

A Reputational Theory of Firm Dynamics[†]

By SIMON BOARD AND MORITZ MEYER-TER-VEHN*

We study the life cycle of a firm that produces a good of unknown quality. The firm manages its quality by investing while consumers learn via public breakthroughs; if the firm fails to generate such breakthroughs, its revenue falls and it eventually exits. Optimal investment depends on the firm's reputation (the market's belief about its quality) and self-esteem (the firm's own belief about its quality), and is single-peaked in the time since a breakthrough. We derive predictions about the distribution of revenue and propose a method to decompose the impact of policy changes into investment and selection effects. (JEL D11, D21, D25, D83, G31, L15)

In many markets, consumers have imperfect information about the quality of the products or services they purchase. A restaurant patron would like to know about the hygiene of the establishment, a patient undergoing surgery would like to know the competence of their surgeon, and an investor would like to know about the returns of a mutual fund. This paper investigates the incentives that firms (e.g., restaurants, surgeons, mutual funds) have to invest in order to maintain quality. In particular, we study investment and exit decisions over their life cycles and derive implications for the dynamics of reputation and the cross-sectional distribution of revenue.

Our paper provides a framework to study investment and selection effects in industries where the market learns about quality over time. In contrast to models of perfect information, like Ericson and Pakes (1995), a firm's most important assets are the market's belief in its quality (its reputation) and the firm's own belief in its quality (its self-esteem). These state variables differ from traditional capital assets in important aspects. First, reputation depends on the market's beliefs about the firm's investment, rather than actual investment. The resulting moral hazard problem dampens incentives and qualitatively changes firm dynamics. Specifically, we suppose customers learn about quality from public breakthroughs; investment then decreases in the time after a breakthrough when investment is publicly observable but is single-peaked when investment is unobservable. Second, in a perfect information model, price-taking firms invest and exit efficiently, leaving no role for policy to

*Board: UCLA (email: sboard@econ.ucla.edu); Meyer-ter-Vehn: UCLA (email: mtv@econ.ucla.edu). Leslie Marx was coeditor for this article. We have received helpful comments from Andy Atkeson, Heski Bar-Isaac, V. Bhaskar, Alessandro Bonatti, Andrew Clausen, Christian Hellwig, Hugo Hopenhayn, Johannes Hörner, Larry Samuelson, Yuliy Sannikov, Connan Snider, Chad Syverson, and seminar audiences at Bocconi University, University of Chicago GSB, Essex, LSE, Princeton University, PSE, University of Toronto, TSE, Warwick, the Duke IO Theory Conference, ES Winter Meetings, Gerzensee, Mannheim Reputation Conference, SAET, SED, and SWET. We gratefully acknowledge financial support from NSF grant 0922321.

[†]Go to <https://doi.org/10.1257/mic.20190376> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

improve welfare. By contrast, our model permits meaningful analysis of disclosure policies that improve the market's information, e.g., requiring restaurants to post health grades or publishing surgeon report cards. To this end, we propose a method to decompose the impact of such policies on investment and selection effects. Third, reputation varies as news arrives and is thus volatile even if the underlying product quality is stable. For example, when a restaurant is featured in the *New York Times*, demand jumps even though the quality does not change. Hence, one needs a model to be able to infer the underlying quality from volatile restaurant revenue; we discuss how our model can do exactly this.

In the model, a long-lived firm sells a product of high or low quality to a continuum of identical, short-lived consumers. The firm chooses how much to invest in the quality of its product; it can also exit at any time. Consumers observe neither the firm's investment history nor the resulting quality. Rather, they learn about quality via public breakthroughs that can only be generated by a high-quality product. We call the market's belief that quality is high, x_t , the firm's *reputation* and assume that it determines revenue. Like the market, the firm does not observe its quality directly, but learns via public breakthroughs; unlike the market, it can also recall its past investments. We call the firm's belief that its quality is high, z_t , the firm's *self-esteem*.

In a pure strategy equilibrium, reputation and self-esteem coincide on-path. However, since the market does not observe deviations from equilibrium, investment incentives are determined by the marginal value of self-esteem off-path. We characterize the firm's optimal investment and exit decisions over its life cycle, focusing on equilibria where reputation is solely a function of the time since the last breakthrough.

Our first result represents the marginal value of self-esteem and hence investment incentives as an integral over future reputational dividends that derive from an increased chance of breakthroughs. Using this characterization, we show that in any pure strategy equilibrium, incentives are single-peaked in the time since the last breakthrough (Theorem 1). Immediately after a breakthrough, the firm is known to be high quality and does not benefit from another breakthrough. At the opposite extreme, when the firm is about to exit, it cuts investment to zero. Intuitively, investment only pays off if investment affects quality *and* quality is revealed via a breakthrough. When the firm is dt from exiting, the joint probability of these two events is of order dt^2 , and therefore the marginal benefit of investment vanishes. Investment incentives are thus maximized at an intermediate time; we formally show that they are single-peaked. We also prove equilibrium exists (Theorem 2), possibly in mixed strategies. The proof defines a topology on strategies that allows us to apply the Kakutani-Fan-Glicksberg fixed-point theorem.¹

To illustrate the applicability of the model, we use the Kolmogorov forward equation to examine how learning, investment, and entry jointly shape the steady-state distribution of reputation. Additionally, we propose a method to decompose the drivers of steady-state revenue into investment and selection effects. This decomposition

¹ Other papers—e.g., Cisternas (2017)—establish existence constructively from the firm's adjoint equation. The problem with this approach in our setting is that the firm's first-order condition for investment is not sufficient for optimality; in particular, investments are dynamic complements, and so multistep deviations may be optimal.

is particularly useful in interpreting comparative statics, such as a disclosure policy that raises the frequency of breakthroughs or a minimum wage that raises operating costs. As a proof of concept, we use a simple simulation to show that a disclosure policy may have little effect on mean steady-state reputation but drastically raises the relative importance of investment over selection.

Section III considers two model variants that isolate the two different economic forces behind the single-peaked incentives in the baseline model. They also illustrate how the information structure in the market affects firms' incentives to invest. In the first variant, we assume the market observes the firm's investment. Thus, both the firm and the market symmetrically learn about the firm's quality, but there is no moral hazard. We show that investment *decreases* in the time since a breakthrough, as the firm approaches bankruptcy (Theorem 3). In the second variant, we assume the firm knows its own quality in addition to its investment, while the market knows neither. We show that investment *increases* in the time since a breakthrough, as the benefit from a breakthrough grows (Theorem 4). In stark contrast to the baseline model, investment is maximized at the exit threshold; crucially, the firm immediately observes when investment raises quality and then chooses to remain in the market. The single-peaked investment in the baseline model combines the decreasing incentives in the first variant and the increasing incentives in the second variant.

Our model provides a natural lens through which to study the reputation and revenue dynamics of restaurants, surgeons, and mutual funds. In Section IV, we discuss how one might identify the model parameters from data in these industries. We also discuss how to generalize the model to accommodate competition and more general learning structures.

Literature.—Our model draws inspiration from two canonical models of firm dynamics. Jovanovic (1982) assumes that firms and consumers symmetrically learn about the firms' quality but abstracts from investment. Ericson and Pakes (1995) assume firms invest in the quality of their products but suppose there is complete information. We combine learning and investment and characterize how these forces jointly determine the evolution of the firm's quality and reputation. Thus, we can study how information policies (e.g., restaurant grades, surgeon reports) affect investment incentives, firm dynamics, and turnover.

There are a variety of other models of firm life cycles. Hopenhayn (1992) assumes firm capabilities change over time according to a Markov process and studies the resulting entry and exit patterns. Cabral (2016) considers a reduced-form model of reputational firm dynamics, where reputation is modeled as a state variable akin to capital stock rather than being derived from Bayes' rule. Gale and Rosenthal (1994) and Rob and Fishman (2005) consider the dynamics of repeated game equilibria where incentives arise from punishment strategies.

The paper is related to the literature on reputation, especially our prior paper (Board and Meyer-ter-Vehn 2013). The current paper introduces exit, allowing us to study firm life cycles; this is of interest for three reasons. First, the possibility of exit qualitatively changes the nature of incentives, as illustrated by the difference between the single-peaked incentives in the baseline model and the increasing incentives when the firm knows its quality. (Without exit, there is no difference

between these models.) Second, the model provides an empirical framework to study industry turnover and productivity dispersion and to measure the impact of disclosure policies and minimum wages. Third, the exit decision introduces new methodological challenges, since we must keep track of both reputation and self-esteem. (Without exit, we need only keep track of reputation.) As such, 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. In particular, Bonatti and Hörner (2011, 2017) consider 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 (2017) analyzes a general model of two-sided learning with moral hazard; his incentive equation is analogous to our "marginal value of self-esteem."

Several papers endogenize the flow of information in reputation models. Halac and Prat (2016) suppose that the probability of a signal depends not only on the "quality of the firm" but also on an additional "action of the market"; after a breakthrough, investment incentives strictly increase and ultimately plateau, in contrast to the eventually vanishing incentives here. Hauser (2021a, 2021b) supposes that firms can speed up good news or hide bad news. Marinovic, Skrzypacz, and Varas (2018) allow firms to pay to certify their quality. And Vellodi (2021) asks which information structure maximizes social welfare.

There is also a literature on "name trading." Tadelis (1999) considers firms with exogenous quality that can sell or abandon their names; Tadelis (2002) and Mailath and Samuelson (2001) introduce moral hazard. These papers show that high-quality firms would like to acquire names with intermediate reputations, where Bayesian updating is fastest. Vial and Zurita (2017) consider a highly tractable model in which types are exogenous, and firms can only abandon bad names. Ultimately, the incentives to purchase a name are quite different from the incentives to invest in quality: a new name boosts reputation for a given firm quality and can be quickly run down if the firm is low-quality, while investment boosts firm quality for a given reputation and thereby changes the future evolution of reputation.²

Our paper is the first to study investment, learning, and exit in a single model. It provides a canonical model of how information policies affect industry dynamics and thereby speaks to a growing empirical literature. Jin and Leslie (2003, 2009) and Luca (2016) study how hygiene grades and Yelp scores affect restaurants. Cutler, Huckman, and Landrum (2004) and Kolstad (2013) study how medical report cards affect the behavior of surgeons. Cabral and Hortaçsu (2010) consider the impact of negative feedback on eBay sellers. There are many more papers that examine the effect of new information on customers' decisions, such as health plan choice (Jin and Sorensen 2006), hospital choice (Pope 2009), doctors' prescriptions (Arrow, Bilir, and Sorensen 2020), school choice (Hastings and Weinstein 2008), college

²Our model has implications for the trade of underlying knowledge, such as patents, over the life cycle of the firm (e.g., Serrano 2010). Indeed, assume that a patent raises a firm's quality with a probability that is constant across firms and known to all firms. Low-reputation firms would then sell their intellectual property to high-reputation firms because of the convexity of firm value $V(t, z)$ in z , as established in Section IIA.

applications (Luca and Smith 2013), caloric intake (Bollinger, Leslie, and Sorensen 2011), and the demand for books (Sorensen 2007).

The paper provides a model of brand dynamics, which are often modeled in an informal way. For example, Bronnenberg, Dubé, and Gentzkow (2012) study the dynamics of brand shares when customers move between cities and their preferences depend on past purchases. Similarly, Foster, Haltiwanger, and Syverson (2016) 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 a customer base or building a reputation.” More broadly, the paper helps explain how R&D investments feed into knowledge and firm’s values (e.g. Hall, Jaffe, and Trajtenberg 2005) and how productivity and profits vary within industries, with some firms investing in their assets and growing, while others disinvest and shrink (e.g. Syverson 2011).

I. Model

Players and Actions: A long-lived firm faces a mass of short-lived consumers, also referred to as *the market*. Time $t \in [0, \infty)$ is continuous. The firm chooses a stochastic process of *investment* levels $A_t \in [0, \bar{a}]$ where $\bar{a} < 1$; it also chooses a stopping time $T \in [0, \infty]$ at which it exits the market.

Technology: 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 $\mu\theta$; that is, breakthroughs only occur when $\theta = H$. We write h^t for histories of breakthrough arrival times before time t , and h for infinite histories.

The firm does not observe product quality, either, but does recall its past actions. Its investment plan $\{A_t\}_{t \geq 0}$ and exit time T are thus progressively measurable with respect to the filtration induced by public histories h^t . In turn, the investments $A := \{A_t\}_{t \geq 0}$ control the distribution of quality $\{\theta_t\}_{t \geq 0}$ and thereby the histories of breakthroughs h ; we write E^A for expectations under this measure and call $Z_t = E^A[\theta_t | h^t]$ the firm’s *self-esteem* at time $t < T$. This reflects the firm’s belief in 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} . The firm’s *reputation* is given by $X_t := E^{\tilde{A}}[\theta_t | h^t]$. For concreteness, we apply this formula at all times, including $t > \tilde{T}$; in other words, the market draws no inference about the firm’s past investments after it fails to exit as

expected.³ Initially, self-esteem and reputation are exogenous and coincide; $X_0 = Z_0$.

Payoffs: The firm has flow profits $\pi(X_t) - c(A_t)$ and discount rate $r > 0$. The firm's income $\pi(\cdot)$ is smooth and strictly increasing with boundaries $\pi(0) < 0$ and $\pi(1) > 0$; we further restrict $\pi(\cdot)$ below in (4). As a running example, we use $\pi(X) = X - k$, where X is revenue and $k \in (0, 1)$ is the firm's operating cost. The firm's investment cost $c(\cdot)$ is smooth, strictly increasing, and convex, with $c(0) = 0$.

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

$$(1) \quad E^A \left[\int_{t=0}^T e^{-rt} (\pi(X_t) - c(A_t)) dt \right].$$

To highlight the distinct roles of market beliefs \tilde{A} and actual investment A in (1), note that \tilde{A} determines the firm's reputation $X_t = E^{\tilde{A}}[\theta_t | h^t]$ for a given history h^t , while A determines the distribution over histories h^t .

A Restriction on 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 strategies that only depend on the time since the last breakthrough. Formally, we assume there exists a deterministic process $a = \{a_t\}$ and exit time $\tau \in [0, \infty]$ such that if the last breakthrough before t was at time s , then $A_t = a_{t-s}$ and $\tau = T - s$. We write such strategies as (a, τ) and the resulting self-esteem as $z = \{z_t\}$, where $z_0 = 1$. Similarly, define beliefs $(\tilde{a}, \tilde{\tau})$ and denote the induced reputation by $x = \{x_t\}$, where $x_0 = 1$. Given reputation $x = \{x_t\}$, the firm's problem resets at a breakthrough and so strategies (a, τ) are indeed optimal, justifying the restriction.

Truncating the integral in (1) at the first breakthrough (which arrives at rate μz_t), the firm's continuation value at time t is

$$(2) \quad V(t, z_t) = \sup_{a, \tau} \int_{s=t}^{\tau} e^{-\int_t^s (r + \mu z_u) du} [\pi(x_s) - c(a_s) + \mu z_s V(0, 1)] ds.$$

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

³This implies that reputation is continuous at \tilde{T} . By contrast, if off-path beliefs were such that reputation jumps down at time \tilde{T} , effort would still be single-peaked (Theorem 1), but the exit time need not satisfy the indifference equation (9). We ignore such equilibria because it is implausible for the market to interpret failure to exit as a signal of low quality. Alternatively, if off-path beliefs induced reputation to jump up at time \tilde{T} , the firm would like to stay in the market, which cannot arise in a pure strategy equilibrium. When we extend the model to mixed beliefs in the context of equilibrium existence in Appendix AD, we account for the market's on-path inferences from the firm's failure to exit and allow for richer off-path inferences.

Reputational Dynamics: Self-esteem is governed by the firm's investment and the history of breakthroughs. 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 is given by

$$(3) \quad g(a, z) = \lambda(a - z) - \mu z(1 - z)$$

as in, for example, Board and Meyer-ter-Vehn (2013). 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 exceeds the firm's self-esteem, $a > z$, and negative otherwise. The second term derives from the absence of breakthroughs and is always negative. Analogously, reputation is governed by believed investment \tilde{a} and the history of breakthroughs, jumping to 1 at a breakthrough and in its absence following $\dot{x}_t = g(\tilde{a}_t, x_t)$.

We want to guarantee that in the absence of a breakthrough, the firm eventually exits. To this end, assume

$$(4) \quad \pi(z^\dagger) + \mu z^\dagger \pi(1)/r < 0,$$

where $z^\dagger \in (0, 1)$ is the unique level of self-esteem where reputational drift vanishes under maximal investment, $g(\bar{a}, z^\dagger) = 0$.⁴ So defined, drift (3) is strictly negative on $[z^\dagger, 1]$ for any beliefs $(\tilde{a}, \tilde{\tau})$, and (absent a breakthrough) reputation and self-esteem eventually drop below z^\dagger . At that point, (4) ensures that the integrand in (2) is negative and the firm exits, where $\pi(1)/r$ serves as an upper bound for $V(0, 1)$.

Remarks: The model makes several assumptions of note. First, as in Board and Meyer-ter-Vehn (2013), we assume that quality at time t is based on investment at the time of the last technology shock. While “upgrades” are natural, there are two ways to think of “downgrades” in our model. One can interpret investment as the choice of absorptive capacity, determining the ability of a firm to adapt to a changing world (Cohen and Levinthal 1990). Alternatively, one can think of a firm's quality as its advantage over a competitive fringe that advances one rung on a ladder of quality at each technology shock. Our specification has two substantive implications. It implies that high-quality and low-quality firms are equally good at investing, so we do not hardwire in any complementarity or substitutability between current quality and investment.⁵ It also means that believed investment affects reputation continuously, via its drift (3); this contrasts with other recent models of endogenous persistent states, where equilibrium beliefs can lead to jumps in reputation (e.g., Dilmé 2019, Halac and Prat 2016).

⁴Note, z^\dagger is well defined since $g(\bar{a}, 0) > 0 > g(\bar{a}, 1)$, and $g(a, z)$ is convex in z . It satisfies $z^\dagger = \lambda(\bar{a} - z^\dagger)/\mu(1 - z^\dagger)$ and is thus smaller than \bar{a} and λ/μ .

⁵We could assume, equivalently, that 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 $\lambda/2$.

Second, we assume that the firm does not know its own quality and learns from the public signals. In the context of our motivating examples, a restaurant owner tries to raise its quality by changing the menu and decor and learns whether this has been successful (in the eyes and taste buds of customers) by reading newspaper or Yelp reviews. A surgeon tries to improve her skills by taking courses and talking to other doctors and learns whether this has worked when she faces a patient with complications. Or, a mutual fund tries to raise its returns by appointing a new team of analysts and learns their effectiveness when it wins an industry award. As a contrast, we study the case of known quality in Section IIIB and discuss intermediate cases in Appendix BE.

Third, we assume “perfect good news learning,” where information arrives via breakthroughs that reveal high quality. This stylized information structure allows us to analyze our model in terms of strategies (a, τ) that depend only on the time since the last breakthrough rather than Markovian strategies that depend on reputation and self-esteem; this is key for our equilibrium existence proof.⁶ The intuition behind our single-peaked incentives extends to imperfect good-news Poisson or Brownian signals. We discuss this extension in Section IV.

Finally, while our formal results concern the life cycle of a single firm, one can think of the firm as operating in a competitive industry in steady state, such as the restaurant industry or the market for surgeons. This market interpretation allows us to consider the implications of the model for the cross-sectional distribution of revenue. We discuss this extension in Section IV.

II. Analysis

Section IIA analyzes the firm’s optimal strategy for general (pure) market beliefs, showing that investment incentives are single-peaked in the time since a breakthrough. The key step is to express the marginal value of self-esteem as an integral of a series of dividends. Section IIB proves equilibrium existence, while Section IIC characterizes the cross-sectional distribution of revenue. Section IID simulates the model and decomposes the drivers of reputation into investment and selection effects. Finally, Section IIE shows how to use the model to quantify the impact of disclosure policies and minimum wages.

A. The Firm’s Problem

We begin with some preliminary observations about reputation and firm value for arbitrary pure market beliefs.

⁶In a pure strategy equilibrium, $\{x_t\}$ is decreasing and thus invertible, so our formulation with state variables (t, z) is equivalent to a standard Markovian formulation with state variables (x, z) . We leverage the (t, z) formulation when considering mixed-strategy equilibria with potentially non-monotonic reputation $\{x_t\}$ in Section IIB. In particular, if the market has mixed beliefs over Markovian investment strategies, reputation x_t is not a sufficient statistic for market beliefs. In this case, one would have to keep track of the market’s belief about z_t —that is, the market’s second-order belief about quality.

LEMMA 1: For any pure beliefs $(\tilde{a}, \tilde{\tau})$,

- (i) reputation $\{x_t\}$ is continuous and strictly decreasing,
- (ii) an optimal strategy (a^*, τ^*) exists and $\tau^* < \infty$, and
- (iii) firm value $V(t, z)$ strictly decreases in t and strictly increases in z .

PROOF:

Continuity of $\{x_t\}$ follows by the fact that $\dot{x} = g(\tilde{a}, x)$, leveraging our assumption that the market makes no inference from failure to exit at $\tilde{\tau}$. Existence of an optimal strategy follows from standard compactness arguments in the proof of equilibrium existence, Theorem 2. Monotonicity of $\{x_t\}$ and exit in finite time, $\tau^* < \infty$, follow from the fact that reputational drift $g(\tilde{a}, x)$ is negative for $x > z^\dagger$ and, by assumption (4), that the firm exits before x reaches z^\dagger . Since $\{x_t\}$ falls in t , so too does $V(t, z)$. Finally, $V(t, z)$ rises in z since self-esteem raises the chance of a breakthrough that raises reputation and income. See Appendix AA for a formal argument on the last point. ■

The optimal strategy is characterized by the firm's Hamilton Jacobi Bellman (HJB) equation

$$(5) \quad rV(t, z) = \max_a \left[\pi(x_t) - c(a) + g(a, z)V_z(t, z) + V_t(t, z) + \mu z(V(0, 1) - V(t, z)) \right]^+,$$

where $y^+ := \max\{y, 0\}$, capturing the firm's ability to exit for a continuation value of 0. Investment raises the firm's quality, its self-esteem, and thus its value. In particular, equation (3) implies that investment raises self-esteem at rate λ , so the marginal benefit of investing is $\lambda V_z(t, z_t^*)$ whenever this derivative exists. Optimal investment $a_t^* \in (0, \bar{a})$ thus satisfies the first-order condition

$$(6) \quad \lambda V_z(t, z_t^*) = c'(a_t^*)$$

with $a_t^* = 0$ if $\lambda V_z(t, z_t^*) < c'(0)$ and $a_t^* = \bar{a}$ if $\lambda V_z(t, z_t^*) > c'(\bar{a})$.

To address the differentiability issue, observe that the firm's value $V(t, z)$ is convex in z . This follows because firm payoff for any fixed strategy (a, τ) is linear in z , and $V(t, z)$ is the upper envelope of such linear functions. Intuitively, convexity captures the value of information. For example, if the firm exits at (t, z) , information about quality is valuable since good news raises self-esteem and value while the option to exit protects the firm from bad news; that is, $0 = V(t, z - \epsilon) = V(t, z) < V(t, z + \epsilon)$. Convexity implies that even if the value function is not differentiable, it admits directional derivatives $V_{z^-}(t, z), V_{z^+}(t, z)$.

Given the first-order condition (6), we need to understand the marginal value of self-esteem. Our next result, the workhorse of this paper, expresses $V_z(t, z_t^*)$ in terms of future reputational dividends.

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

$$(7) \quad \Gamma(t) := \int_t^{T^*} e^{-\int_t^s (r + \lambda + \mu(1 - z_u^*)) du} \mu [V(0, 1) - V(s, z_s^*)] ds.$$

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

PROOF.

Equation (7) is an integral version of the adjoint equation for the firm's control problem. It follows by differentiating firm payoff for a fixed strategy (a^*, τ^*) with respect to the firm's initial self-esteem z . This coincides with the (directional) derivative of the value function by the envelope theorem. See Appendix AB for details. ■

Intuitively, quality and self-esteem raise the probability of a breakthrough and, since they are persistent, pay off dividends over time. That is, incremental self-esteem dz raises the probability of a breakthrough in $[s, s + ds]$ by $\mu dz ds$, while the value of a breakthrough equals $V(0, 1) - V(s, z_s^*)$. 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_u^* , a breakthrough arrives at $u \in [t, s]$, self-esteem jumps to 1, and the increment disappears; third, reputational drift (3) is not constant in z , and its derivative equals $g_z(a_u^*, z_u^*) = -(\lambda + \mu(1 - 2z_u^*))$. Summing these three components yields the discounting term in (7).

If there is a single optimal strategy (a^*, τ^*) , the derivative $V_z(t, z_t^*)$ exists and coincides with $\Gamma(t)$. If there are multiple optimal strategies, then $V_z(t, z_t^*)$ does not exist and each solution gives rise to a different trajectory of self-esteem z_t^* and a different value of $\Gamma(t)$. However, for any (a^*, τ^*) , equation (7) is well-defined, bounded by the directional derivatives of V , and describes the firm's investment incentives given (a^*, τ^*) . This implies the following necessary condition for optimal investment.

LEMMA 3 (Optimal Investment): *Any optimal strategy (a^*, τ^*) satisfies*

$$(8) \quad \lambda \Gamma(t) = c'(a_t^*)$$

with $a_t^ = 0$ if $\lambda \Gamma(t) < c'(0)$ and $a_t^* = \bar{a}$ if $\lambda \Gamma(t) > c'(\bar{a})$ for almost all t .⁷*

⁷ Changing investment at a measure-zero set of times does not affect payoffs, so any statements about optimal investment hold only almost always; hereafter, we omit this qualifier. It can also be eliminated by restricting the firm to forward-continuous strategies (see Board and Meyer-ter-Vehn 2013).

PROOF:

For a fixed optimal strategy (a^*, τ^*) with associated trajectory $\{z_t^*\}$, equation (7) captures the marginal benefit of increasing z_t^* . So, if $\lambda\Gamma(t) > c'(a_t^*)$ for a positive measure of times t , then the firm could raise its payoff by raising its investment a_t at those times. ■

THEOREM 1: *For any pure beliefs $(\tilde{a}, \tilde{\tau})$, optimal investment $\{a_t^*\}$ is single-peaked in the time since a breakthrough t ;⁸ at the exit threshold, $a_{\tau^*}^* = 0$. The optimal exit time τ^* satisfies*

$$(9) \quad \pi(x_{\tau^*}) + \mu z_{\tau^*} V(0, 1) = 0.$$

PROOF:

We wish to show that $\Gamma(t)$ is single-peaked in t with boundary conditions $\Gamma(0) > 0$, $\dot{\Gamma}(0) > 0$ and $\Gamma(\tau^*) = 0$. Taking the derivative of investment incentives (7) and setting $\rho(t) := r + \lambda + \mu(1 - z_t^*)$ yields the adjoint equation

$$(10) \quad \dot{\Gamma}(t) = \rho(t)\Gamma(t) - \mu(V(0, 1) - V(t, z_t^*)).$$

Now assume that $\rho(t)$ and $V(t, z_t^*)$ are differentiable. Then $\dot{\rho}(t) = -\mu\dot{z}_t^*$ and $(d/dt)V(t, z_t^*) = \dot{z}_t^*\Gamma(t) + V_t(t, z_t^*)$; in Appendix AC, we show that these functions are indeed absolutely continuous and extend our arguments to that case. Differentiating (10),

$$(11) \quad \ddot{\Gamma}(t) = \rho(t)\dot{\Gamma}(t) + \dot{\rho}(t)\Gamma(t) + \mu\frac{d}{dt}V(t, z_t^*) = \rho(t)\dot{\Gamma}(t) + \mu V_t(t, z_t^*).$$

Since $V_t(t, z_t^*) < 0$, $\dot{\Gamma}(t) = 0$ implies $\ddot{\Gamma}(t) < 0$; hence $\Gamma(t)$ is single-peaked, increasing or decreasing. To see that it is actually single-peaked, observe that when $t = 0$, equation (7) implies $\Gamma(0) > 0$ because the integrand $\mu(V(0, 1) - V(s, z_s^*))$ is strictly positive for $s \in [0, \tau^*]$. Equation (10) then implies $\dot{\Gamma}(0) = \rho(0)\Gamma(0) > 0$. Finally, when $t = \tau^*$, equation (7) immediately implies that $\Gamma(\tau^*) = 0$.

To understand the exit condition (9), recall that the firm's value is given by (2) and that $\{x_t\}$ is continuous and decreasing. When the firm shirks, its flow payoff is $\pi(x_t)$ and its option value of staying in the market has a flow value of $\mu z_t^* V(0, 1)$. Thus, if $\pi(x_t) + \mu z_t^* V(0, 1) > 0$, the firm can raise its payoff by shirking and staying in the market. Conversely, if $\pi(x_{\tau^*}) + \mu z_{\tau^*}^* V(0, 1) < 0$, then $\pi(x_t) - c(a_t) + \mu z_t^* V(0, 1) < 0$ for t just before τ^* and any investment a_t , and the firm would have been better off exiting a little earlier. ■

The evolution of investment incentives is shaped by two countervailing forces. Just after a breakthrough, an additional breakthrough has no value and the reputational dividend is zero; investment incentives (7) depend on current and future

⁸ Investment a_t^* is *single-peaked* if it increases on $[0, s]$ and decreases on $[s, \tau^*]$ for some $s \in (0, \tau^*)$.

dividends, so $\Gamma(0)$ is small but positive. As time progresses, future larger dividends draw closer and $\Gamma(t)$ rises. At the other extreme, investment vanishes at the exit time τ^* because there is no time left for the investment to pay off. That is, the benefit of investment is of the second order because both a technology shock and a breakthrough must arrive in the remaining time interval for the investment to avert exit. In Sections IIIA and IIIB, we present two model variants where we switch off each of these forces in turn and show that investment incentives are, respectively, decreasing and increasing in the time since a breakthrough.

This intuition suggests that investment incentives are small at the extremes. More strongly, Theorem 1 shows that incentives are single-peaked. A more precise intuition is as follows. As t rises, the firm foregoes the reputational dividends over $[t, t + dt]$, as captured by the second term in (10). This negative effect becomes more important over time as the reputational dividend increases; this is captured by the negative term $\mu(d/dt)V(t, z_t^*)$ in (11). On the upside, an increase in t brings future and larger dividends closer, as captured by the first term in (10). 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.

When exit is imminent, the firm ceases to invest, accelerating its demise. This force can matter in practice: for example, Goldfarb (2007) argues that the brewer Schlitz realized that the rise of Miller would have a large impact on its future profitability. This led it to disinvest in the brand by changing the preservatives, switching to lower-quality accelerated batch fermentation, and firing much of its marketing team. This pattern is not seen in classic models of industry dynamics like Jovanovic (1982) and Bar-Isaac (2003) that abstract from investment, meaning the reputation of the firm is a martingale.

B. Equilibrium

So far we have studied a firm's optimal strategy for arbitrary, pure beliefs $(\tilde{a}, \tilde{\tau})$. In this section, we close the model by assuming that market beliefs are correct. Such equilibrium dynamics can take different forms. To illustrate this, suppose cost is linear: $c(a) = ca$ for $a \in [0, \bar{a}]$. By Lemma 3, optimal investment is bang-bang with $a_t^* = \bar{a}$ when $\lambda\Gamma(t) > c$ and $a_t^* = 0$ when $\lambda\Gamma(t) < c$. In any pure strategy equilibrium, Theorem 1 tells us that investment incentives $\Gamma(t)$ are single-peaked, so there are two cases.

When costs are low, $\Gamma(0) > c = \Gamma(t_1)$ for some t_1 , the firm chooses $a = \bar{a}$ on $[0, t_1]$ and $a = 0$ on $[t_1, \tau]$. We call this a "probationary equilibrium" since the market assumes a firm invests for a fixed period of time after each breakthrough but then grows suspicious if no breakthrough is forthcoming. The firm's reputation initially drifts down slowly, as the favorable beliefs about investment offset the bad news conveyed by the lack of breakthroughs. After enough time without a breakthrough, market beliefs turn against the firm and the perceived disinvestment hastens the firm's decline.

When costs are high, $\Gamma(0) < c = \Gamma(t_0) = \Gamma(t_1)$ for some $t_0 < t_1$, the firm chooses $a = 0$ on $[0, t_0]$, $a = \bar{a}$ on $[t_0, t_1]$, and $a = 0$ on $[t_1, \tau]$. Here, the firm's

initial incentives $\Gamma(0)$ are insufficient to motivate effort. After a breakthrough, the 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.

We now prove equilibrium existence. An *equilibrium* consists of a distribution over investment and exit strategies (a, τ) and a reputation trajectory $\{x_t\}$ such that given $\{x_t\}$, all equilibrium strategies (a, τ) solve the firm's problem (2), and reputation $\{x_t\}$ is derived from the firm's strategy by Bayes' rule whenever possible.⁹ This definition allows for mixed strategies; Bayesian updating is more subtle with mixed beliefs since the market draws inferences from the fact that the firm has not exited. We spell out how to calculate the firm's reputation with mixed strategies and define mixed-strategy equilibrium in Appendix AD. In contrast to the case of the pure beliefs in Section IIB, reputation $\{x_t\}$ need not monotonically fall in t , meaning that Theorem 1 may not apply.¹⁰

Fixing a candidate strategy (a, τ) , the marginal value of self-esteem (7) and the first-order condition (8) are necessary for optimality. However, these equations are not sufficient since they correspond to checking only "one-step deviations on-path," i.e., checking the HJB equation (5) for all (t, z_t) rather than for all (t, z) . We must take this problem seriously since the firm's investment exhibits dynamic complementarity: investment today raises tomorrow's marginal benefit of self-esteem and thus optimal investment. Formally, this follows from the convexity of value V in z ; intuitively, today's investment raises the firm's time horizon and thereby its value from potential future breakthroughs. We thus do not use equations (7), (8) to construct an equilibrium, but rather establish equilibrium existence with an abstract fixed-point argument.

THEOREM 2: *An equilibrium exists.*

PROOF.

The best-response correspondence maps reputation $\{x_t\}$ to sets of optimal strategies (a^*, τ^*) . Bayes' rule maps distributions over strategies (a, τ) into reputation $\{x_t\}$. The Kakutani-Fan-Glicksberg Theorem then yields a fixed point. The key step in the proof is to define the appropriate weak topology that renders the strategy space compact and the two correspondences continuous. See Appendix AD for details. ■

⁹This definition does not impose sequential optimality of strategy (a, τ) and thus corresponds to Nash equilibrium rather than perfect Bayesian equilibrium. This is for notational convenience and is without loss. Indeed, the firm's investment is unobservable, so deviations do not affect beliefs and reputation. Thus, any equilibrium is outcome-equivalent to a perfect Bayesian equilibrium. In fact, all of the analysis in Section IIA starting at states t, z_t^* extends immediately to optimal strategies starting at any state t, z .

¹⁰To understand the non-monotonicity, assume that the firm mixes between two strategies: one with high investment $\{\bar{a}_t\}$ and a late exit time $\bar{\tau}$, and another with low $\{\underline{a}_t\}$ and an early exit time $\underline{\tau}$. When the firm fails to exit at $\underline{\tau}$, the market infers that the firm has been investing according to $\{\bar{a}_t\}$ and reputation jumps up. This particular mixed strategy cannot be an equilibrium since exit at time $\underline{\tau}$ is incompatible with the upward discontinuity of reputation. But for mixed strategies where the low $\{\underline{a}_t\}$ firm randomizes smoothly over exit times, reputation rises gradually after $\underline{\tau}$, which does not contradict equilibrium.

C. Steady-State Distribution

The reputational dynamics shape the distribution of reputation and income in the industry. Consider a continuum of price-taking firms, such as that studied above, and assume that new firms enter the market continuously at rate ϕ with reputation drawn with density h on $[x^e, 1]$, where $x^e = x_{\tau^*}$ is the reputation at the exit time. Writing $g(x) = g(a^*(x), x)$ for equilibrium drift, the density of firms $f(x, t)$ with reputation $x \in [x^e, 1]$ at time t is governed by the Kolmogorov forward equation¹¹

$$(12) \quad \frac{\partial}{\partial t} f(x, t) = -\frac{\partial}{\partial x} [f(x, t)g(x)] - \mu x f(x, t) + \phi h(x).$$

The second term captures the measure of firms experiencing a breakthrough, which takes their reputation from x to 1. The third term captures firms entering at reputation x . To understand the subtler first term, suppose for simplicity that the drift is constant and negative, $g(x) \equiv g < 0$, and that the density $f(x, t)$ rises in x . As more firms drift toward x from above than drift away below, the number of firms at x rises. The same argument obtains when the density is constant, $f(x, t) \equiv f(t)$, but the downward drift increases in x , $-g'(x) > 0$.

In steady state, the density of firms is constant, $f(x, t) \equiv f(x)$, and the measure of entering firms exactly compensates for the measure of exiting firms whose reputation drifts into the exit threshold, $\phi = -g(x^e)f(x^e)$. The total measure of firms depends on the choice of ϕ , which we normalize so that $\int_{x^e}^1 f(x) dx = 1$. Setting the LHS of (12) to 0 and rearranging, we get

$$(13) \quad f'(x) = \underbrace{\frac{g'(x)}{-g(x)} f(x)}_{\text{drift } (-/+)} + \underbrace{\frac{\mu x}{-g(x)} f(x)}_{\text{jumps } (+)} - \underbrace{\frac{\phi h(x)}{-g(x)}}_{\text{entry } (-)}.$$

To understand this equation, suppose as a benchmark that firms drift down at a constant rate $g(x) \equiv g < 0$, there are no breakthroughs, and all entry is at $x = 1$. Thus all 3 terms on the right-hand-side of (13) vanish, and so $f'(x) = 0$. Considering the three terms in sequence, the first captures the idea that firms accumulate where the downward drift is slowest, e.g., where the firm invests. The second term captures breakthroughs; these jumps prevent firms from reaching lower levels of reputation, and so $f'(x) > 0$. The third term captures entry; so long as firms enter with reputations below 1, the downward drift implies that firms accumulate at lower levels of reputation, and so $f'(x) < 0$.

D. Simulation

In this section we simulate an equilibrium (see Figure 1). This serves two purposes. First, it illustrates the empirical features that the model can generate. Second, we use it as a test case for policy experiments, proposing a method to decompose changes of average steady-state reputation into investment and selection effects.

¹¹ See, e.g., Kolmogorov (1931, eq. 179), Gardiner (2009, eq. 3.4.22), or Gabaix (2009, eq. 19), all of which include an additional Brownian term.

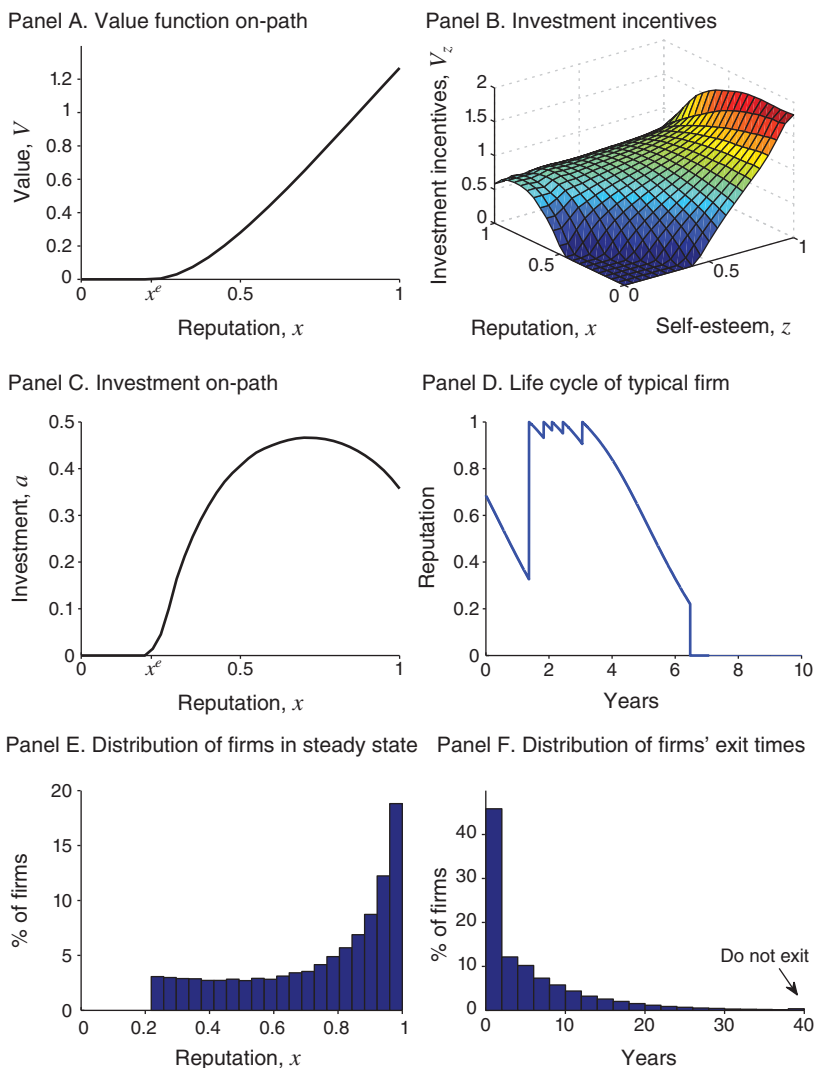


FIGURE 1. EQUILIBRIUM SIMULATION

Notes: Panels A and C show a firm's value and investment on-path, where $x_t = z_t$. Panel B shows investment incentives as a function of reputation x_t and self-esteem z_t . Panel D shows a typical firm life cycle. Panels E and F show the steady-state distribution of firms' reputations and exit times, assuming firms enter with reputation uniform on $[x^e, 1]$. We assume $\pi = x - k$, $k = 0.5$, $r = 0.2$, $c(a) = 0.2a^2$, $\bar{a} = 0.8$, $\lambda = 0.2$, and $\mu = 1$.

This simulation considers a restaurant that has revenue of $\$x$ million a year, operating costs of $k = \$500,000$, quadratic investment cost $c(a) = 0.2a^2$, an upper bound on investment $\bar{a} = 0.8$, 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 a high-quality restaurant is reviewed positively once a year, on average. Finally, we set $\lambda = 0.2$, so that technology shocks arrive every

5 years, on average. To make the figures easier to interpret, we adopt a Markovian perspective by replacing the firm's state variable t with its time- t reputation x_t .¹²

First, consider a single firm. Panel A illustrates the value function, which appears linear apart from the smooth pasting at the exit point, $x^e = 0.22$; a firm with perfect reputation has value $V = 1.27$. Panel B shows investment incentives V_z on the whole-state space (x, z) . One can see that V_z increases in z , illustrating the convexity of the value function. Panel C shows investment along the equilibrium path where $x = z$, which corresponds to the 45° line of Panel B; one can see that investment is single-peaked (as in Theorem 1). Panel D then shows the typical life cycle of a firm. Over the first two years its reputation falls despite the firm investing; shortly before it would have exited, a breakthrough boosts its reputation to one. Over the next two years, the firm continues to maintain its high quality and obtains a sequence of breakthroughs. Eventually, this string of successes comes to an end (perhaps because it switches to low quality, or perhaps because it is unlucky), and its reputation declines monotonically; ultimately, it exits $\tau^* = 3.4$ years after its last breakthrough.

We next consider the cross-sectional distribution of revenue and longevity, assuming firms enter with reputation equal to self-esteem and uniformly distributed on $[x^e, 1]$. Panel E illustrates the distribution of revenue, showing that there are lots of firms at the top, with reputation $x \approx 1$, and also slightly more firms at the exit point than in the middle. We can understand this using equation (13). The downward drift is relatively constant (see panel D), so the first term can be ignored; intuitively, the positive drift from investment $\lambda(a - x)$ and the negative drift from market learning $\mu x(1 - x)$ are both fastest in the middle and offset each other. The breakthroughs lead to lots of firms with reputation $x \approx 1$, while the entrants with low reputation raise the number of firms with reputation x^e . Panel F then illustrates the distribution of exit times: 45 percent of all firms fail in the first 2 years.¹³

Quality and reputation derive from two sources: investment by incumbent firms and replacement of exiting firms by new firms. Since learning does not affect reputation in expectation, equation (3) implies that the average reputation of incumbent firms changes at rate $\lambda(a - x)$, which is negative in steady state. This decline must be offset by replacement, which is the product of the replacement rate $\phi = -f(x^e)g(x^e)$ times the expected boost in reputation $(1 - x^e)/2$ when a new entrant with reputation $x \sim U[x^e, 1]$ replaces an exiting firm with reputation x^e . Rearranging, we can decompose average reputation into the investment and replacement effects

$$(14) \quad E[X] = E[A] + \frac{\phi}{\lambda} \frac{1 - x^e}{2},$$

where the expectation is taken with respect to the steady-state distribution f and $A = a(X)$. In our simulation, the average reputation $E[X]$ is 0.74, with

¹²To numerically solve for equilibrium, we start with a candidate investment strategy $\{a^0(x, z)\}$ and calculate reputation $\{x_t\}$ given correct beliefs and the resulting payoffs $\{V(x, z)\}$ via (2); we then derive optimal investment $\{a^1(t, z)\}$ via the first-order condition (8). We iterate this process until it converges; the fixed point satisfies the HJB equation $(r + \mu z)V(x, z) = \sup_a [\pi(x_t) - c(a) + g(\bar{a}, x)V_x(x, z) + g(a, z)V_z(x, z) + \mu zV(1, 1)]^+$ and has correct beliefs: $\bar{a} = a$. Hence, it is an equilibrium.

¹³Coincidentally, Parsa et al. (2005) report that 40–50 percent of restaurants fail over their first 2 years.

52 percent coming from investment and 48 percent coming from replacement. Breaking down the latter term, the replacement rate equals $\phi = 0.18$, implying an approximate life expectancy of $1/\phi = 5.6$ years,¹⁴ while the mean jump from replacement is 0.39.

E. Applications

Our model can help evaluate the impact of disclosure policies that aim to improve information in the marketplace. In the United States, these policies started with the 1906 Pure Food and Drug Act and the 1933 Securities Exchange Act and have become increasingly popular as a form of “light touch” regulation.¹⁵ For example, Jin and Leslie (2003, 2009) show that the introduction of restaurant health grades led to an increase in the quality of restaurants, especially for independent establishments. Similarly, Kolstad (2013) shows that the introduction of surgeon report cards prompted surgeons to improve the quality of their services. The model can also be used to assess the impact of third-party information providers that have substantial impacts on their respective industries, such as Moody’s (bonds), TripAdvisor (hotels), Edmunds (cars), Zillow (houses), and Greatschools (schools). For example, Luca (2016) shows that the introduction of Yelp pushed demand toward good restaurants and affected the selection of restaurants in the market.

Our paper provides a method to evaluate the impact of such changes. As a test case, recall the simulation in Section IID and consider a 20 percent increase in the frequency of breakthroughs μ from 1 to 1.2. The increased monitoring reduces moral hazard, raising maximum firm value V from 1.27 to 1.46; the higher option value then lowers the exit threshold x^e from 0.22 to 0.18. Surprisingly, average reputation $E[X]$ rises only slightly, from 0.74 to 0.77. But the decomposition (14) tells a different story. The faster learning boosts investment $E[A]$ from 0.38 to 0.48, while the exit rate ϕ drops from 0.18 to 0.14. Investment now accounts for 63 percent of the mean reputation, whereas replacement accounts for 37 percent (compared to 52–48 percent before). In addition to showing how faster learning raises the importance of investment over replacement, it provides a cautionary tale about inferring investment from the steady-state reputation; we discuss how one can identify the model in Section IV.

The model is also useful to analyze other policies. For example, Luca and Luca (2018) find that an increase in the minimum wage raises the exit rates of restaurants with low ratings on Yelp, but has no effect on highly rated restaurants. Higher minimum wages correspond to an increase in operating costs in our model, which has a direct effect on selection by raising the exit threshold and also an indirect effect on investment by lowering profits; the reduction in believed investment then hastens the decline in reputation between breakthroughs and further shortens life expectancy. For example, a 20 percent rise in operating cost k —from 0.5 to 0.6—lowers the

¹⁴This $1/\phi$ formula is approximate, as it assumes that exit is i.i.d. along a firm’s life cycle, but it will be useful to understand the effect of policy changes.

¹⁵Fung, Graham, and Weil (2007) discuss over 130 rules that were introduced by the federal government over 1995–2005, such as policies about disclosing the likelihood of a car rolling over or the presence of toxic substances in workplaces.

maximum firm value V from 1.27 to 0.68 and raises the exit threshold x^e from 0.22 to 0.36. Again, the mean steady-state reputation is about the same (0.75 before and after), but the mean investment is much smaller (0.19 versus 0.38) and the exit rate is substantially higher (0.34 versus 0.18). Investment now accounts for 25 percent of mean reputation, whereas replacement accounts for 75 percent (compared to 52–48 percent in the benchmark).

III. Two Model Variants

In this section, we consider two natural variants of our baseline model; these help us understand the two economic forces underlying the single-peaked incentives seen in Theorem 1. In Section IIIA, we suppose the market observes the firm's investment. Thus, both the firm and the market symmetrically learn about the firm's quality, but there is no moral hazard. We show that investment *decreases* in the time since a breakthrough, as the firm approaches bankruptcy. In Section IIIB, we suppose instead the firm knows its own quality in addition to its investment, while the market knows neither. We show that investment *increases* in the time since a breakthrough, as the benefit from a breakthrough grows.

A. Consumers Observe Firm's Investment

Consider the baseline model and suppose that the market observes both 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. As in equation (2), we truncate the firm's flow payoffs at a breakthrough, yielding

$$(15) \quad \hat{V}(z_t) = \sup_{a, \tau} \int_t^\tau e^{-\int_t^s (r + \mu z_u) du} [\pi(z_s) - c(a_s) + \mu z_s \hat{V}(1)] ds.$$

Since the firm controls both self-esteem and reputation, the analysis of the equilibrium reduces to a decision problem. Write $(\hat{a}, \hat{\tau})$ for the optimal strategy and $\hat{z} = \{\hat{z}_t\}$ for the associated self-esteem.

The analysis follows Section II. The value function is strictly convex with derivative¹⁶

$$(16) \quad \hat{V}'(z_t) = \hat{\Gamma}(t) := \int_t^{\hat{\tau}} e^{-\int_t^s (r + \lambda + \mu(1 - \hat{z}_u)) du} [\pi'(z_s) + \mu(\hat{V}(1) - \hat{V}(\hat{z}_s))] ds.$$

In any optimal strategy, investment satisfies

$$\lambda \hat{\Gamma}(t) = c'(\hat{a}_t)$$

with $\hat{a}_t = 0$ if $\lambda \hat{\Gamma}(t) < c'(0)$ and $\hat{a}_t = \bar{a}$ if $\lambda \hat{\Gamma}(t) > c'(\bar{a})$.

¹⁶ See Appendix BA for a proof. The subsequent analysis implies that V is indeed differentiable.

THEOREM 3: *Investment \hat{a}_t decreases in the time since a breakthrough; at the exit threshold, $\hat{a}_{\hat{\tau}} = 0$. Moreover, the optimal exit time satisfies*

$$(17) \quad \pi(\hat{z}_{\hat{\tau}}) + \mu \hat{z}_{\hat{\tau}} \hat{V}(1) = 0.$$

PROOF:

Given assumption (4), 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, $\hat{\Gamma}(t) = \hat{V}(z_t)$ strictly decreases in t . ■

Intuitively, as the firm gets closer to the exit time, any investment pays off over a shorter horizon, reducing incentives. When compared to the baseline model with moral hazard, investment directly raises reputation, and hence income, in addition to raising self-esteem and the chance of future breakthroughs. This additional direct benefit of investment is captured by the additional “direct dividend” of $\pi'(z)$ in (16), which is absent in (7). In the baseline model, investment initially increases over time after a breakthrough as the reputational dividends grow. Here we see incentives monotonically decrease over time as the loss of direct dividends outweighs the growth of reputational dividends. This intuition also suggests that the elimination of moral hazard raises the level of investment and the exit time relative to the baseline model. In Appendix BB, we verify this intuition when $\pi(x) = x - k$ and k is sufficiently large.

To illustrate the optimal strategy, we simulate it for the same parameters as the baseline model.¹⁷ The elimination of moral hazard raises the firm’s maximum value from 1.27 to 1.38 while lowering the exit threshold from 0.22 to 0.21. Panel A of Figure 2 shows that investment is increasing in reputation (see Theorem 3) and higher than investment in the baseline model (Figure 1, panel C) at every level of reputation. Together with the lower exit threshold, the higher investment lowers the downward drift and extends the time to exit τ from 3.4 to 5.4 years. Panel B shows the steady-state distribution of revenue: the frequent breakthroughs and low downward reputational drift at the top skew the distribution of firms toward high levels of reputation. There is little entry, so in contrast to the baseline model, the density monotonically decreases as we approach the exit point.

We can understand the quantitative differences with the baseline model using the reputation decomposition (14). Average steady-state reputation is 0.87, with 89 percent coming from investment and 11 percent from replacement. Breaking down the replacement term, the exit rate drops to 0.04 (implying an approximate life expectancy of 22 years) while the average jump from replacement is 0.40. All told, the elimination of moral hazard substantially increases investment and also reduces the exit rate, both of which raise the importance of investment over replacement.

¹⁷To simulate the equilibrium, we guess a post-breakthrough value $\hat{V}(1)$ and use the exit condition (17) to derive x^e . We then use the HJB equation to calculate values at each reputation, yielding a post-breakthrough value $\hat{V}(1)$. In equilibrium, $\hat{V}(1)$ and $\bar{V}(1)$ coincide.




FIGURE 2. CONSUMERS OBSERVE FIRMS' INVESTMENT

Notes: Panel A shows a firm's investment. Panel B shows the steady-state distribution of firms' reputations, assuming firms enter with reputation uniform on $[\hat{x}^e, 1]$. Parameters are the same as Figure 1.

These forces also account for the higher firm value, with the increase in life expectancy being more important for the average firm.¹⁸

B. Firm Knows Its Own Quality

Consider the baseline model and suppose the firm observes its own quality. As before, the market learns from public breakthroughs and maintains beliefs about the firm's investment and quality. Crucially, the firm's exit decision is now a signal of its quality, since high-quality firms are more valuable and stay in business longer.

We focus on strategies that depend on the time since the last breakthrough, t , and on current quality. Formally, such a strategy consists of an investment plan $a^\theta = \{a_t^\theta\}$ and an exit time τ^θ for $\theta = L, H$. To analyze firm value for an arbitrary trajectory of reputation $\{x_t\}$, we truncate its cash flow expansion at the first technology shock, obtaining

(18)

$$\begin{aligned} \check{V}(t, \theta) = \sup_{a^\theta, \tau^\theta} \int_t^{\tau^\theta} e^{-(r+\lambda)(s-t)} & \left[\pi(x_s) - c(a_s^\theta) + \lambda \left(a_s^\theta \check{V}(s, H) + (1 - a_s^\theta) \check{V}(s, L) \right) \right. \\ & \left. + \mu \theta (\check{V}(0, H) - \check{V}(s, H)) \right] ds, \end{aligned}$$

where the last term captures the value of breakthroughs with present value $\check{V}(0, 1) - \check{V}(s, 1)$ and arrival rate $\mu\theta$.

¹⁸As a back-of-the-envelope calculation, in the benchmark model the average value is approximately $E[x - k - c(a)] / (r + \phi) = 0.21 / 0.38 = 0.55$, where the expectation is the mean steady-state value. With observed investment, we have $E[x - k - c(a)] / (r + \phi) = 0.25 / 0.24 = 1.02$. Thus, the majority of the increase in the average firm value comes from the increase in life expectancy.

Writing $\check{\Gamma}(t) = \check{V}(t, H) - \check{V}(t, L)$ for the *value of quality*, optimal investment \check{a}_t^θ is thus characterized by the first-order condition

$$(19) \quad \lambda \check{\Gamma}(t) = c'(\check{a}_t^\theta),$$

with $\check{a}_t^\theta = 0$ if $\lambda \check{\Gamma}(t) < c'(0)$ and $\check{a}_t^\theta = \bar{a}$ if $\lambda \check{\Gamma}(t) > c'(\bar{a})$. Importantly, optimal investment \check{a}_t^θ is independent of the firm's quality, allowing us to drop the superscript θ . Intuitively, investment only pays off if there is a technology shock, in which case the firm's current quality is irrelevant.

Equilibrium strategies $(\check{a}, \check{\tau}^\theta)$ and reputation $\{x_t\}$ are defined as in Section IIB. We restrict attention to equilibria where $\{x_t\}$ is continuous and weakly decreasing.¹⁹

THEOREM 4: *There exists an equilibrium with continuous and weakly decreasing reputation $\{x_t\}$ and pure investment $\{\check{a}_t\}$. In any such equilibrium, the exit time of the low-quality firm $\check{\tau}^L$ has support $[\underline{\tau}, \infty)$ for some $\underline{\tau} > 0$. Reputation and firm value are constant for $t \in [\underline{\tau}, \infty)$ and satisfy*

$$(20) \quad \pi(x_t) + \max_a [a \lambda \check{V}(t, H) - c(a)] = 0.$$

The high-quality firm never exits; i.e., $\check{\tau}^H = \infty$. Investment \check{a}_t increases over $[0, \underline{\tau}]$ and remains constant thereafter.

PROOF:

Exit behavior follows from Bar-Isaac (2003, Proposition 2). Low-quality firms start to exit at some time $\underline{\tau}$: if all low-quality firms exited at $\underline{\tau}$, reputation would jump to 1, undermining equilibrium. Hence, low-quality firms randomize, exiting at a constant rate ψ ; this places a lower bound on reputation, allowing high-quality firms to stay in the market. See Appendix BC for details. Existence follows by the Kakutani-Fan-Glicksberg fixed-point theorem, as shown in Appendix BD.

We now show that investment increases over time. Since a low-quality firm is indifferent at the exit threshold, it can achieve its value by remaining in the market indefinitely, thereby following the same strategy as the high-quality firm. Subtracting the value of high-value and low-value firms (18), we obtain the following expression for the equilibrium value of quality:

$$(21) \quad \check{\Gamma}(t) = \check{V}(t, H) - \check{V}(t, L) = \int_t^\infty e^{-(r+\lambda)(s-t)} \mu [\check{V}(0, H) - \check{V}(s, H)] ds.$$

The integrand in (21) represents the *reputational dividend of quality*: high quality 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 . As $s \in [0, \underline{\tau}]$ rises, the firm's value $\check{V}(s, 1)$ falls and reputational

¹⁹ As in footnote 3, the substantial part of this restriction is to rule out downward jumps in reputation. Indeed, for any t , there exists an equilibrium where the firm exits at time t and failure to exit is punished by off-path market beliefs that quality is low, dropping reputation to 0 and justifying the firm's exit. We ignore such equilibria because it is implausible for the market to interpret failure to exit as a signal of low quality.

dividends $\check{V}(0, 1) - \check{V}(s, 1)$ grow. Hence, an increase in t leads to an increase in the value of quality (21) and in investment via the first-order condition (19). ■

Intuitively, breakthroughs are most valuable to a firm with low reputation since a breakthrough takes the firm from its current reputation to $x = 1$. These increasing investment incentives are in sharp contrast to the single-peaked, eventually vanishing investment incentives in Theorem 1. In the baseline model, the firm gives up near the exit threshold and coasts into bankruptcy; with privately known quality, the firm fights until the bitter end. Recall that with unknown quality, the firm's investment at times $t \in [\tau^* - dt, \tau^*]$ pays off only if a technology shock arrives and a breakthrough arrives that averts exit. The probability of this joint event is of order dt^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 the boost in quality is immediately observed by the firm, averting exit. Thus, investment incentives are of order dt at all times and are actually maximized when the firm is about to exit.²⁰

To illustrate the equilibrium, Figure 3 simulates the model for the same parameters as in Figure 1.²¹ Panel A shows the value functions for the high-quality and low-quality firms; these are only defined on $[\check{x}^e, 1]$, since reputation never falls below $\check{x}^e = x_{\tau}$. Interestingly, the value after a breakthrough is lower in the baseline model (1.14 versus 1.27), while the average steady-state value is somewhat higher (0.77 versus 0.55).²² Intuitively, high-reputation firms are better off in the baseline model since investment is high and they are a long way from the exit threshold, whereas the average firm is better off in the known-quality model since the signaling effect of remaining in the market props up firm value. Panel B shows that investment falls in reputation (see Theorem 4), and that investment at a given reputation is a little lower than in the baseline model. Panel C shows a typical life cycle. The firm's reputation quickly hits the exit threshold ($\check{x}^e = 0.47$), where reputation is constant, as the quitting of low-quality firms offsets the negative inference from lack of a breakthrough; it then obtains a sequence of breakthroughs before its reputation once again declines, and it eventually exits. At the threshold, low-quality firms exit at rate $\psi = 0.49$. From the market's perspective, firms at \check{x}^e exit at rate $(1 - \check{x}^e)\psi = 0.26$ and receive breakthroughs at rate $\check{x}^e\mu = 0.47$ and so remain at the exit threshold for 1.37 years. The overall exit rate multiplies the conditional exit rate ψ with the number of low firms at the exit

²⁰One may wonder how to bridge these starkly different predictions about investment behavior at the exit threshold. In Appendix BE, we propose a model with imperfect private information that includes known and unknown quality as extreme cases. As with unknown quality, we argue that as long as private quality information is imperfect, investment vanishes at the exit threshold. However, the "invest until the bitter end" insight of the known-quality variant is also robust in the sense that as the private information becomes perfect, the time interval over which investment vanishes shrinks to zero.

²¹To simulate the equilibrium, we guess a post-breakthrough value $\check{V}(0, H)$ and use the HJB equation and the exit condition (20) to derive the exit threshold and the values at x^e . We then use the HJB equation to calculate values at each reputation, yielding a post-breakthrough value $\check{V}(0, H)$. In equilibrium, $\check{V}(0, H)$ and $\bar{V}(0, H)$ coincide.

²²These average steady-state numbers come from the following approximation. In the baseline model, the average value is roughly $E[x - k - c(a)] / (r + \phi) = 0.21 / 0.38 = 0.55$. With known quality, we have $E[x - k - c(a)] / (r + \phi) = 0.23 / 0.29 = 0.77$, so higher average values primarily result from the lower exit rate.

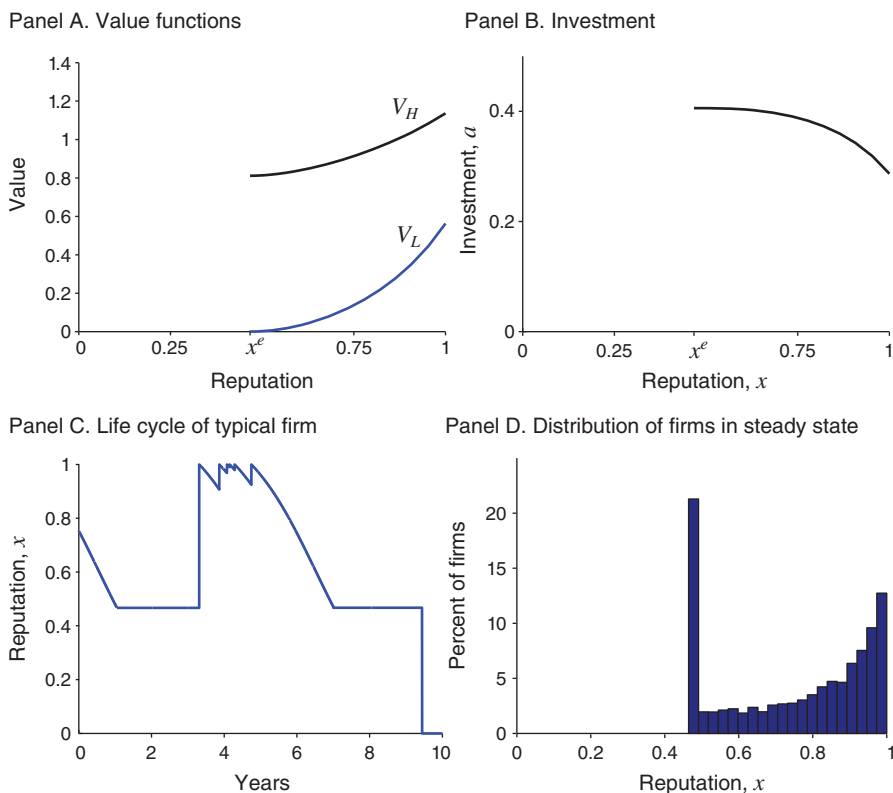


FIGURE 3. FIRM KNOWS ITS QUALITY

Notes: Panels A and B show a firm's value and investment. Panel C shows a typical firm life cycle. Panel D shows the steady-state distribution of firms' reputations, assuming firms enter with reputation uniform on $[\check{x}^e, 1]$. Parameters are the same as Figure 1.

point, $\phi = (1 - \check{x}^e)\psi F(\check{x}^e) = 0.10$, implying an approximate life expectancy of 10 years. Finally, panel D shows the steady-state distribution of firms' reputations. Equation (13) implies that density is decreasing in reputation over $(\check{x}^e, 1]$ as a result of the breakthroughs; the 0 drift at \check{x}^e then results in a mass of firms at the exit threshold.

We can understand the quantitative differences with the benchmark model by decomposing average reputation. Relative to equation (14), there is now a third, signaling term supporting the firm's reputation. Since low-quality firms at the exit threshold exit at rate ψ , failure to exit is good news; this induces an upward drift of $\psi\check{x}^e(1 - \check{x}^e)$ on the $F(\check{x}^e)$ firms at the cutoff. All told,

$$E[X] = E[A] + \frac{\phi}{\lambda} \frac{1 - \check{x}^e}{2} + \frac{\psi}{\lambda} \check{x}^e (1 - \check{x}^e) F(\check{x}^e).$$

In Figure 3, the average steady-state reputation is 0.75, with 49 percent coming from investment, 17 percent from replacement, and 34 percent coming from signaling.

IV. Discussion

We have proposed a model in which firms make optimal investment and exit decisions while the market learns about the quality of the firm's product. We characterize investment incentives and show they are single-peaked in the firm's reputation. This yields predictions about the distribution of firms' reputations and the turnover rate. The model follows the spirit of Ericson and Pakes (1995), and we hope that it can be used as a framework for empirical work. For example, it could be used to study the rise and fall of new restaurants or the incentives for surgeons to invest in their skills. Here we briefly discuss how one might identify the model and two important extensions: competition and Brownian learning.

Identification: To illustrate how to identify the model parameters, suppose we observed data generated by the baseline model with $\pi(x) = x - k$. Given a time series of x_t (i.e., reputation or revenue data), one can identify μ from the frequency of jumps, λ from the drift at x^e , and on-path investment $a(x)$ via the drift of x_t elsewhere. Calibrating the interest rate, one can then uniquely identify the operating cost k from the exit condition (9) and the investment cost function $c(\cdot)$ from the first-order condition (8). See Appendix BF for details. Alternatively, one could use accounting data to shed light on costs or aggregate data like the distribution of firms' reputations and the distribution of exit rates.

Competition: Our firm operates in isolation, but the analysis extends to a competitive market that is in steady state, such as the restaurant industry. There are several ways of modeling this. In Atkeson, Hellwig, and Ordoñez (2014), agents consume a little of each good. Income is then $\pi(x) = px_t - k$, where the price $p = P(\int x_{i,t} di)$ is determined by aggregate reputation. One can then model entry by having an exogenous inflow or letting potential entrants pay a cost and draw an initial reputation. In Vial and Zurita (2017) there is mass one of agents, each of whom consumes the output of one firm. In equilibrium, mass one of firms operate with exiting firms balancing the flow of new entrants who, for example, arrive exogenously with random reputation. Income is then $\pi(x) = p_0 + x - k$, where p_0 is determined by the zero-profit condition of the exiting firm. Alternatively, following Dixit and Stiglitz (1977), one could embed our reputation model into monopolistic competition, where firm i 's reduced-form flow profits $\pi_i(x_i, x_{-i})$ depend on its own reputation x_i and the (stationary) distribution of its competitor's reputations x_{-i} . While competition does not affect the shape of incentives (i.e., Theorem 1), it can change the impact of policies via the endogenous prices. For example, a disclosure policy that would previously raise average reputation and reduce exit would lower prices, raise the exit threshold, and attract entry.

Alternatively, one can explicitly consider strategic interaction in the product market via a model of oligopoly. For example, consider a duopoly where a firm's flow-profits $\pi(x, y)$ depend on its competitor's reputation y as well as own reputation x . Such reduced-form profits might be derived from static Bertrand competition via $\pi(x, y) = (x - y)^+$ or a logit demand model (e.g., Anderson, De Palma, and Thisse 1992). Adopting a Markovian perspective, assume that firms are

believed to follow symmetric strategies $\tilde{a}(x, y)$. Firm value is then a function of both firms' reputations and its own self-esteem and is thus governed by the following HJB equation:

$$\begin{aligned} rV(x, y, z) = & \max_a \left[\left[\pi(x, y) - k - c(a) \right] + \mu y \left[V(x, 1, z) - V(x, y, z) \right] \right. \\ & + \mu z \left[V(1, y, 1) - V(x, y, z) \right] + g(\tilde{a}(x, y), x) V_x(x, y, z) \\ & \left. + g(\tilde{a}(y, x), y) V_y(x, y, z) + g(a, z) V_z(x, y, z) \right]^+. \end{aligned}$$

Thus, the firm's optimal investment $a(x, y, z)$ satisfies the usual first-order condition, $c'(a(x, y, z)) = \lambda V_z(x, y, z)$, while beliefs in a (pure strategy) equilibrium are given by $\tilde{a}(x, y) = a(x, y, x)$. The qualitative nature of investment incentives derived for the monopoly model generalize to this oligopoly model. Specifically, one can show that a firm's incentives increase after a breakthrough—since an immediate, additional breakthrough has no value—and vanish when reputation is close to the exit boundary and the firm's life span is short.

Information Structure: Our paper assumes that the market learns via a perfect good news process, which is particularly tractable: it allows us to prove equilibrium existence and delivers robust intuitions about how investment varies with reputation (see below). It also provides reasonable predictions about some aggregate distributions (e.g., the distribution of firms' Yelp scores and exit rates). However, revenue data with a right-skewed distribution or continuous time paths might call for more continuous models of learning. Such models are easy to simulate and have similar qualitative properties to the perfect good news model.

To be specific, consider a model variant where instead of the breakthroughs, the market and firm observe Brownian signals of quality $d\xi_t = \mu_B \theta_t + dW_t$; otherwise, the model is as in Section I. A firm's reputation and self-esteem evolve according to

$$\begin{aligned} dx_t &= g(a_t, x_t, z_t) dt + \sigma(x_t) dW_t, \\ dz_t &= g(a_t, z_t, z_t) dt + \sigma(z_t) dW_t, \end{aligned}$$

with drift $g(a, x, z) = \lambda(a - x) + \mu_B x(1 - x)(z - x)$ and volatility $\sigma(x) = \mu_B x(1 - x)$. Given Markovian beliefs $\tilde{a}(x)$, the value function is determined by the HJB equation

$$\begin{aligned} rV(x, z) = & \max_a \left[\left[\pi(x) - k - c(a) \right] + g(\tilde{a}, x, z) V_x(x, z) + g(a, z, z) V_z(x, z) \right. \\ & \left. + \frac{1}{2} \sigma(x)^2 V_{xx}(x, z) + \frac{1}{2} \sigma(z)^2 V_{zz}(x, z) + \sigma(x) \sigma(z) V_{xz}(x, z) \right]^+. \end{aligned}$$

Again, the firm's optimal investment $a(x, z)$ satisfies the usual first-order condition, $c'(a(x, z)) = \lambda V_z(x, z)$, while beliefs in a (pure strategy) equilibrium are given by $\tilde{a}(x) = a(x, x)$. The qualitative nature of investment incentives derived for the

perfect good news model generalize to this Brownian learning process. Specifically, one can show that a firm's incentives are decreasing in x when its reputation is high, $x \approx 1$, and vanish when its expected life span is short, $x \approx x^e$.²³

Alternatively, one could assume a Poisson signal structure with imperfectly revealing signals. A Poisson learning process with imperfect good news is qualitatively similar to perfect good news or Brownian learning. By contrast, Poisson learning with bad news may exhibit investment at low reputations. Intuitively, reputational drift at the exit cutoff may be positive (in the absence of a signal), meaning that life expectancy does not vanish. Ultimately, the appropriate model depends on the news structure that is generating the data and the policies one wishes to investigate. For example, a Poisson model allows one to study the effect of giving awards to the best surgeon (adding good news signals) or flagging the worst (adding bad news signals).

APPENDIX A. MATHEMATICAL APPENDIX: PROOFS FROM SECTION III

A. Monotonicity of Value Function in Lemma 1

Here we adapt arguments from Board and Meyer-ter-Vehn (2013) to show that for strictly decreasing reputation $\{x_t\}$, firm value $V(t, z)$ strictly decreases in t and strictly increases in z on $\{(t, z) : V(t, z) > 0\}$.

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 firm's 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.²⁴ Given $\{t_i, \zeta_i\}$ and $(\{A_\sigma^*\}, T^*)$, we can compute the reputations 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

$$(A1) \quad X_\sigma \leq X'_\sigma \quad \text{and} \quad Z_\sigma \leq Z'_\sigma$$

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 (A1) for $\sigma \in [0, t_1)$. At $\sigma = t_1$, the low (high) firm experiences a breakthrough if $\zeta_1 \leq Z_{\sigma-}$ ($\zeta_1 \leq Z'_{\sigma-}$).

²³The key step in this argument is to generalize the representation of investment incentives in terms of reputational dividends (7) to these learning processes. Just as with perfect good news learning, at the top, $x = 1$, the reputational dividend vanishes, and so investment falls; at the bottom, $x = x^e$, the time to exit τ^e vanishes, and so, too, does investment.

²⁴Note that, unlike everywhere else in this paper, this strategy for the high firm is not a solely a function of the time since its last breakthrough. Rather, it is a function of the time since the low firm's last breakthrough, while the high firm may have experienced additional breakthroughs in the meantime due to its high quality.

As $Z_{\sigma-} \leq Z'_{\sigma-}$, we get (A1) for $\sigma = t_1$. Inductive application of these steps yields (A1) 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 profits $\pi(X'_\sigma) - c(A^*_\sigma)$ at all times σ , implying $V(t', z') \geq V(t, z)$. As long as firm value is strictly positive and the firms do not exit immediately, the inequality $X_\sigma \leq X'_\sigma$ is strict for a positive measure of times with positive probability, implying $V(t', z') > V(t, z)$.

B. 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

$$(A2) \quad \Pi(t, z_t) = \int_{s=t}^{\tau} e^{-\int_t^s (r + \mu z_u) du} [\pi(x_s) - c(a_s) + \mu z_s \Pi(0, 1)] ds$$

be the firm's continuation value, where the integral of the cash flows is truncated at the first breakthrough as in (2). We will show that $\Pi(t, z)$ is differentiable in z with derivative

$$(A3) \quad \Pi_z(t, z_t) = \int_{s=t}^{\tau} e^{-\int_t^s (r + \lambda + \mu(1 - z_u)) du} \mu [\Pi(0, 1) - \Pi(s, z_s)] ds.$$

Equation (7) then follows, by the envelope theorem, Milgrom and Segal (2002, Theorem 1). To show (A3), we first recall two facts from Board and Meyer-ter-Vehn (2013).

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

$$(A4) \quad \psi(t) = \int_t^{\tau} e^{-\int_t^s \rho(u) du} \phi(s) ds$$

is the unique solution to the integral equation

$$(A5) \quad f(t) = \int_{s=t}^{\tau} (\phi(s) - \rho(s)f(s)) ds.$$

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).

PROOF OF LEMMA 2:

Setting $\psi(s) = e^{-r(s-t)} \Pi(s, z_s)$, $\rho(s) = \mu z_s$, and $\phi(s) = e^{-r(s-t)} (x_s - c(a_s) - k + \mu z_s \Pi(0, 1))$, equation (A2) becomes (A4). Applying Claim 1, we get (A5), which becomes

$$\Pi(t, z_t) = \int_{s=t}^{\tau} e^{-r(s-t)} [\pi(x_s) - c(a_s) + \mu z_s (\Pi(0, 1) - \Pi(s, z_s))] ds.$$

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 (A5). Applying Claim 1, we get (A4), 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 (A3). ■

C. Differentiability in Proof of Theorem 1

This Appendix relaxes the assumption that value functions are differentiable in the proof of Theorem 1. We first establish that whenever the partial derivative $V_t(t, z_t^*)$ exists, it is equal to

$$(A6) \quad \Psi(t) := \int_t^{\tau^*} e^{-\int_t^s (r + \mu z_u^*) du} d\pi(x_s).$$

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

Rewrite the firm's continuation value (2) by writing $\sigma = s - t$ for the time since t and $(\{a_\sigma^*\}, \zeta^*)$ for the optimal strategy starting at t . Then,

$$V(t, z_t^*) = \int_{\sigma=0}^{\zeta^*} e^{-\int_0^\sigma (r + \mu z_{t+\sigma}^*) du} [\pi(x_{t+\sigma}) - c(a_\sigma^*) + \mu z_{t+\sigma}^* V(0, 1)] d\sigma.$$

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 (A6). As $\pi' > 0$ and $\{x_t\}$ strictly decreases, $d\pi(x_s) < 0$ and hence $\Psi(t)$ must be negative.

Next, the discount rate $\rho(t) = r + \lambda + \mu(1 - z_t^*)$ is Lipschitz continuous, with derivative μz_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 $(d/dt)V(t, z_t^*) = \Psi(t) + \dot{z}_t^* \Gamma(t)$ for almost all t .

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

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

Since $\Psi(s) < 0$, $\dot{\Gamma}(t) \leq 0$ and $\dot{\Gamma}(s)$ and $\rho(s)$ are continuous, the integrand is strictly negative for small ε , so $\dot{\Gamma}$ strictly decreases on some small interval $[t, t + \varepsilon]$. If Γ 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.

D. Proof of Theorem 2

Preliminaries: We first define mixed strategies. Write mixed beliefs over the firm's investment $\tilde{a} = \{\tilde{a}_t\}$ and exit time $\tilde{\tau}$ as $F = F(\tilde{a}, \tilde{\tau})$. Let $\tau(F) := \min\{t : F(\tilde{\tau} \leq t) = 1\}$ be the first time at which the market expects the firm to exit with certainty. Writing E^F for expectations under F , the firm's reputation is given by $x_t = E^F[\theta_t | h^t, \tilde{\tau} > t]$ for all $t < \tau(F)$. Since breakthroughs arrive with intensity $\mu z_s(\tilde{a})$, 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

$$(A7) \quad x_t = \frac{E^F[z_t(\tilde{a}) w_t(\tilde{a}) \mathbf{1}_{\{t < \tilde{\tau}\}}]}{E^F[w_t(\tilde{a}) \mathbf{1}_{\{t < \tilde{\tau}\}}]}, \quad \text{for all } t < \tau(F).$$

When the firm fails to exit—that is, at times $t \geq \tau(F)$ —the market revises its beliefs about quality x_t in an arbitrary, measurable manner consistent with the timing of the last breakthrough

$$(A8) \quad \underline{x}_t \leq x_t \leq \bar{x}_t, \quad \text{for all } t \geq \tau(F),$$

where \underline{x}_t solves the law of motion for zero investment $\dot{x} = g(0, x)$ and \bar{x}_t for full investment $\dot{x} = g(\bar{a}, x)$. The restriction (A8) prevents the firm from signaling what it doesn't know to the market.

An *equilibrium* consists of a distribution over investment and exit strategies $F = F(a, \tau)$ and a reputation trajectory $x = \{x_t\}$ such that (i) given $\{x_t\}$, any strategy (a, τ) in the support of F solves the firm's problem (2), and (ii) reputation $\{x_t\}$ is derived from F by Bayes' rule via (A7) for $t < \tau(F)$ and satisfies (A8) for $t \geq \tau(F)$.

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

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

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

$$BR(x) = \operatorname{argmax}_{a, \tau} \Pi(a, \tau; x)$$

and the Bayesian updating formula \mathcal{B} defined by (A7)–(A8) 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 Topological Spaces: By (4), the firm's optimal exit time is bounded above by some finite $\bar{\tau}$, allowing us to truncate the domain of all pertinent functions at $\bar{\tau}$. So motivated, embed the space B of measurable, bounded functions $[0, \bar{\tau}] \rightarrow [0, 1]$ in the (rescaled) unit ball of $L^2([0, \bar{\tau}], \mathbb{R})$. We interpret both investment $\{a_t\}$ and reputation $\{x_t\}$ as elements of B . 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 $\int_0^{\bar{\tau}} (a_t^n - a_t) \xi_t dt \rightarrow 0$ for all test functions $\xi \in L^2([0, \bar{\tau}], \mathbb{R})$. This topology is coarse enough to make B compact but also fine enough for the trajectory $\{z_t\}_{t \in [0, \bar{\tau}]}$ to be continuous (in the sup-norm) in $\{a_t\}_{t \in [0, \bar{\tau}]}$ (see Theorem 43.5 of Davis 1993). Thus, firm value (2) is continuous in $(\{a_t\}, \{x_t\}, \tau)$ and is maximized by some $(\{a_t^*\}, \tau^*)$. As for the firm's mixed strategies, we equip $\Delta(B \times [0, \bar{\tau}])$ with the topology of convergence in distribution. Standard arguments (e.g., Theorem 14.11 of Aliprantis and Border 1999) show that this space is compact and locally convex.

Upper Hemicontinuity of Bayes' Rule: We now prove that the correspondence $\mathcal{B}: \Delta(B \times [0, \bar{\tau}]) \rightarrow B$ mapping beliefs F to the set of measurable trajectories $\{x_t\}$ that satisfy (A7)–(A8) is upper hemicontinuous. Consider a sequence of beliefs F^n (with expectation E^n) that converges to F in distribution. Here, $\mathcal{B}(F^n)$ consists of all measurable trajectories $\{x_t^n\}$ that satisfy (A7) for $t < \tau(F^n)$ (when replacing E^F by E^n) and (A8) for $t \geq \tau(F^n)$. As F assigns probability less than one to the event $\{\bar{\tau} < t\}$ for any $t < \tau(F)$, so does F^n for sufficiently large n ; thus, $\lim_{n \rightarrow \infty} \tau(F^n) \geq \tau(F)$.

We now show that $x_t^n \rightarrow x_t$ for all $t < \tau(F)$ at which the marginal distribution $F(\bar{\tau})$ is continuous.²⁵ Consider the numerator of (A7). (The argument for the denominator is identical.) The integrand $\chi_t^-(\tilde{a}, \tilde{\tau}) := z_t(\tilde{a}) w_t(\tilde{a}) \mathbf{1}_{\{\tilde{\tau} > t\}}$ is continuous in \tilde{a} (see Theorem 43.5 of Davis 1993) and lower semicontinuous

²⁵ It can have at most countably many discontinuities.

in $\tilde{\tau}$; similarly, $\chi_t^+(\tilde{a}, \tilde{\tau}) := z_t(\tilde{a}) w_t(\tilde{a}) \mathbf{1}_{\{\tilde{\tau} \geq t\}}$ is continuous in \tilde{a} and upper semicontinuous in $\tilde{\tau}$. The portmanteau theorem thus implies $\liminf E^n[\chi_t^-(\tilde{a}, \tilde{\tau})] \geq E^F[\chi_t^-(\tilde{a}, \tilde{\tau})]$ and $E^F[\chi_t^+(\tilde{a}, \tilde{\tau})] \geq \limsup E^n[\chi_t^+(\tilde{a}, \tilde{\tau})]$. 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 $E^F[\chi_t^-(\tilde{a}, \tilde{\tau})] = E^F[\chi_t^+(\tilde{a}, \tilde{\tau})]$ and $\limsup E^n[\chi_t^+(\tilde{a}, \tilde{\tau})] = \limsup E^n[\chi_t^-(\tilde{a}, \tilde{\tau})]$. Thus,

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

and so $\lim E^n[\chi_t^-(\tilde{a}, \tilde{\tau})]$ exists and equals $E^F[\chi_t^-(\tilde{a}, \tilde{\tau})]$ 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 $\tau(F)$ subject to (A8), all trajectories $\{x_t^n\}_{t \in [0, \tau]} \in \mathcal{B}(F^n)$ are uniformly close to $\mathcal{B}(F)$; that is, \mathcal{B} is upper hemicontinuous.

Upper Hemicontinuity 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 $F = F(a, \tau)$ (Theorem 14.5 of Aliprantis and Border 1999); thus, Berge's maximum theorem implies that the best response mapping $BR : B \rightarrow \Delta(B \times [0, \bar{\tau}])$ is upper hemicontinuous.

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

APPENDIX B. PROOFS FROM SECTION IV

A. Derivation of Equation (16)

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

$$\hat{\Pi}(z_t) = \int_{s=t}^{\tau} e^{-\int_t^s (r + \mu z_u) du} [\pi(z_s) - c(a_s) + \mu z_s \hat{\Pi}(1)] ds.$$

Setting $\psi(s) = e^{-r(s-t)} \hat{\Pi}(z_s)$, $\rho(s) = \mu z_s$, and $\phi(s) = e^{-r(s-t)} (z_s - c(a_s) + \mu z_s \hat{\Pi}(1))$ yields equation (A4). Applying Claim 1, equation (A5) becomes

$$\hat{\Pi}(z_t) = \int_{s=t}^{\tau} e^{-r(s-t)} [\pi(z_s) - c(a_s) + \mu z_s (\hat{\Pi}(1) - \hat{\Pi}(z_s))] ds.$$

Taking the derivative and applying Claim 2, we get

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

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 (A5). Applying Claim 1, equation (A4) becomes

$$\hat{\Pi}'(z_t) = \int_{s=t}^{\tau} e^{-\int_t^s (r+\lambda+\mu(1-2z_u)) du} \mu \left[\pi(z_s) + \hat{\Pi}(1) - \hat{\Pi}(z_s) \right] ds.$$

The envelope theorem then implies equation (16).

B. Observability Raises Investment

Here we argue that if $\pi(x) = x - k$ and k is sufficiently large, investment and the exit time are higher under observable investment than under the benchmark model, $\hat{a}_t > a_t^*$ and $\hat{\tau} > \tau^*$.

First, dropping the positive terms $\mu(\hat{V}(1) - \hat{V}(\hat{z}_s))$ in the integrand and $\mu \hat{z}_u$ in the exponent of (16) implies $\hat{\Gamma}(\hat{\tau} - s) > (1 - e^{-(r+\lambda+\mu)s})/(r + \lambda + \mu)$. Similarly, dropping the negative $-\mu(1 - \hat{z}_u)$ in the exponent in (7) and bounding the integrand above by $V(0, 1) - V(s, z_s^*) < V(0, 1) < (1 - k)/r$ implies $((1 - k)/r)((1 - e^{-(r+\lambda)s})/(r + \lambda)) > \Gamma^*(\tau^* - s)$. Then, for k close to 1, $\hat{\Gamma}(\hat{\tau} - s) > \Gamma^*(\tau^* - s)$ and so $\hat{a}_{\hat{\tau}-s} > a_{\tau^*-s}^*$.

Next, observe that $\hat{V}(1) > V^*(0, 1)$, so the exit conditions (9) and (17) imply $\hat{z}_{\hat{\tau}} < z_{\tau^*}^*$. Together with $\hat{a}_{\hat{\tau}-s} > a_{\tau^*-s}^*$ and $\hat{z}_0 = z_0^* = 1$, this implies $\hat{\tau} > \tau^*$. Since observable investment falls over time, we conclude that $\hat{a}_{\tau^*-s} > \hat{a}_{\hat{\tau}-s} > a_{\tau^*-s}^*$ and hence, $\hat{a}_t > a_t^*$ for all $t = \tau^* - s$.

Investment is also higher at a given level of reputation $x = z$. Indeed, $\hat{a}_{\hat{\tau}-s} > a_{\tau^*-s}^*$ together with $\hat{z}_{\hat{\tau}} < z_{\tau^*}^*$ imply $\hat{z}_{\hat{\tau}-s} < z_{\tau^*-s}^*$, so defining $s' < s$ via $\hat{z}_{\hat{\tau}-s'} = z_{\tau^*-s}^* = z$, investment at z is higher when investment is observable since $\hat{a}(z, z) := \hat{a}_{\hat{\tau}-s'} > \hat{a}_{\hat{\tau}-s} > a_{\tau^*-s}^* =: a^*(z, z)$.

C. Exit Behavior in Theorem 4

Here we establish that a low-quality firm exits at time $\tilde{\tau}^L \in [\underline{\tau}, \infty)$, while the high-quality firm never exits.

Reputation $\{x_t\}$ continuously decreases in t , so a mimicking argument analogous to that in Appendix AA implies that firm value (18) continuously decreases in t . Optimal investment $\{\hat{u}_t\}$ maximizes the integrand in (18) pointwise. For the low firm, $\theta = 0$, the “breakthrough term” is 0 and this integrand also decreases in t . Thus, exiting is optimal for the low firm exactly when the integrand vanishes; since $\hat{V}(\tilde{\tau}^L, L) = 0$, this simplifies to (20). The high-quality firm also benefits from the positive breakthrough term, so the integrand for the high-quality firm

exceeds the integrand for the low-quality firm. Thus, 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 $\underline{\tau}$, recall the upper bound on equilibrium exit times $\bar{\tau}$ from the proof of Theorem 2. Unless the market expects the low-quality firm to start exiting and draws a positive inference from its failure to exit, reputational drift $g(\bar{a}, x) \leq g(\bar{a}, x)$ is strictly negative for $x \in [z^\dagger, 1]$ and takes reputation below z^\dagger at or before $\bar{\tau}$. 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 1, 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 $\psi := -g(a_{\underline{\tau}}, x_{\underline{\tau}})/x_{\underline{\tau}}(1 - x_{\underline{\tau}})$ to offset the negative reputational drift due to learning.

D. Equilibrium Existence in Theorem 4

The existence proof of Theorem 2 does not apply as stated to Theorem 4, since the time-to-exit τ is not bounded above, and the theorem statement is stronger, establishing continuity and monotonicity of x_t . To achieve this stronger result, we first apply the fixed-point argument to the set of *pure strategies* and subsequently use this pure strategy fix point to define an equilibrium with pure investment and randomized exit, as in Appendix BC.

Let L be the space of decreasing, equi-Lipschitz-continuous functions from $[0, \bar{\tau}]$ to $[0, 1]$ with the topology of uniform convergence; L is compact by the Arzela-Ascoli theorem. As in the proof of Theorem 2, we interpret both investment $a = \{a_t\}_{t \in [0, \bar{\tau}]}$ and reputation $x = \{x_t\}_{t \in [0, \bar{\tau}]}$ as elements of L .

For any $x = \{x_t\} \in L$, let $BR(x) \subseteq L \times [0, \bar{\tau}]$ be the set of optimal investment and exit strategies (a^*, τ^*) , assuming that the firm must exit at the latest by time $\bar{\tau}$. Conversely, for any $(a, \tau) \in L \times [0, \bar{\tau}]$, let $B(a, \tau) = x \in L$ be the reputation trajectory defined by $x_0 = 1$ and $\dot{x}_t = g(a_t, x_t)$ for $t \leq \tau$ and $x_t = x_\tau$ for $t \geq \tau$. Standard arguments show that the function B is continuous, and the correspondence BR is upper hemicontinuous and compact valued. To see that the set of pure best responses $BR(x)$ is convex,²⁶ note first that optimal investment a is unique. As for exit, if times $\tau < \tau'$ are both optimal, then reputation x_t must be flat on $[\tau, \tau']$ and, hence, all exit times $\tau'' \in [\tau, \tau']$ are also optimal.

Thus, the correspondence $BR \times B$, mapping $L \times (L \times [0, \bar{\tau}])$ to itself, satisfies the conditions of the Kakutani-Fan-Glicksberg theorem and admits a fixed

²⁶This is the key difference to the baseline model with unknown quality. There, the set of pure best responses is not convex, requiring us to apply the fixed-point argument to the set of mixed strategies, which does not deliver monotonicity and continuity of equilibrium investment.

point $(x, (a, \tau))$. This fixed point must feature exit $\tau < \bar{\tau}$, for otherwise x_t drops below z^\dagger by $\bar{\tau}$, prompting exit.

This fixed point is not yet the desired equilibrium, since investment a_t and reputation x_t are only defined for $t \in [0, \bar{\tau}]$, and a “hard exit” by low-quality firms at τ is not compatible with equilibrium, as argued in Appendix BC. But this is easily fixed: extend investment and reputation by setting $a_t = a_\tau$ and $x_t = x_\tau$ for $t > \bar{\tau}$ and consider the random exit time with rate $\psi = -g(a_\tau, x_\tau)/x_\tau(1 - x_\tau)$ on $[\tau, \infty)$. This constitutes an equilibrium since investment and exit behavior justifies the market beliefs x_t for all $t \geq 0$ and are also optimal given $\{x_t\}$.

E. Firm Observes Private Signals about Its Quality

Here we sketch a model variant where the firm receives additional private signals about its quality; this bridges the models with unknown and known quality and sheds light on the different findings regarding investment levels around the exit threshold. Specifically, suppose the firm observes private breakthroughs that arrive at rate ν when quality is high.

This model resets at a public breakthrough, with both reputation and self-esteem jumping to one. We can thus restrict attention to strategies that specify investment and exit as a function of the time since the last public breakthrough t and the firm’s self-esteem (which now depends on the history of private breakthroughs since the last public breakthrough).

At a private breakthrough self-esteem jumps to one. Absent either breakthrough, self-esteem is governed by $\dot{z}_t = g(a_t(z_t), z_t)$ with

$$g(a, z) = \lambda(a - z) - (\mu + \nu)z(1 - z).$$

Since the market does not observe the firm’s private breakthroughs, it perceives self-esteem z_t as a random variable with expectation x_t . For $\nu = 0$, we recover the baseline model with unknown quality, whereas large ν approach the model with known quality.

We truncate the firm’s cash flow expansion at either kind of breakthrough to obtain

(B1)

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

where the additional term $\nu z_s^* V(s, 1)$, cf (2), captures the firm’s continuation value after a private breakthrough. Analogous to Lemma 7, investment incentives are given by

(B2)

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

These disappear at the exit time τ^* as in Theorem 1, so the firm shirks close to the exit threshold. Thus, even for large ν , i.e., close to the known-quality case, investment vanishes at the exit time, in contrast to Theorem 4. However, the known-quality result that the firm fights until the end is robust in the following sense: the integrand in (B2) increases in ν , so fixing $\tau^* - t$ and considering the limit $\nu \rightarrow \infty$, the integral converges to the known-quality value of quality $\check{V}(t, 1) - \check{V}(t, 0) = \check{I}(t)$, which is boundedly positive and rises in t .

F. Identification

In the text, we discuss that the reputation process $\{X_t\}$ uniquely identifies (μ, λ) and investment $\{a_t\}$. In this Appendix, we show that the exit threshold x^e , revenue $\{x_t\}$, and investment $\{a_t\}$ uniquely identify the cost parameters $(k, c(a))$. Recall $\{a_t\}$ is single-peaked in t —hence, maximized at some \hat{t} and decreasing on $[\hat{t}, \tau]$. Since $x_{\hat{t}} = z_{\hat{t}}$, and x_t is known, we write the firm's on-path value as $V(t)$.

First, we fix an arbitrary maximum value, $V(0)$. The operating cost k then follow from the exit equation,

$$x_{\tau} - k + \mu x_{\tau} V(0) = 0.$$

For the cost function $c(a)$, the first-order condition (8) and the marginal value of effort (7) imply that

$$(B3) \quad c'(a_t) = \lambda \int_t^{\tau} e^{-\int_t^s (r + \lambda + \mu(1 - x_u)) du} \mu [V(0) - V(s)] ds$$

where firm value $V(s)$ is given by

$$(B4) \quad V(s) = \int_{v=s}^{\tau} e^{-\int_s^v (r + \mu x_u) du} [x_v - k - c(a_v) + \mu x_v V(0)] dv.$$

Taking $V(0)$ as exogenous and $\{x_t, a_t\}$ as data, $c'(a_t)$ is expressed as a function of $\{c(a_s) : s > t\}$. Using $a_{\tau} = c(a_{\tau}) = 0$ as boundary condition, $c(a_t)$ is then identified from $t = \tau$ to $t = \hat{t}$.

Next, we identify $V(0)$. Suppose there are two maximum values, $V_1(0) - V_2(0) = \Delta > 0$, consistent with the data. The exit condition together with $V_1(\tau) = V_2(\tau) = 0$ imply $k_1 > k_2$. Let $t < \tau$ be the latest time with $V_1(t) - V_2(t) = \Delta$. Since $V_1(0) - V_1(s) > V_2(0) - V_2(s)$ for all $s > t$, (B3) implies $c'_1(a_t) > c'_2(a_t)$. Since $c_1(a_{\tau}) = c_2(a_{\tau}) = 0$, we have $c_1(a_s) > c_2(a_s)$ for all $s > t$. But then (B4) together with the higher costs in scenario 1 imply that $V_1(t) < V_2(t)$ for the latest time t , and thus $V_1(0) < V_2(0)$, contradicting our original assumption. Intuitively, an increase in the maximum value raises incentives requiring higher costs to justify the observed actions, yet higher costs are inconsistent with a higher value.

This identification approach is very sensitive to the details of the model. In practice, one might rather adopt a more robust, semiparametric approach.

For example, given $\{a_t, x_t\}$, one could calculate value functions (1) for any cost parameters $(k, c(a))$ and then choose those parameter values that minimize the gains from deviations measured by, say, violations of the first-order condition $c'(a_t) = \lambda \Gamma(t)$.

REFERENCES

- Aliprantis, Charalambos D., and Kim C. Border. 1999. *Infinite Dimensional Analysis*. 2nd ed. Berlin: Springer-Verlag.
- Anderson, Simon P., Andre de Palma, and Jacques-Francois Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. Cambridge, MA: MIT Press.
- Arrow, Kenneth J., L. Kamran Bilir, and Alan Sorensen. 2020. "The Impact of Information Technology on the Diffusion of New Pharmaceuticals." *American Economics Journal: Applied Economics* 12 (3): 1–39.
- Atkeson, Andrew, Christian Hellwig, and Guillermo Ordoñez. 2014. "Optimal Regulation in the Presence of Reputation Concerns." *Quarterly Journal of Economics* 130 (1): 415–64.
- Bar-Isaac, Heski. 2003. "Reputation and Survival: Learning in a Dynamic Signalling Model." *Review of Economic Studies* 70 (2): 231–51.
- Board, Simon, and Moritz Meyer-ter-Vehn. 2013. "Reputation for Quality." *Econometrica* 81 (6): 2381–2462.
- Board, Simon, and Moritz Meyer-ter-Vehn. 2022. "Replication Data for: A Reputational Theory of Firm Dynamics." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E118970V1>.
- Bollinger, Bryan, Phillip Leslie, and Alan Sorensen. 2011. "Calorie Posting in Chain Restaurants." *American Economic Journal: Economic Policy* 3 (1): 91–128.
- Bonatti, Alessandro, and Johannes Hörner. 2011. "Collaborating." *American Economic Review* 101 (2): 632–63.
- Bonatti, Alessandro, and Johannes Hörner. 2017. "Career Concerns with Exponential Learning." *Theoretical Economics* 12 (1): 425–75.
- Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow. 2012. "The Evolution of Brand Preferences: Evidence from Consumer Migration." *American Economic Review* 102 (6): 2472–2508.
- Cabral, Luís. 2016. "Living up to Expectations: Corporate Reputation and Persistence of Firm Performance." *Strategy Science* 1 (1): 2–11.
- Cabral, Luís, and Ali Hortaçsu. 2010. "The Dynamics of Seller Reputation: Evidence from eBay." *Journal of Industrial Economics* 58 (1): 54–78.
- Cisternas, Gonzalo. 2017. "Two-Sided Learning and the Ratchet Principle." *Review of Economic Studies* 85 (1): 307–51.
- Cohen, Wesley M., and Daniel A. Levinthal. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation." *Administrative Science Quarterly* 35 (1): 128–52.
- Cutler, David M., Robert S. Huckman, and Mary Beth Landrum. 2004. "The Role of Information in Medical Markets: An Analysis of Publicly Reported Outcomes in Cardiac Surgery." *American Economic Review* 94 (2): 342–46.
- Davis, Mark. 1993. *Markov Models and Optimization*. Chapman and Hall.
- Dilmé, Francisc. 2019. "Reputation Building through Costly Adjustment." *Journal of Economic Theory* 181: 586–626.
- Dixit, Avinash K., and Joseph E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 67 (3): 297–308.
- Ericson, Richard, and Ariel Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62 (1): 53–82.
- Foster, Lucia, John Haltiwanger, and Chad Syverson. 2016. "The Slow Growth of New Plants: Learning about Demand?" *Economica* 83 (329): 91–129.
- Fung, Archon, Mary Graham, and David Weil. 2007. *Full Disclosure: The Perils and Promise of Transparency*. Cambridge, UK: Cambridge University Press.
- Gabaix, Xavier. 2009. "Power Laws in Economics." *Annual Review of Economics* 1: 255–94.
- Gale, Douglas, and Robert Rosenthal. 1994. "Price and Quality Cycles for Experience Goods." *RAND Journal of Economics* 25 (4): 590–607.
- Gardiner, Crispin W. 2009. *Stochastic Methods*. Berlin: Springer.

- Goldfarb, Avi.** 2007. "Schlitz: Why Schlitz Hit the Fan." In *Industry and Firm Studies*, 4th ed., edited by Victor J. Tremblay and Carol Horton Tremblay, 321–41. New York: Routledge.
- Halac, Marina, and Andrea Prat.** 2016. "Managerial Attention and Worker Performance." *American Economic Review* 106 (10): 3104–32.
- Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg.** 2005. "Market Value and Patent Citations." *RAND Journal of Economics* 36 (2): 16–38.
- Hastings, Justine S., and Jeffrey M. Weinstein.** 2008. "Information, School Choice, and Academic Achievement: Evidence from Two Experiments." *Quarterly Journal of Economics* 123 (4): 1373–1414.
- Hauser, Daniel.** 2021a. "Censorship and Reputation." Unpublished.
- Hauser, Daniel.** 2021b. "Promoting a Reputation For Quality." Unpublished.
- Hopenhayn, Hugo A.** 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5): 1127–50.
- Jin, Ginger Zhe, and Alan T. Sorensen.** 2006. "Information and Consumer Choice: The Value of Publicized Health Plan Ratings." *Journal of Health Economics* 25 (2): 248–75.
- Jin, Ginger Zhe, and Phillip Leslie.** 2003. "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards." *Quarterly Journal of Economics* 118 (2): 409–51.
- Jin, Ginger Zhe, and Phillip Leslie.** 2009. "Reputational Incentives for Restaurant Hygiene." *American Economic Journal: Microeconomics* 1 (1): 237–67.
- Jovanovic, Boyan.** 1982. "Selection and the Evolution of Industry." *Econometrica* 50 (3): 649–70.
- Kolmogorov, Andrey.** 1931. "Über die analytischen Methoden der Wahrscheinlichkeitsrechnung." *Mathematische Annalen* 104: 415–58.
- Kolstad, Jonathan T.** 2013. "Information and Quality When Motivation Is Intrinsic: Evidence from Surgeon Report Cards." *American Economic Review* 103 (7): 2875–2910.
- Luca, Michael.** 2016. "Reviews, Reputation, and Revenue: The Case of Yelp.com." Harvard Business School Working Paper 12-016.
- Luca, Dara Lee, and Michael Luca.** 2018. "Survival of the Fittest: The Impact of the Minimum Wage on Firm Exit." Harvard Business School Working Paper 17-088.
- Luca, Michael, and Jonathan Smith.** 2013. "Salience in Quality Disclosure: Evidence from the U.S. News College Rankings." *Journal of Economics and Management Strategy* 22 (1): 58–77.
- Mailath, George J., and Larry Samuelson.** 2001. "Who Wants a Good Reputation." *Review of Economic Studies* 68 (2): 415–41.
- Marinovic, Iván, Andrzej Skrzypacz, and Felipe Varas.** 2018. "Dynamic Certification and Reputation for Quality." *American Economic Journal: Microeconomics* 10 (2): 58–82.
- Milgrom, Paul, and Ilya Segal.** 2002. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica* 70 (2): 583–601.
- Parsa, H.G. John T. Self, David Njite, and Tiffany King.** 2005. "Why Restaurants Fail." *Cornell Hotel and Restaurant Administration Quarterly* 46 (3): 304–22.
- Pope, Devin G.** 2009. "Reacting to Rankings: Evidence from 'America's Best Hospitals'." *Journal of Health Economics* 28 (6): 1154–65.
- Rob, Rafael, and Arthur Fishman.** 2005. "Is Bigger Better? Customer Base Expansion through Word-of-Mouth Reputation." *Journal of Political Economy* 113 (5): 1146–75.
- Sannikov, Yuliy.** 2013. "Moral Hazard and Long-Run Incentives." Unpublished.
- Serrano, Carlos J.** 2010. "The Dynamics of the Transfer and Renewal of Patents." *RAND Journal of Economics* 41 (4): 686–708.
- Sorensen, Alan T.** 2007. "Bestseller Lists and Product Variety." *Journal of Industrial Economics* 55 (4): 715–38.
- Syverson, Chad.** 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2): 326–65.
- Tadelis, Steven.** 1999. "What's in a Name? Reputation as a Tradeable Asset What's in a Name? Reputation as a Tradeable Asset." *American Economic Review* 89 (3): 548–63.
- Tadelis, Steven.** 2002. "The Market for Reputations as an Incentive Mechanism the Market for Reputations as an Incentive Mechanism." *Journal of Political Economy* 110 (4): 854–82.
- Vellodi, Nikhil.** 2021. "Ratings Design and Barriers to Entry." Unpublished.
- Vial, Bernardita, and Felipe Zurita.** 2017. "Entrants' Reputation and Industry Dynamics." *International Economic Review* 58 (2): 529–59.