1)-\Pi(t, z_s))$, and $f(s) = e^{-\int_t^s r+\lambda+\mu(1-2z_u)du}\Pi_z(s, z_s)$, the previous equation becomes (A.4). Applying Claim 1, we get (A.3) which becomes

$$\Pi_z(t, z_t) = \int_t^{\tau} e^{-\int_t^s \mu z_u du} e^{-\int_t^s r+\lambda+\mu(1-2z_u)du}\mu(\Pi(0,1) - \Pi(t, z_s))ds,$$

implying (A.2).

A.2 Monotonicity of Value Function in Section 3

Lemma 6 *If $\{x_t\}$ strictly decreases, then $V(t, z)$ strictly decreases in t and strictly increases in z on $\{(t, z) : V(t, z) > 0\}$.*

Proof. Fix $t \geq t'$ and $z \leq z'$ and consider a 'low' firm with initial state (t, z) and a 'high' firm with initial state (t', z'). We can represent the firms' uncertainty as an increasing sequence of potential breakthrough times $\{t_i\}_{i\in\mathbb{N}}$ that follow a Poisson distribution with parameter λ, and a sequence of uniform $[0, 1]$ random variables $\{\zeta_i\}_{i\in\mathbb{N}}$, with the interpretation that the firm experiences an actual breakthrough after time σ (that is at time $t+\sigma$ for the 'low' firm and at time $t'+\sigma$ for the 'high' firm) if $\sigma = t_i$ for some i and $\zeta_i \leq Z_{t-}$. Fixing any realization of uncertainty $\{t_i, \zeta_i\}_{i\in\mathbb{N}}$, let $(\{A^*_\sigma\}, T^*)$ be the 'low' firm's optimal strategy given this realization, and assume that the 'high' firm mimics this strategy. Note that this strategy is in general not recursive for the 'high' firm. Given $\{t_i, \zeta_i\}$ and $(\{A^*_\sigma\}, T^*)$, we can compute revenue and self-esteem of the 'low' and 'high' firms (X_σ, Z_σ) and (X'_σ, Z'_σ), respectively, for any $\sigma \geq 0$. We now argue inductively that

$$X_\sigma \leq X'_\sigma \text{ and } Z_\sigma \leq Z'_\sigma \tag{A.5}$$

for any $\sigma < t_i$ and any $i \in \mathbb{N}$. For $i = 1$ ($\sigma \in [0, t_1)$) we have $X_\sigma = x_{t+\sigma} < x_{t'+\sigma} = X'_\sigma$ because $\{x_t\}$ decreases, and the self-esteem trajectories Z_σ, Z'_σ are governed by the ODE $\dot{z} = g(a, z)$, implying (A.5) for $\sigma \in [0, t_1)$. At $\sigma = t_1$, the 'low' (resp 'high') firm experiences a breakthrough if $\zeta_1 \leq Z_{\sigma-}$ (resp $\zeta_1 \leq Z'_{\sigma-}$). As $Z_{\sigma-} \leq Z'_{\sigma-}$, we get (A.5) for $\sigma = t_1$. Inductive application of these steps yields (A.5) for all σ. Thus by mimicking the 'low' firm's optimal strategy $(\{A^*_\sigma\}, T^*)$ for any realization $\{t_i, \zeta_i\}$, the 'high' firm can guarantee itself weakly higher cash-flows $X'_\sigma - cA^*_\sigma - k$ at all times σ, implying $V(t', z') \geq V(t, z)$. As long as firm value is strictly positive and the firms don't exit immediately, the inequality $X_\sigma \leq X'_\sigma$ is strict for a positive measure of times with positive probability, implying $V(t', z') > V(t, z)$. □

A.3 Proof of Theorem 2

The discount rate $\rho(t) = r + \lambda + \mu(1 - z_t^*)$ is Lipschitz continuous, with derivative $\mu \dot{z}_t^*$ where $\dot{z}_t^* = g(a_t^*, z_t^*) = \lambda(a_t^* - z_t^*) - \mu z_t^*(1 - z_t^*)$ for almost all t. Firm value as a function of time $t \mapsto V(t, z_t^*)$ is also Lipschitz continuous with derivative $\frac{d}{dt}V(t, z_t^*) = \Psi(t) + \dot{z}_t^*\Gamma(t)$ for almost all t.

Now assume that $\dot{\Gamma}(t) \leq 0$. Then

$$
\begin{aligned}
\dot{\Gamma}(t+\varepsilon) - \dot{\Gamma}(t) &= \int_t^{t+\varepsilon} \frac{d}{ds}\left[\rho(s)\Gamma(s) - \mu(V(0,1) - V(s, z_s^*))\right] ds \\
&= \int_t^{t+\varepsilon} \left[\rho(s)\dot{\Gamma}(s) + \dot{\rho}(s)\Gamma(s) + \mu \frac{d}{ds} V(s, z_s^*)\right] ds \\
&= \int_t^{t+\varepsilon} \left[\rho(s)\dot{\Gamma}(s) - \mu \dot{z}_s^*\Gamma(s) + \mu(\Psi(s) + \dot{z}_s^*\Gamma(s))\right] ds \\
&= \int_t^{t+\varepsilon} \left[\rho(s)\dot{\Gamma}(s) + \mu\Psi(s)\right] ds.
\end{aligned}
$$

As $\dot{\Gamma}(t) \leq 0$ and $\Psi(t) < 0$ and both of these functions and $\rho(t)$ are continuous, the integrand is strictly negative for small ε, so $\dot{\Gamma}$ strictly decreases on some small interval $[t, t+\varepsilon)$. If $\dot{\Gamma}$ did not strictly decrease on $[t, \tau^*]$ there would exist $t' > t$ with $\dot{\Gamma}(t') < 0$ and $\dot{\Gamma}(t'+\varepsilon) \geq \dot{\Gamma}(t')$ for arbitrarily small ε, which is impossible by the above argument.

A.4 Proof of Theorem 4

Proof strategy: The firm's payoff from strategy (a, τ) is given by

$$
\Pi(a, \tau; x) = \frac{\int_0^\tau e^{-\int_{s=0}^t r + \mu z_s ds} (x_t - c a_t - k)\, dt}{1 - \int_0^\tau e^{-\int_{s=0}^t r + \mu z_s ds} \mu z_t dt}
$$

The proof idea is to show that the firm's best-response correspondence

$$
BR(x) = \arg\max_{a,\tau} \Pi(a, \tau; x)
$$

and the Bayesian updating formula $\mathcal{B}$ defined by (4.1) admit a fixed point.

To establish existence of a fixed point we define topologies on the space of mixed strategies F and reputation trajectories $\{x_t\}$ with the property that both spaces are compact, locally convex, and Hausdorff, and both correspondences are upper-hemicontinuous. Then the existence of the fixed point follows by the Kakutani-Fan-Glicksberg theorem.

Defining the topologies: In the proof of Lemma 1 we interpreted investment strategies $\{a_t\}_{t \in [0, \overline{\tau}]}$ as elements of a space B, endowed with a weak topology under which B is compact. We now also interpret revenue trajectories $\{x_t\}$ as elements of B with this topology. As for the firm's mixed strategies, we equip $\Delta(B \times [0, \overline{\tau}])$ with the topology of convergence in distribution. Standard arguments (Aliprantis and Border (1999, Theorem 14.11)) show that this space is compact. By definition it is locally convex.

Upper hemi continuity of Bayes' rule: We now prove that the correspondence $\mathcal{B} : \Delta(B \times [0, \overline{\tau}]) \to B$ mapping beliefs F to the set of measurable trajectories $\{x_t\}$ that satisfy (4.1) for $t < \tau(F)$ is upper hemi continuous. Consider a sequence of beliefs F^n (with expectation $\mathbb{E}^n$) that converges to F in distribution. $\mathcal{B}(F^n)$ consists of all measurable trajectories $\{x_t^n\}$ that satisfy (4.1) (when replacing $\mathbb{E}^F$ by $\mathbb{E}^n$) for $t < \tau(F^n)$. As F assigns probability less than one to the event $\{\tilde{\tau} < t\}$ for any $t < \tau(F)$, so does F^n for sufficiently large n; thus, $\lim_{n\to\infty} \tau(F^n) \geq \tau(F)$.

We now show that $x_t^n \to x_t$ for all $t < \tau(F)$ at which the marginal distribution $F(\tilde{\tau})$ is continuous.[16] Consider the numerator of (4.1) (the argument for the denominator is identical). The integrand $\chi_t^-(\tilde{a}, \tilde{\tau}) := z_t(\tilde{a}) w_t(\tilde{a}) \mathbb{I}_{\{\tilde{\tau} > t\}}$ is continuous in $\tilde{a}$ (see, Davis (1993, Theorem 43.5)) and lower semi-continuous in $\tilde{\tau}$; similarly, $\chi_t^+(\tilde{a}, \tilde{\tau}) := z_t(\tilde{a}) w_t(\tilde{a}) \mathbb{I}_{\{\tilde{\tau} \geq t\}}$ is continuous in $\tilde{a}$ and upper semi-continuous in $\tilde{\tau}$. The portmanteau theorem thus implies $\liminf \mathbb{E}^n[\chi_t^-(\tilde{a}, \tilde{\tau})] \geq \mathbb{E}^F[\chi_t^-(\tilde{a}, \tilde{\tau})]$ and $\mathbb{E}^F[\chi_t^+(\tilde{a}, \tilde{\tau})] \geq \limsup \mathbb{E}^n[\chi_t^+(\tilde{a}, \tilde{\tau})]$. As χ_t^- and χ_t^+ are bounded and disagree only for $\tilde{\tau} = t$, which happens with probability zero under F and thus with vanishing probability under F^n, we have $\mathbb{E}^F[\chi_t^-(\tilde{a}, \tilde{\tau})] = \mathbb{E}^F[\chi_t^+(\tilde{a}, \tilde{\tau})]$ and $\limsup \mathbb{E}^n[\chi_t^+(\tilde{a}, \tilde{\tau})] = \limsup \mathbb{E}^n[\chi_t^-(\tilde{a}, \tilde{\tau})]$. Thus,

$$\liminf \mathbb{E}^n[\chi_t^-(\tilde{a}, \tilde{\tau})] \geq \mathbb{E}^F[\chi_t^-(\tilde{a}, \tilde{\tau})] = \mathbb{E}^F[\chi_t^+(\tilde{a}, \tilde{\tau})] \geq \limsup \mathbb{E}^n[\chi_t^+(\tilde{a}, \tilde{\tau})] = \limsup \mathbb{E}^n[\chi_t^-(\tilde{a}, \tilde{\tau})]$$

and so $\lim \mathbb{E}^n[\chi_t^-(\tilde{a}, \tilde{\tau})]$ exists and equals $\mathbb{E}^F[\chi_t^-(\tilde{a}, \tilde{\tau})]$ as desired. Thus $\{x_t^n\}$ converges to $\{x_t\}$ pointwise for almost all $t \in [0, \tau(F)]$, and therefore in the $L^2([0, \tau(F)], [0, 1])$-norm and a fortiori in the weak topology. As the set $\mathcal{B}(F)$ allows for any measurable trajectories after $\tau(F)$, all trajectories $\{x_t^n\}_{t \in [0, \overline{\tau}]} \in \mathcal{B}(F^n)$ are uniformly close to $\mathcal{B}(F)$; that is, $\mathcal{B}$ is upper hemi continuous.

Upper hemi continuity of the firm's best responses: Self-esteem z_t is continuous in $a \in B$, so the firm's payoff $\Pi(a, \tau; x)$ is continuous in a, τ and x, and thus also continuous in $F = F(a, \tau)$ (Aliprantis and Border (1999, Theorem 14.5)); thus Berge's maximum theorem implies that the best response mapping $BR : B \to \Delta(B \times [0, \overline{\tau}])$ is upper hemi continuous.

Summary: We have shown that $(x, F) \mapsto (\mathcal{B}(F), BR(x))$ is an upper hemi continuous, convex-valued mapping of the compact, locally convex, Hausdorff space $B \times \Delta(B \times [0, \overline{\tau}])$ to itself. The Kakutani-Fan-Glicksberg theorem therefore implies that this mapping has a fixed point; this fixed point constitutes an equilibrium.

[16] It can have at most countably many discontinuities.

A.5 Proof of Equation (5.2)

This proof is analogous to the proof of Lemma 2, given in Appendix A.1. With observable investment, the firm's payoff from strategy (a, τ) is given by

$$\hat{\Pi}(z_t) = \int_{s=t}^{\tau} e^{-\int_t^s r+\mu z_u du}(z_s - ca_s - k + \mu z_s \hat{\Pi}(1))ds.$$

Setting $\psi(s) = e^{-r(s-t)}\hat{\Pi}(z_s)$, $\rho(s) = \mu z_s$ and $\phi(s) = e^{-r(s-t)}(z_s - ca_s - k + \mu z_s \hat{\Pi}(1))$ yields equation (A.3). Applying Claim 1, equation (A.4) becomes

$$\hat{\Pi}(z_t) = \int_{s=t}^{\tau} e^{-r(s-t)}(z_s - ca_s - k + \mu z_s(\hat{\Pi}(1) - \hat{\Pi}(z_s))ds.$$

Taking the derivative and applying Claim 2 we get

$$\begin{aligned} \hat{\Pi}'(z_t) &= \int_{s=t}^{\tau} e^{-r(s-t)}\frac{dz_s}{dz_t}(1 + \mu(\hat{\Pi}(1) - \hat{\Pi}(z_s)) - \mu z_s \hat{\Pi}'(z_s))ds \\ &= \int_{s=t}^{\tau} e^{-\int_t^s r+\lambda+\mu(1-2z_u)du}(1 + \mu(\hat{\Pi}(1) - \hat{\Pi}(z_s)) - \mu z_s \hat{\Pi}'(z_s))ds. \end{aligned}$$

Setting $\rho(s) = \mu z_s$ and $\phi(s) = e^{-\int_t^s r+\lambda+\mu(1-2z_u)du}\mu(\hat{\Pi}(1) - \hat{\Pi}(z_s))$, $f(s) = e^{-\int_t^s r+\lambda+\mu(1-2z_u)du}\hat{\Pi}'(z_s)$ satisfies (A.4). Applying Claim 1, equation (A.3) becomes

$$\hat{\Pi}'(z_t) = \int_{s=t}^{\tau} e^{-\int_t^s r+\lambda+\mu(1-z_u)du}\mu(1 + \hat{\Pi}(1) - \hat{\Pi}(z_s))ds.$$

The envelope theorem then implies equation (5.2).

A.6 Proof of Theorem 6

Suppose otherwise, that $t^* \geq \hat{t}$. Then $\hat{z}_{\hat{t}} = z_{\hat{t}}(\overline{a}) \geq z^*_{\hat{t}}$ because $\hat{a}_t = \overline{a}$ on $[0, \hat{t}]$; since reputational drift $g(a, z)$ is negative for $z \geq \hat{z}_{\hat{t}}$ and any investment $\{a_t\}$, this implies $\hat{z}_{\hat{t}} \geq z^*_{t^*}$ and there exists $\tilde{t} \geq \hat{t}$ such that $\hat{z}_{\tilde{t}} = z^*_{t^*} =: z$. By definition both firms shirk after reaching self-esteem z and so the trajectories coincide, $\hat{z}_{\tilde{t}+t} = z^*_{t^*+t}$. For convenience we write this joint trajectory as z_t; that is, we restart the clock at $t = 0$ when self-esteem reaches z. Writing $x^*_t = x_{t^*+t}$ for the revenue of the unobservable firm, we have $x^*_t \leq z_t$ because the equilibrium strategy that invests the longest, until t^*, leads to the highest self-esteem z_t while x^*_t is a weighted average of z_t and the self-esteem resulting from equilibrium strategies with lower investment.

We now show that the observable firm with self-esteem z has strictly higher investment incentives than the unobservable firm $\hat{\Gamma}(\tilde{t}) > \Gamma(t^*) = c/\lambda$. This contradicts our assumption that the observable firm weakly prefers to shirk at time $\tilde{t} \geq \hat{t}$. The investment incentives of the two firms are given by

$$\hat{\Gamma}(\tilde{t}) = \int_0^{\hat{\tau}-\tilde{t}} e^{-\int_0^t r+\lambda+\mu(1-z_s)ds} \left[1+\mu(\hat{V}(1)-\hat{V}(z_t))\right] dt, \tag{A.6}$$

$$\Gamma(t^*) = \int_0^{\tau^*-t^*} e^{-\int_0^t r+\lambda+\mu(1-z_s)ds} \mu(V(0,1)-V(t^*+t,z_t))dt. \tag{A.7}$$

First, we argue that the 'observable' firm exits later than the 'unobservable' firm, i.e. $\hat{\tau}-\tilde{t} > \tau^*-t^*$. To see this, note that the former enjoys higher flow payoffs than the latter,

$$z_t - k + \mu z_t \hat{V}(1) > x_t^* - k + \mu z_t V(0,1),$$

because $z_t \geq x_t^*$ and the 'observable' firm's value at a breakthrough is higher than the 'unobservable' firm's, $\hat{V}(1) \geq V(0,1)$; the latter follows because the 'observable' firm can ensure itself weakly higher flow payoffs by mimicking the equilibrium strategy of the 'unobservable' firm.

Next, consider the integrands in (A.6) and (A.7). We can write the value functions as

$$\hat{V}(z_t) = \int_{s=t}^{\hat{\tau}-\tilde{t}} e^{-\int_t^s r+\mu z_u du} \left[z_s - k + \mu z_s \hat{V}(1)\right] ds$$

$$V(t^*+t,z_t) = \int_{s=t}^{\tau^*-t^*} e^{-\int_t^s r+\mu z_u du} \left[x_s^* - k + \mu z_s V(0,1)\right] ds$$

The second value decreases if we force the firm to exit at the suboptimal time $\hat{\tau}-\tilde{t}$ and also omit the positive revenue x_s^*, i.e.

$$V(t^*+t,z_t) > \int_{s=t}^{\hat{\tau}-\tilde{t}} e^{-\int_t^s r+\mu z_u du} \left[-k + \mu z_s V(0,1)\right] ds.$$

Taking differences

$$\hat{V}(z_t) - \hat{V}(t^*+t,z_t) < \int_{s=t}^{\hat{\tau}-\tilde{t}} e^{-\int_t^s r+\mu z_u du} z_s \left[1+\mu(\hat{V}(1)-V(0,1))\right] ds < \frac{1}{\mu} + \hat{V}(1) - V(0,1),$$

where the last inequality follows from

$$\int_{s=t}^{\hat{\tau}-\tilde{t}} e^{-\int_t^s r+\mu z_u du} z_s ds < \int_{s=t}^{\infty} -\frac{1}{\mu}\frac{d}{ds}\left(e^{-\int_t^s \mu z_u du}\right) ds = \frac{1}{\mu}\left[1-\exp\left(\int_t^{\infty} \mu z_u du\right)\right] \leq \frac{1}{\mu}.$$

Rearranging, we get

$$1+\mu(\hat{V}(1)-\hat{V}(z_t)) > \mu(V(0,1)-V(t^*+t,z_t)).$$

Thus both the integrand and the integration domain are larger in (A.6) than in (A.7), implying $\hat{\Gamma}(\tilde{t}) > \Gamma(t^*)$ and completing the proof.

References

Abito, J.-M., D. Besanko, and D. Diermeier (2012): "Corporate Reputational Dynamics, Private Regulation, and Activist Pressure," Working paper, Northwestern.

Aliprantis, C. D., and K. C. Border (1999): *Infinite Dimensional Analysis*, vol. 32006. Springer.

Atkeson, A., C. Hellwig, and G. Ordonez (2012): "Optimal Regulation in the Presence of Reputation Concerns," Working paper, Penn.

Bar-Isaac, H. (2003): "Reputation and Survival: learning in a dynamic signalling model," *Review of Economic Studies*, 70(2), 231–251.

Board, S., and M. Meyer-ter-Vehn (2013): "Reputation for Quality," *Econometrica*, 81(6), 2381–2462.

Bonatti, A., and J. Hörner (2011): "Collaborating," *American Economic Review*, 101, 632–663.

——— (2013): "Career Concerns and Market Structure," Working paper, MIT.

Bronnenberg, B. J., J.-P. H. Dubé, and M. Gentzkow (2012): "The Evolution of Brand Preferences: Evidence from Consumer Migration," *American Economic Review*, 102(6), 2472–2508.

Cabral, L. (2014): "Living Up To Expectations: Corporate Reputation and Sustainable Competitive Advantage," Working paper, NYU.

Cisternas, G. (2014): "Two-Sided Learning and Moral Hazard," Working paper, MIT.

Davis, M. (1993): *Markov Models and Optimization*. Chapman & Hall.

Diermeier, D. (2011): *Reputation Rules*. McGraw Hill.

Ericson, R., and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62(1), 53–82.

Foster, L., J. Haltiwanger, and C. Syverson (2013): "The Slow Growth of New Plants: Learning about Demand," Working paper, Chicago.

Gale, D., and R. Rosenthal (1994): "Price and Quality Cycles for Experience Goods," *RAND Journal of Economics*, 25(4), 590–607.

Goldfarb, A. (2007): "Schlitz: Why Schlitz Hit the Fan," in *Industry and Firm Studies, 4th edition*, ed. by V. J. Tremblay, and C. H. Tremblay, pp. 321–341. M. E. Sharpe Inc.

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–182.

Hopenhayn, H. A. (1992): "Entry, Exit, and firm Dynamics in Long Run Equilibrium," *Econometrica*, 60(5), 1127–1150.

Jovanovic, B. (1982): "Selection and the Evolution of Industry," *Econometrica*, 50(3), 649–670.

Keller, G., and S. Rady (2014): "Breakdowns," *Theoretical Economics*, forthcoming.

Kovrijnykh, A. (2007): "Career Uncertainty and Dynamic Incentives," Working paper, University of Chicago.

Mailath, G., and L. Samuelson (2001): "Who Wants a Good Reputation," *Review of Economic Studies*, 68(2), 415–441.

Milgrom, P., and I. Segal (2002): "Envelope theorems for arbitrary choice sets," *Econometrica*, 70(2), 583–601.

Rob, R., and A. Fishman (2005): "Is Bigger Better? Customer Base Expansion through Word-of-Mouth Reputation," *Journal of Political Economy*, 113(5), 1146–1175.

Sannikov, Y. (2014): "Moral Hazard and Long-Run Incentives," Working paper, Princeton University.

Syverson, C. (2011): "What Determines Productivity?," *Journal of Economic Literature*, 49(2), 326–365.