Reputation for Quality*

Simon Board[†] and Moritz Meyer-ter-Vehn[‡]

This Version: January 8, 2010
First Version: June 2009.

Abstract

We propose a new model of firm reputation where product quality is persistent and determined by the firm's past investments. Reputation is then modeled directly as the market belief about quality. We analyze how investment incentives depend on the firm's reputation and derive implications for reputational dynamics.

Reputational incentives depend on the specification of market learning. When consumers learn about quality through good news events, incentives decrease in reputation and there is a unique work-shirk equilibrium with convergent dynamics. When learning is through bad news events, incentives increase in reputation and there is a continuum of shirk-work equilibria with divergent dynamics. Across all imperfect learning processes with Brownian and Poisson signals, we show that when costs are low there exists a work-shirk equilibrium with convergent dynamics. This equilibrium is essentially unique if market learning contains a good news or Brownian component.

1 Introduction

In most industries firms can invest into the quality of their products through human capital investment, research and development, and organizational change. While imperfect monitoring by consumers gives rise to a moral hazard problem, the firm can share in the created value by building a reputation for quality, justifying premium prices. This paper analyzes the investment incentives in such a market, characterizing how they depend on the current reputation of the firm and the market learning process.

*We thank Andy Atkeson, Heski Bar-Isaac, Dirk Bergemann, Willie Fuchs, Christian Hellwig, Hugo Hopenhayn, Boyan Jovanovic, David Levine and Ron Siegel for their ideas. We have also received helpful comments from many others, including seminar audiences at ASSA 2010, ASU, Bonn, CETC 2009, Chicago, Cowles 2009, Frankfurt, Munich, NYU, Penn, Pitt, SED 2009, Stern, SWET 2009, Toulouse, UCI, UCL, UCSD, UNC-Duke, UT-Austin. We gratefully acknowledge financial support from NSF grant 0922321. Keywords: Reputation, R&D, Dynamic Games, Imperfect Monitoring, Industry Dynamics. JEL: C73, L14

[†]Department of Economics, UCLA. http://www.econ.ucla.edu/sboard/
[‡]Department of Economics, UCLA. http://www.econ.ucla.edu/people/Faculty/Meyer-ter-Vehn.html

Our key innovation over classical models of reputation and repeated games is to suppose product quality is a function of past investments rather than just current effort. From a modeling perspective, the introduction of persistence turns quality into a state variable and allows us to model reputation directly as the market's belief about quality. Therefore, our firm works to actually build the quality that underlies reputation, rather than to signal an exogenous type or to avoid punishments by a counterparty. From an economic perspective, persistence alleviates the firm's moral hazard problem. Investment affects quality in a lasting way, and yields rewards even if the firm is believed to be shirking in the future. As a result, the model gives rise to simple Markovian equilibria that explain when a firm will build a reputation, when it will work to maintain its reputation, and when it will choose to run it down.

In the model, illustrated in Figure 1, one long-lived firm sells a product of high or low quality to a continuum of identical short-lived consumers. Product quality is a stochastic function of the firm's past investments. Consumers observe neither quality nor investment directly, but experienced utility acts as an imperfect market signal of the underlying product quality. At each point in time, consumers' willingness to pay is determined by the market belief that the quality is high, x_t, which we call the *reputation* of the firm. This reputation changes over time as a function of (a) the equilibrium beliefs of the firm's investments, and (b) market learning via consumers' realized utility.

Motivated by the Levy Decomposition Theorem, we suppose that market learning has two components: (1) a Brownian motion capturing continuous information such as consumer reports, and (2) Poisson processes capturing discrete events such as product failures. A Poisson event is a *good news* signal if it indicates high quality, and a *bad news* signal if it indicates low quality. Market learning is *imperfect* if no Poisson signal perfectly reveals the firm's quality.[1]

Investment is incentivized by the difference in value between a high and low quality firm, which we call the *(asset) value of quality.* If the value of quality is decreasing in the firm's reputation, then the incentive of a low-reputation firm to build a reputation exceeds the incentive of a high-reputation firm to maintain its reputation. Equilibria are then *work-shirk*, in that the firm works if its reputation lies below a cutoff x^*; and dynamics are *convergent*, in that the limit distribution of the firm's reputation is independent of its initial reputation. Conversely, if the value of quality is increasing in the firm's reputation, equilibria are *shirk-work* and dynamics are *divergent.*

Quality derives its value by increasing expected utility to consumers and thereby the firm's reputation. Crucially, as quality is persistent, this reputational payoff does not take the form of an immediate one-off reputational boost, but accrues to the firm as a stream of future *reputational*

[1] There are many examples of these learning processes. Continuous updating may occur as drivers learn about the build-quality of a car, as clients learn about the skills of a consultancy, and as callers learn about the customer service of a telephone service provider. Good news signals may occur in academia when a paper becomes famous, in the bio-tech industry when a trial succeeds, and for actors when they win an Oscar. Bad news signals may occur in the computer industry when batteries explode, in the financial sector when a borrower defaults, and for doctors when they are sued for medical malpractice.

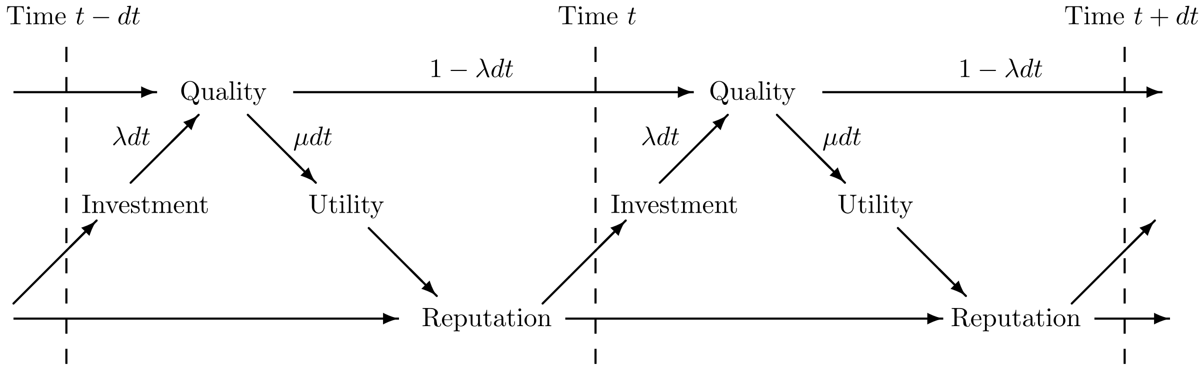

Figure 1: **Gameform.** Quality is a function of the firm's investment and its past quality. Consumer utility is an imperfect signal of quality that the market uses to update the firm's reputation.

dividends. Theorem 1 formalizes this idea by writing the asset value of quality as the present value of its future reputational dividends.

The shape of reputational dividends, and therefore the value of quality, depend on the specification of the underlying market learning process. For perfect good news, where high quality gives rise to product *breakthroughs* that boost reputation to one, reputational dividends decrease in the firm's reputation. For perfect bad news signals, where high quality insures the firm against product *breakdowns* that destroy its reputation, reputational dividends increase in the firm's reputation. For imperfect learning processes, where Bayesian learning ceases for extreme reputations, reputational dividends are hump-shaped.

Using Theorem 1, it is straightforward to characterize equilibria with perfect poisson learning. Under perfect good news, there is a unique work-shirk equilibrium. Intuitively, a breakthrough takes the firm's reputation to 1, so high quality is more valuable to a firm with low reputation. Conversely, under perfect bad news, there is a continuum of *shirk-work* equilibria. Intuitively, a breakdown takes the firm's reputation to 0, so high quality is more valuable to a firm with high reputation.

Theorem 2 analyzes imperfect learning processes when the cost of investment is sufficiently small. While the hump-shaped reputational dividends suggest a shirk-work-shirk equilibrium, we surprisingly show existence of a work-shirk equilibrium. We also provide a condition on the learning process for this equilibrium to be unique. The work-shirk result relies on a fundamental asymmetry. For $x \approx 1$, work is not sustainable: If the firm is believed to work, its reputation stays high and reputational dividends stay small, undermining incentives to invest. For $x \approx 0$, work is sustainable: If the firm is believed to work, its reputation drifts up and reputational dividends increase, generating incentives to invest. Crucially, a firm with a low reputation works not because of the immediate, small reputational dividends but because the larger future dividends it expects after its reputation has drifted up. Thus, persistent quality and the endogenous reputational drift jointly introduce an asymmetry that is responsible for the work-shirk equilibrium.

The condition for uniqueness, (HOPE), requires that reputation drifts up with positive probability, even if the firm is believed to be shirking; it is satisfied if market learning has a Brownian, or good news component. This condition unravels putative shirk-work-shirk equilibria, as the possibility of enjoying the favorable beliefs in the work-region guarantees high investment incentives for a firm around the shirk-work cutoff. To the contrary, without (HOPE), adverse beliefs below a shirk-work cutoff are self-fulfilling and support a continuum of shirk-work-shirk equilibria.

Finally, we connect our analysis to repeated games models where quality is chosen in every period, by taking the quality obsolescence rate λ to infinity. With complete information, an increase in λ front-loads the returns to investment and increases investment incentives. With incomplete information, there is a countervailing effect: For large values of λ, equilibrium beliefs dominate market learning in determining reputational dynamics. In a work-shirk profile, market beliefs rapidly drift towards the cutoff and reputational dividends vanish. Thus, there are no work-shirk equilibria for high λ, but full shirking is an equilibrium. In contrast, a shirk-work cutoff induces divergent reputational dynamics and can therefore incentivize investment when λ is high. Accordingly, we find that investment disappears under perfect good news, while any shirk-work cutoff can be sustained under perfect bad news.

1.1 Literature

The key feature distinguishing our paper from classical models of reputation and repeated games is that product quality is a function of past investments rather than current effort. This difference is important. In classical models, the firm exerts effort to convince the market that it will also exert effort in the future. In our model, a firm's investment increases its quality and future revenue independent of market beliefs about future investment, since quality is persistent.

The two reputation models closest to ours are Mailath and Samuelson (2001) and Holmström (1999), which both model reputation as the market's belief about some exogenous state variable. The mechanisms linking effort, type and utility are depicted in Figure 2. In Mailath and Samuelson (2001) a competent firm, that can choose to work or shirk, tries to distinguish itself from an incompetent type, that always shirks. A reputation for competence benefits the firm to the degree that the market expects a competent firm to work. With imperfect monitoring, firms with high reputations shirk because updating is slow. This causes effort to unravel from the top as a firm just below a putative work-shirk cutoff finds it unprofitable to further invest into its reputation. In our model, persistent quality prevents this unraveling because current investment affects the firm's future reputation and revenue irrespective of beliefs about its future investments.

Holmstrom's (1999) signal-jamming model is similar to ours in that the firm's type directly affects consumers' utility. In this model, the firm works to induce erroneous market beliefs that its exogenous ability type is higher than in reality. This is in stark contrast to our model, where a firm obtains a high reputation by actually building a high quality.

In both of these papers, learning about a fixed type eventually vanishes and so do reputational

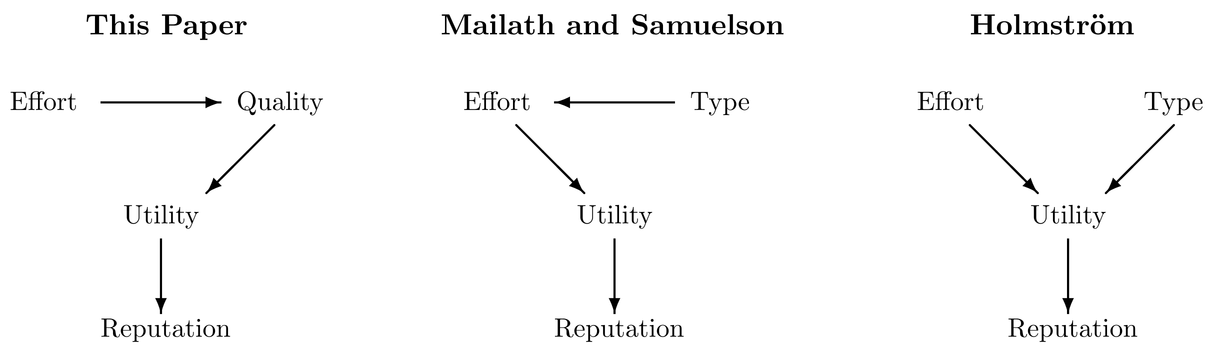

Figure 2: **Comparison of Reputation Models.** The literature usually models reputation as belief over some exogenous type. This type affects consumer utility either directly, as in Holmstrom (1999) or indirectly through the cost of effort, Mailath and Samuelson (2001). In contrast, our firm has no exogenous type, but quality is persistent and reputation is directly the market belief about quality.

incentives. These are instances of a more general theme: Cripps, Mailath, and Samuelson (2004) show that with imperfect monitoring and fixed types, reputation is always a short-run phenomenon. A long-run analysis of reputation requires "... some mechanism by which the uncertainty about types is continually replenished".[2] Our stochastic investment into quality is a natural candidate for this mechanism. Unlike in models of exogenous shocks, such as Holmström (1999), in which reputation simply trails these shocks, the reputational dynamics of our model are driven by the forward-looking reputational incentives.[3]

There is a wider literature on lifecycle effects in reputation models, as surveyed in Bar-Isaac and Tadelis (2008). Some of these results can be understood through our trichotomy of good, bad and Brownian learning: With good news learning, firms with low reputations try to build, or buy a reputation (Tadelis (1999)). With bad news learning, firms with high reputations have high incentives to maintain them (Diamond (1989)). With Brownian learning, reputational incentives are hump-shaped (Benabou and Laroque (1992), Mailath and Samuelson (2001)).

In contrast to the repeated games literature (e.g. Fudenberg, Kreps, and Maskin (1990)), our model is distinguished by an evolving state variable. Investment directly feeds through to future reputation and revenue in our model, rather than preventing deliberate punishment by a counterparty. However, the models are connected by a common limit: As the frequency of play in a repeated game increases, it approaches a continuous-time game where quality is chosen at any $t \in \mathbb{R}$. This is also the limit of our reputation model as we take the quality obsolescence rate λ to ∞.

In a repeated prisoners' dilemma with frequent actions, Abreu, Milgrom, and Pearce (1991) and Sannikov and Skrzypacz (2007) show that only discrete, "bad news" signals that indicate defection

[2]The same theme appears in the dynamic learning literature, e.g. Keller and Rady (1999).
[3]Liu (2009) gives an alternative explanation of long-run reputational dynamics that is driven by imperfect, costly recall and lack of a public posterior.

can sustain cooperation. Brownian and good news signals are too noisy to deter defections without destroying all surplus by punishments on the equilibrium path. Thus, sustained cooperation depends on the learning process in the same way as in our model. While the common limit already suggests this analogy, our model highlights an alternative mechanism that distinguishes the role of bad news signals in overcoming moral hazard, namely divergent reputational dynamics.[4]

Our model has clear empirical predictions concerning the dynamics of reputations. While there is a growing empirical literature concerning reputation (Bar-Isaac and Tadelis (2008)), most of these papers are static, focusing on quantifying the value of reputation. One notable exception is Cabral and Hortaçsu (2009) which shows that an eBay seller who receives negative feedback becomes more likely to receive additional negative feedback, and is more likely to exit. This is consistent with our bad news case where a seller who receives negative feedback stops investing.

2 Model

Overview: There is one firm and a continuum of consumers. Time $t \in [0, \infty)$ is continuous and infinite; the common interest rate is $r \in (0, \infty)$. At time t the firm produces one unit of a product that can have high or low quality, $\theta_t \in \{L = 0, H = 1\}$. The expected instantaneous value of the product to a consumer equals $\theta_t dt$. The market belief about product quality $x_t = \Pr(\theta_t = H)$ is called the firm's reputation. The firm chooses investment $\eta_t \in [0, 1]$ at cost $c\eta_t dt$.

Technology: Product quality θ_t is a function of past investments $(\eta_s)_{0 \leq s \leq t}$ via a Poisson process with arrival rate λ that models quality obsolescence. Absent a shock quality is constant, $\theta_{t+dt} = \theta_t$; at a shock previous quality becomes obsolete and quality is determined by the level of investment, $Pr(\theta_{t+dt} = H) = \eta_t$.[5] Quality at time t is then a geometric sum of past investments,

$$\Pr(\theta_t = H) = \int_0^t \lambda e^{\lambda(s-t)} \eta_s ds + e^{-\lambda t} \Pr(\theta_0 = H). \tag{2.1}$$

Information: Investment η_t and actual product quality θ_t are observed only by the firm. Consumers learn about quality through a stochastic process $dZ_t = dZ(\theta_t, \varepsilon_t)$, with noise ε_t that is

[4] Our model is distantly related to other literatures. In the contract design literature, models with persistent effort have been studied by Fernandes and Phelan (2000) and Jarque (2008). The contrast between classical reputation models and our model, is also analogous to the difference between models of industry dynamics with exogenous types (Jovanovic (1982), Hopenhayn (1992)) and those with endogenous capital accumulation (Ericson and Pakes (1995)).

[5] This formulation provides a tractable way to model product quality as a function of past investments. One can interpret investment as the choice of absorptive capacity, determining the ability of a firm to recognise new external information and apply it to commercial ends (Cohen and Levinthal (1990)). Equivalently, one could assume the firm observes arrivals of technology shocks, and then chooses whether to adopt the new technology at cost $k = c/\lambda$. Yet another equivalent interpretation is that a low-quality firm chooses the arrival rate of high quality from $[0, \lambda]$ at marginal cost c/λ, and a high-quality firm can abate the intensity λ of the arrival of low quality at marginal cost c/λ.

i.i.d. across time and independent of the quality obsolescence process. The signal dZ_t subsumes information from experienced consumer utility.[6] Motivated by the Levy Decomposition Theorem, we suppose Z_t is generated by a Brownian motion and a finite number of Poisson processes. The Brownian motion $Z_{B,t}$ is given by

$$dZ_{B,t} = \mu_B \theta_t dt + dW_t$$

where the drift depends on quality, and W_t is the Wiener process. The Poisson process are indexed by $y \in \{1, \ldots, Y\}$ with quality-dependent arrival rates $\mu_{\theta,y}$. Signal y is *good news* if the net arrival rate $\mu_y := \mu_{H,y} - \mu_{L,y}$ is positive, *perfect good news* if $\mu_{L,y} = 0$, *bad news* if $\mu_y < 0$, and *perfect bad news* if $\mu_{H,y} = 0$. Market learning is *imperfect* if there is no perfect Poisson signal.

Reputation Updating: The reputation increment $dx_t = x_{t+dt} - x_t$ is governed by the signal dZ_t and market beliefs about investment $\widetilde{\eta}_t$. By independence, dx_t can be decomposed additively:

$$dx_t = \lambda(\widetilde{\eta}_t - x_t)dt + x_t(1 - x_t)\frac{\Pr(dZ_t|H) - \Pr(dZ_t|L)}{x_t \Pr(dZ_t|H) + (1 - x_t) \Pr(dZ_t|L)}. \tag{2.2}$$

The first term is the differential version of equation (2.1); if expected quality after a technology shock $\widetilde{\eta}_t$ exceeds current expected quality x_t reputation drifts up. The second term is the standard Bayesian increment, which we evaluate explicitly in Section 4. When conditioning on quality θ, we write $d_\theta x_t$ for the increments.

Profit and Consumer Surplus: The firm and consumers are risk-neutral. At time t the firm sets price equal to the expected value x_t, so consumers' expected utility is 0. The firm's profit is $(x_t - c\eta_t)dt$ and its discounted present value is thus given by:

$$V_\theta\left(x; \eta, \widetilde{\eta}\right) := \int_{t=0}^{\infty} e^{-rt} \mathbb{E}_{\theta_0=\theta, x_0=x, \eta, \widetilde{\eta}}\left[x_t - c\eta_t\right] dt. \tag{2.3}$$

Markov-Perfect-Equilibrium: We assume Markovian beliefs $\widetilde{\eta} = \widetilde{\eta}\left(x\right)$ and show below that optimal investment $\eta = \eta\left(x\right)$ is independent of history and current product quality θ.[7] A Markov-Perfect-Equilibrium $\langle \eta, \widetilde{\eta} \rangle$ consists of a Markovian investment function $\eta : [0,1] \to [0,1]$ for the

[6]For example, with pure Brownian learning, we can let dZ_t/μ_B measure the instantaneous utility of the agent.

[7]In principle, the firm's effort choice η as well as market beliefs $\widetilde{\eta}$ could depend on the entire public history $Z^t = (Z_s)_{0 \le s < t}$, as well as the private history $\theta^t = (\theta_s)_{0 \le s < t}$ and time t. We assume that market beliefs $\widetilde{\eta}$ are Markovian because we think of the continuum of consumers as sharing their experience in a sufficient, yet incomplete manner, e.g. through consumer reports. For Markovian beliefs $\widetilde{\eta}$, all payoff relevant parameters at time t depend on the history only via the current product quality θ_t and the firm's reputation x_t. Thus, the optimal effort choice of the firm only depends on these two parameters.

firm and Markovian market beliefs $\widetilde{\eta} : [0,1] \to [0,1]$ such that 1) investment maximizes firm value, $\eta \in \eta^*(\widetilde{\eta}) := \arg\max_\eta \{V_\theta(x; \eta, \widetilde{\eta})\}$, and 2) market beliefs are correct, $\widetilde{\eta} = \eta$. In a Markovian equilibrium η, we write the firm's value as a function of its quality and its reputation, $V_\theta(x)$.

2.1 Optimal Investment Choice

The benefit of investment over $[t, t+dt]$ is the probability of a technology shock hitting, λdt, times the difference in value functions $\Delta(x) := V_H(x) - V_L(x)$, which we call the *value of quality.* The marginal cost of investment is c, so optimal investment $\eta(x)$ is given by

$$\eta(x) = \begin{cases} 1 & \text{if } c < \lambda\Delta(x), \\ 0 & \text{if } c > \lambda\Delta(x). \end{cases} \tag{2.4}$$

Quality after the shock is independent of current quality, so optimal investment is independent of current quality. Lemma 1 summarizes this discussion:

Lemma 1 *For Markovian beliefs $\widetilde{\eta}(x)$ there is an optimal Markovian investment function $\eta(x)$ that depends on the firm's reputation but not on its product quality. Additionally, $\eta(x)$ satisfies the "bang-bang" equation (2.4).*

Equation (2.4) makes the model tractable and is the reason that we assume the cost of investment to be independent of product quality and past investment. An implication of equation (2.4) is that our results are not driven by the asymmetric information about product quality θ, but solely by the unobserved investment η into future quality.

2.2 Cutoff Equilibria and Reputational Dynamics

We call an equilibrium *work-shirk*, if there exists a cutoff x^* such that a firm with low reputation $x < x^*$ invests, $\eta(x) = 1$, whereas a firm with a high reputation $x > x^*$ does not, $\eta(x) = 0$. The opposite case, where low reputations shirk and high reputations work, is called a *shirk-work* equilibrium.

Work-shirk equilibria and shirk-work equilibria have opposite reputational dynamics. Net of market learning, dynamics $dx = \lambda(\widetilde{\eta}_t - x_t)dt$ are *convergent* in a work-shirk equilibrium, i.e. $dx > 0$ for $x < x^*$ and $dx < 0$ for $x > x^*$, but *divergent* in a shirk-work equilibrium. We will see in Section 6 that for high values of λ incentives disappear in work-shirk profiles, but not in shirk-work profiles.

We consider two investment profiles η, η' as close if their value functions are cloze. Formally, let $dist(\eta, \eta') := \sup_{x \in (0,1), \theta} \{|V_{\theta,\eta}(x) - V_{\theta,\eta'}(x)|\}$. This pseudo-metric captures a fundamental asymmetry between work-shirk and shirk-work cutoffs. The *full-work* profile, with $\eta(x) = 1$ for all x, is close to a work-shirk profile with cutoff $x^* \approx 1$ but not to a shirk-work profile with cutoff $x^* \approx 0$. In the former two profiles the reputational trajectories x_t underlying the value function are always close; to the contrary in the shirk-work profile the reputational trajectory x_t can be trapped in the shirk region $[0, x^*]$ when $x_0 \approx 0$.

2.3 Welfare

Suppose product quality is publicly observed. The benefit of investing equals the obsolescence rate λ, times the price differential 1, divided by the effective discount rate $r + \lambda$. Thus first-best investment is given by:

$$\eta = \begin{cases} 1 & \text{if } c < \frac{\lambda}{r+\lambda} \\ 0 & \text{if } c > \frac{\lambda}{r+\lambda} \end{cases} . \tag{2.5}$$

There is no equilibrium with positive investment if $c > \lambda/\left(r + \lambda\right)$: Investment decreases welfare and consumers receive zero utility in equilibrium, so firm profits must be negative. The firm therefore prefers to shirk at all levels of reputation, thereby guaranteeing itself a non-negative payoff.

We thus restrict attention in the paper to the case $c < \lambda/\left(r + \lambda\right)$.

3 Value of Quality

The firm's value $V_\theta(x)$ is a function of its reputation x and its quality θ. While reputation directly determines revenue, quality derives its value indirectly through its effect on reputation. More precisely, Theorem 1 shows that the value of quality can be written as a present value of future reputational dividends.

We first establish that equilibrium value $V_\theta(x)$ is increasing in reputation x. To do so, we need to account for the contingency that a firm with a high initial reputation x' may shirk, lose its product quality, and fall behind a firm with a low initial reputation x''. However, one feasible strategy for the high reputation firm is to mimic the investment of the low reputation firm, i.e. invest after a history of public signals dZ^t exactly when the firm with initial valuation x'' would invest. The value V of the high reputation firm following this strategy is (1) strictly higher than $V_\theta(x'')$ because it incurs the same investment costs and its reputation never falls behind, and (2) lower than $V_\theta(x')$ by the optimality of the equilibrium strategy. To summarize:

Lemma 2 *Given an optimal response to market beliefs $\eta^*(\widetilde{\eta})$, the value function of the firm $V_\theta(x; \eta^*(\widetilde{\eta}), \widetilde{\eta})$ is strictly increasing in reputation x and increasing in market beliefs $\widetilde{\eta}$.*

Lemma 2 implies that for two equilibria η, η', with $\eta'(x) \geq \eta(x)$ for all x, the firm's value is increasing in investment $V_\theta(x; \eta', \eta') \geq V_\theta(x; \eta, \eta)$.

To analyze the asset value of quality $\Delta(x) = V_H(x) - V_L(x)$, we now develop the value functions into current profits and continuation values. Current profits cancel because both current revenue and costs depend only on reputation but not on quality. The continuation values are discounted at both the interest rate r and the quality obsolescence rate λ.

$$\Delta(x) = (1 - rdt)(1 - \lambda dt)\mathbb{E}\left[V_H(x + d_H x) - V_L(x + d_L x)\right] \tag{3.1}$$

In this expression, the reputational dividend is captured in the fact that the reputation of the high quality firm has increased by $d_H x - d_L x$ relative to the reputation of the low quality firm. For a recursive formulation of $\Delta(x)$ we need to evaluate the value functions at the same levels of future reputation, and do so by adding and subtracting a term $V_H(x + d_L x)$ to obtain

$$\Delta(x) = (1 - rdt - \lambda dt)\left(D_H(x)\, dt + \mathbb{E}\left[\Delta(x + d_L x)\right]\right),$$

where

$$D_\theta(x) := \mathbb{E}[V_\theta(x + d_H x) - V_\theta(x + d_L x)]/dt$$

is the *reputational dividend* of quality. Rather than affecting current profits, quality pays off through a positive reputational drift that increases future profits.

Integrating (3.1) yields equation (3.2) in Theorem 1, which expresses the asset value of quality as the discounted sum of future reputational dividends. This expression serves as a work-horse throughout the paper.

Theorem 1 *Fix any Markovian beliefs $\widetilde{\eta}$ and a Markovian best response $\eta^*(\widetilde{\eta})$. Then two closed-form expressions for the value of quality $\Delta(x)$ are given by:*

$$\begin{aligned} \Delta(x) &= \int_0^\infty e^{-(r+\lambda)t}\mathbb{E}_{x_0=x,\theta^t=L}[D_H(x_t)]dt, && (3.2)\\ &= \int_0^\infty e^{-(r+\lambda)t}\mathbb{E}_{x_0=x,\theta^t=H}[D_L(x)]dt, && (3.3)\end{aligned}$$

where $\theta^t = L$ is short for $\theta_s = L$ for all $s \in [0,t]$.

Proof. To integrate up (3.1), fix x and set $\psi(t) := \mathbb{E}_{x_0=x,\theta^t=L}\left[\Delta(x_t)\right]$. Up to terms of order $o(dt)$ we have

$$\begin{aligned} -d\left(\psi(t)\, e^{-(r+\lambda)t}\right) &= -e^{-(r+\lambda)t}\left(\psi(t+dt) - \psi(t) - (r+\lambda)\, dt\psi(t)\right)\\ &= e^{-(r+\lambda)t}\mathbb{E}_{x_0=x,\theta^t=L}\left[-\mathbb{E}_{x_t}\left[\Delta(x_t + d_L x_t)\right] + \left(1 + (r+\lambda)\, dt\right)\Delta(x_t)\right]\\ &= e^{-(r+\lambda)t}\mathbb{E}_{x_0=x,\theta^t=L}\left[D_H(x_t)\right] dt \end{aligned}$$

and (3.2) follows.

Equation (3.3) follows from the alternative decomposition of (3.1) when we add and subtract $V_L(x + d_H x)$ instead of $V_H(x + d_L x)$. $\square$

While standard reputation models incentivize effort by an immediate effect on the firm's reputation, investment in our model pays off through quality with a delay. Once quality is established, it is persistent and generates a stream of reputational dividends until it becomes obsolete. We must therefore evaluate the reputational incentives at future levels of reputation x_t, rather than just at the current level x.

Corollary 1 *Fix any Markovian beliefs $\widetilde{\eta}$ and a Markovian best response $\eta^*(\widetilde{\eta})$. For a given reputation x, a high-quality firm has a higher value than a low-quality firm, i.e. $V_H(x) \geq V_L(x)$.*

Proof. By the updating equation (2.2) we have $d_H x \geq d_L x$, by Lemma 2 we get $D_\theta(x) = V_\theta(x_t + d_H x_t) - V_\theta(x_t + d_L x_t) \geq 0$. Finally by Theorem 1 we get $\Delta(x) \geq 0$. □

4 Imperfect Learning

Theorem 2 shows that for imperfect market learning processes and sufficiently small costs, there exists a work-shirk equilibrium and, under condition (HOPE), this equilibrium is essentially unique. We prove Theorem 2 by evaluating the reputational dividends that constitute the value of quality and incentivize investment. These reputational dividends are the sum of a Brownian component (4.2) and a Poisson component (4.4).

For the Brownian component of learning, updating evolves according to

$$d_\theta x = \mu_B x(1-x)\left(\mu_B(\theta - x)\,dt + dW\right). \tag{4.1}$$

High quality induces a positive reputational drift of $\mu_B(1-x)$ while low quality induces a negative reputational drift of $-\mu_B x$. Bayesian updating is slow for extreme reputations, where the evolution is dampened by a factor $x(1-x)$.

To calculate the reputational dividend we develop value functions with Itô's formula to obtain:

$$\mathbb{E}_x[V_H(x + d_\theta x)] = V_H(x) + \mu_B^2 x(1-x)(\theta - x)\,V_H'(x)dt + \frac{(\mu_B x(1-x))^2}{2} V_H''(x)dt.$$

The Brownian component of the reputational dividend is thus the difference in reputational drift between a high quality and a low quality firm:

$$D_H(x) = \mathbb{E}_x[V_H(x + d_H x) - V_H(x + d_L x)]/dt = \mu_B^2 x(1-x)V_H'(x). \tag{4.2}$$

The dividend declines to zero for extreme reputations as the factor $x(1-x)$ dampens reputational updating.

For the Poisson component of learning, recall that $\mu_y = \mu_{H,y} - \mu_{L,y}$ is the net arrival rate of good news. The reputational increment is given by

$$dx = \sum_y \mu_y x(1-x) \begin{cases} (\mu_{H,y}x + \mu_{L,y}(1-x))^{-1} & \text{at arrival } y, \\ -dt & \text{otherwise.} \end{cases} \tag{4.3}$$

If y signals good news because $\mu_y > 0$, the reputation jumps upwards at an arrival of y. Absence of such a good news event is bad news, and reputation drifts down.

The Poisson component of the reputational dividend is thus the incremental probability μ_y of an event times the reputational increment $d_y x := \mu_y x(1-x)/(\mu_{H,y}x + \mu_{L,y}(1-x))$ at the arrival of this event

$$D_H(x) = \mathbb{E}_x[V_H(x + d_H x) - V_H(x + d_L x)]/dt = \sum_y \mu_y [V_H(x + d_y x) - V_H(x)]. \tag{4.4}$$

For good news signals this represents an increased probability of an upward jump, and for bad news a decreased probability of a downward jump. If market learning is imperfect $\mu_{\theta,y} \in (0,\infty)$, then the reputational increment $d_y x$ is dampened by a factor $x(1-x)$ and the Poisson component of the dividend declines to zero for extreme reputations, just like the Brownian component.

The most robust feature of equilibrium under imperfect learning is a shirk-region at the top. If the firm is believed to be working in an interval around $x = 1$, it is all but certain to have a high reputation in the future, undermining incentives to actually invest. To go beyond this local result and prove Theorem 2, we will focus on the case of sufficiently small costs c.

To concisely state Theorem 2 we say that the learning process dZ satisfies (HOPE), when a firm with any non-zero reputation has a chance of its reputation increasing even when it is believed to be shirking:

$$\Pr[dx > 0 | \tilde{\eta} = 0] > 0 \qquad \text{for all } x > 0 \tag{HOPE}$$

This holds if (a) there is a Brownian component, (b) there is a good news signal, or (c) there are only bad news signals but $-\sum_y \mu_y > \lambda$, so that the absence of bad news outweighs adverse equilibrium beliefs. Theorem 2 shows that, under (HOPE), equilibrium is *essentially unique*. That is, for any $\epsilon > 0$ any two equilibria η, η' are close in the metric defined in Section 2.2, i.e. $dist(\eta, \eta') < \epsilon$.

Theorem 2 *For any imperfect learning process, there exists $c > 0$ such that for all $c^* \in (0, c)$:*

(a) There exists a work-shirk equilibrium with cutoff $x^* \in (0,1)$.

(b) Reputational dynamics are convergent in such an equilibrium.

(c) With (HOPE), equilibrium is essentially unique.

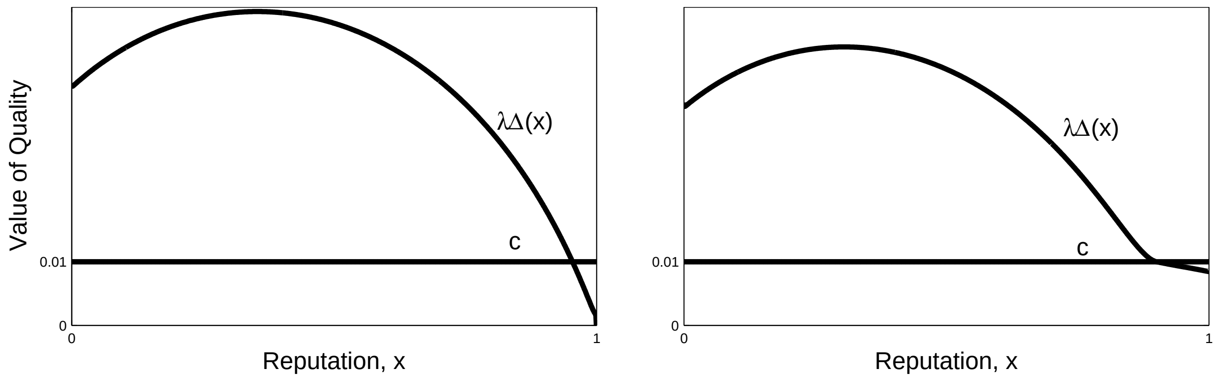

Figure 3: **Value of Quality under Full Effort (left) and in Work-Shirk Equilibrium (right).** This figure assumes that $\mu = 1$, $\lambda = 1$, $r = 1$ and $c = 0.01$. In the work-shirk equilibrium, the resulting cutoff is is $x^* = 0.900$.

Proof. See Appendix B. $\square$

Theorem 2 shows that, when costs are low, investment is sustainable at the bottom but not at the top. This fundamental asymmetry is illustrated in Figure 3 for the case of Brownian motion. Intuitively, when the firm is believed to be working, the value of quality is zero at $x = 1$ since current dividends are zero and, as the firm's reputation stays at $x = 1$, future dividends are zero. In contrast, the value of quality is positive at $x = 0$ since favorable equilibrium beliefs push the firm's reputation into the interior of $(0, 1)$, where dividends are high. The firm thus invests at $x = 0$ not because of the immediate reputational dividends, which are close to 0, but because of future dividends, when the firm's reputation is sensitive to actual quality.

Imperfect learning about low quality is a necessary condition for Theorem 2. Perfect bad news signals invalidate the argument for "shirking at the top" because reputation drops to zero at the arrival of perfect bad news. We study perfect bad news learning in Section 5.2. Imperfect learning about high quality, on the other hand, is merely assumed for symmetry and tractability in the proofs. In Section 5.1 we replicate the results of Theorem 2 for perfect good news.

Figure 3 is almost a proof of Theorem 2(a). Let $\Delta_{x^*}(x)$ be the value of quality for a firm with reputation x in a work-shirk profile with cutoff x^*. It can be shown that $\Delta_1(x)$ is strictly positive on $[0, 1)$, and monotonically decreasing on $[1 - \varepsilon, 1]$ with limit $\Delta_1(1) = 0$. Thus for small c there exists x^* such that

$$\lambda\Delta_1(x) \begin{cases} > c & \text{for } x < x^* \qquad \text{(Low reputations work)} \\ = c & \text{for } x = x^* \qquad (x^* \text{ is indifferent}) \\ < c & \text{for } x > x^* \qquad \text{(High reputations shirk).} \end{cases} \tag{4.5}$$

To prove part (a) we essentially want to replace Δ_1 on the left-hand-side with Δ_{x^*}. This argument assumes not only that $\Delta_{x^*}(\cdot)$ is close to $\Delta_1(\cdot)$, but also that $\Delta_{x^*}(\cdot)$ is monotonically decreasing close to 1, so that high reputations above cutoff x^* prefer to shirk. This step is not immediate. Suppose x^* is close to 1, and consider the marginal value of reputation $V'_\theta(x)$ to a firm with $x > x^*$. A reputational increment dx is valuable to the firm only as long as $x_t > x^*$: When $x_t = x^*$ the increment vanishes because of the convergent drift at x^*. Therefore, $V'_\theta(x)$ and $D_\theta(x)$ may be minimized at the cutoff x^*, and we need to take seriously the possibility that $\Delta_{x^*}(x)$ could be minimized at $x = x^*$ as well.

To show that $\lambda\Delta_{x^*}(x^*) > \lambda\Delta_{x^*}(x)$ for $x \in (x^*, 1]$, we need a better understanding of the reputational dynamics dx and the marginal values $V'_{\theta,x^*}(x)$ for $x, x^* \approx 1$. Assume for notational simplicity that learning is pure Brownian without Poisson signals. Then, the dynamics of $(1-x)$ approximate a geometric Brownian motion which is reflected at $(1-x^*)$ by the large relative difference in the drift terms. For the high quality firm,

$$
\begin{aligned}
d_H(1-x) &= -\lambda(\eta - x)\,dt - \mu_B^2 x(1-x)^2\,dt + \mu_B x(1-x)\,dW \\
&\approx \begin{cases} -\lambda(1-x)\,dt - \mu_B(1-x)\,dW & \text{for } x < x^* \\ \lambda x dt & \text{for } x > x^* \end{cases}
\end{aligned}
$$

and likewise for $d_L(1-x)$.

This has two implications. First, while the dividend may be minimized at x^*, the value of quality at the cutoff $\Delta_{x^*}(x^*)$ is largely determined by the dividends at $x < x^*$. Second, the marginal value of reputation and the dividend at $x > x^*$ are small in relation to those at $x < x^*$. This is because a reputational increment essentially disappears when $x_t = x^*$ and this happens much sooner for initial reputations $x_0 > x^*$ than for $x_0 < x^*$. Hence for $x > x^*$, $\Delta_{x^*}(x)$ is an average of low dividends while $x_t > x^*$, and a continuation value $\Delta_{x^*}(x^*)$ when x_t hits x^*. This average comes to less than $\Delta_{x^*}(x^*)$, as required.

Slow reputational updating at $x \approx 0$ and $x \approx 1$ suggests another, *shirk-work-shirk* type of equilibrium where a firm works when its reputation is between two cutoffs, $x \in [\underline{x}, \overline{x}]$, and shirks elsewhere. A firm with a low reputation is trapped in a lower "shirk-hole" in which market learning is too slow to incentivize investment, while a firm above $\underline{x}$ drifts towards $\overline{x}$.

Theorem 2(c) asserts that such shirk-work-shirk equilibria disappear for small costs. Intuitively, reputational dividends and the value of quality are bounded below on any interval $[\varepsilon, 1-\varepsilon]$, so when costs are small all intermediate reputations prefer to work. While reputational incentives in a shirk-work-shirk profile disappear for x close to zero, incentives are bounded below at the shirk-work cutoff $\underline{x}$, where working is profitable if the firm will escape from the lower shirk region with positive probability. That is, if (HOPE) is satisfied. This shows that with (HOPE) and small costs, the firm works at all low and intermediate levels of reputation $[0, 1-\epsilon]$, and shirks at some

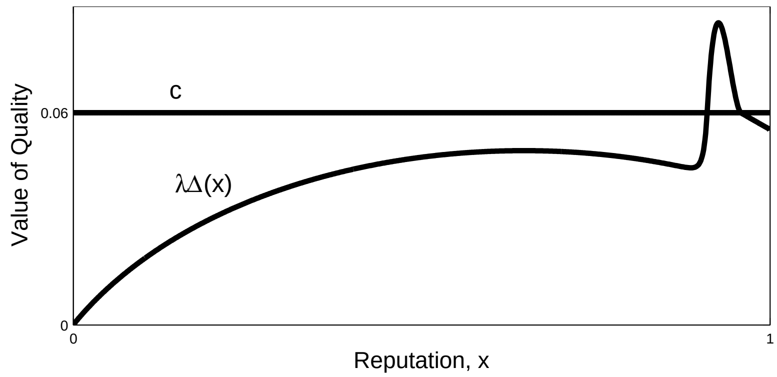

Figure 4: **Shirk-Work-Shirk Equilibrium.** This figure illustrates the value of quality in a work-shirk equilibrium, $\Delta_{x^*}(x)$. This figure assumes that $\mu = 1$, $\lambda = 1$, $r = 1$ and $c = 0.06$. The resulting cutoffs are $\underline{x} = 0.910$ and $\overline{x} = 0.958$.

high levels of reputation. All such equilibria are close in the metric $dist(\cdot, \cdot)$ and the equilibrium is essentially unique. [8]

If (HOPE) is violated the work-shirk equilibrium is not unique but coexists with a continuum of shirk-work-shirk equilibria. Intuitively, the failure of (HOPE) implies that a firm whose reputation drops into the shirk hole will remain there forever. This creates a discontinuity in the value function that incentivizes investment above the cutoff but not in the shirk-hole just below.[9,10]

In such an equilibrium, investment incentives are greatest just above the shirk-work cutoff where the divergent reputational drift makes value functions discontinuous. This captures the intuition that one product breakdown can put a reputable firm in the "hot-seat" where one more breakdown would finish off the firm. In such an equilibrium a firm that fails once will tries, but a firm that fails repeatedly gives up.

The above analysis relies on the assumption of low costs c to ensure work for intermediate reputations $x \in [\varepsilon, 1 - \varepsilon]$. For higher costs, numerical simulations indicate that shirk-work-shirk equilibria exist also in the pure Brownian case where (HOPE) holds (see Figure 4). As discussed above, investment incentives in this equilibrium (with $c = 0.06$) are much higher than in the above

[8] Beyond this we can show that there is a only one work-shirk equilibrium, but there may be other "work-shirk-work-shirk" equilibria with additional work-regions at the very top. For example, given the parameters in Figure 3, there is another equilibrium with working on $[0, 0.900]$, shirking on $[0.900, 0.944]$, working on $[0.944, 0.9605]$ and shirking on $[0.9605, 1]$.

[9] To construct shirk-work-shirk equilibria, we can first choose the lower, shirk-work cutoff low enough so as to discourage work in the shirk-hole, and then reapply the arguments in Appendix B to prove existence of the upper, work-shirk cutoff with the required properties.

[10] This argument implies that the equilibrium correspondence is not lower hemi-continuous. Learning processes with (HOPE) approximate learning processes without (HOPE), e.g. by taking the Brownian component of learning μ_B to zero. While equilibria along the convergent sequence are work-shirk, additional shirk-work-shirk equilibria discontinuously appear in the limit. These shirk-work-shirk equilibria are not close to the work-shirk equilibrium in the metric $dist(\cdot, \cdot)$.

work-shirk equilibrium (with $c = 0.01$ in Figure 3). A firm at the work-shirk cutoff has more to lose when a sequence of bad utility draws can push its reputation into a shirk-region, where it may be stuck forever. This argument makes it unlikely that these two types of equilibrium co-exist for the same parameter values.

5 Perfect Poisson Learning

In this Section we consider Poisson processes that can perfectly reveal the firm's quality. These cases are highly tractable and allow for an explicit equilibrium characterization, as shown in Appendix D. Moreover, these learning models are natural for many applications, making the results interesting in their own right. The perfect good news case illustrates and extends Theorem 2, in that equilibrium is work-shirk and unique. The perfect bad news case highlights the limitations of Theorem 2, in that equilibria are shirk-work and not unique.

5.1 Perfect Good News

Assume that consumers learn about quality θ_t from infrequent product *breakthroughs* that reveal high quality $\theta = H$ with arrival rate μ. Absent a breakthrough, updating evolves deterministically according to

$$\frac{dx}{dt} = \lambda(\eta(x) - x) - \mu x(1 - x). \tag{5.1}$$

Let x_t be the deterministic solution of the ODE (5.1) with initial value x_0.

The reputational dividend is the value of having a high quality in the next instant. This equals the value of increasing the reputation from x to 1 times the probability of a breakthrough:

$$D_H(x) = \mathbb{E}\left[V_H(x + d_H x) - V_H(x + d_L x)\right]/dt = \mu(V_H(1) - V_H(x)).$$

Using equation (3.2), the value of quality is:

$$\Delta(x_0) = \int_0^\infty e^{-(r+\lambda)t}\mu[V_H(1) - V_H(x_t)]dt \tag{5.2}$$

The reputational dividend $V_H(1) - V_H(x_t)$ is decreasing in x_t, so that $\Delta(x_0)$ is decreasing in x_0. Intuitively, a breakthrough that increases the firm's reputation to 1 is most valuable for a firm with a low reputation. Thus, investment incentives decrease in reputation and any equilibrium must be work-shirk.

The work-shirk beliefs imply that reputational dynamics converge to a cycle. Absent a breakthrough, the firm's reputation converges to a stationary point $\hat{x} = \min\{\lambda/\mu, x^*\}$ where the firm works with positive probability. When a breakthrough occurs, the firm's reputation jumps to 1. The firm is then believed to be shirking, so its reputation drifts down to $\hat{x}$, absent another break-

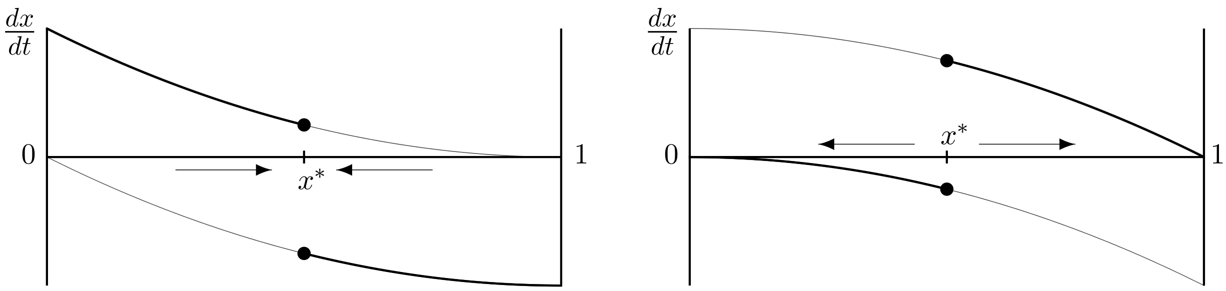

Figure 5: **Reputational dynamics in Good News (left) and Bad News (right).** This figure illustrates how the reputational drift dx/dt, absent a breakthrough, changes with the reputation of the firm, x. These pictures assume $\lambda = \mu$ and $x^* = 1/2$. The dark line shows the drift in equilibrium.

through. In the long-run, the firm's reputation therefore cycles over the range $[\hat{x}, 1]$. In case that $\lambda \geq \mu$, the firm's reputation drifts up whenever it is believed to be working (see Figure 5). Here, the dynamics are stationary at $\hat{x} = x^*$, at which point the firm chooses to work with intensity $\eta(x^*) = x^* \left(1 + \frac{\mu}{\lambda}(1 - x^*)\right)$.

Proposition 1 *Under perfect good news learning:*

(a) Every equilibrium is work-shirk.

(b) Reputational dynamics converge to a non-trivial cycle.

(c) If $\lambda \geq \mu$, the equilibrium is unique.

Proof. Part (a). Reputation x_t follows (5.1), so an increase in x_0 raises x_t at each point in time. Lemma 2 says that $V_H(x)$ is strictly increasing in x, so equation (5.2) implies that $\Delta(x_0)$ is decreasing in x_0. Part (b) follows from (a).

Part (c). Given $\lambda \geq \mu$, the process x_t starting at $x_0 = 1$, falls until it becomes stuck at x^*. Equation (D.4) in Appendix D.1 derives a closed-form expression for $\Delta_{x^*}(x^*)$. One can verify that $\Delta_{x^*}(x^*)$ is decreasing in the cutoff, implying the equilibrium is unique. □

To understand the uniqueness result of Proposition 1(c), suppose the market believes the cutoff is $\tilde{x}$, and denote the firm's best response by $x^*(\tilde{x})$. An increase in $\tilde{x}$ means the firm's reputation will not drift down as far, absent a breakthrough. This change benefits low-quality firms more than high-quality firms, reducing $\Delta(x)$. As a result, $x^*(\tilde{x})$ is decreasing in $\tilde{x}$ and there is a unique fixed point where $x^*(\tilde{x}) = \tilde{x}$.

This result corroborates Theorem 2: Generally, investment incentives are stronger at the bottom than at the top; with perfect good news, incentives are monotonically decreasing.

5.2 Perfect Bad News

Assume that x_t is generated by *breakdowns* that reveal low quality $\theta = L$ with arrival rate μ. Absent a breakdown, updating evolves deterministically according to

$$\frac{dx}{dt} = \lambda(\eta(x) - x) + \mu x(1 - x). \tag{5.3}$$

Let x_t be the deterministic solution of ODE (5.3) with initial value x_0.

The reputational dividend is the value of having a high quality in the next instant. Quality insures the firm against a breakdown, and the reputational dividend thus equals:

$$D_L(x) = \mathbb{E}\left[V_L(x_t + d_H x_t) - V_L(x_t + d_L x_t)\right]/dt = \mu(V_L(x_t) - V_L(0)).$$

Using equation (3.3), the value of quality is:

$$\Delta(x_0) = \int_0^\infty e^{-(r+\lambda)t}\mu[V_L(x_t) - V_L(0)]dt. \tag{5.4}$$

The jump size $V_L(x_t) - V_L(0)$ is increasing in x_t, so that $\Delta(x)$ is increasing in x. Intuitively, a breakdown that destroys the firm's reputation entirely is most damaging for a firm with a high reputation. Thus, incentives to invest increase in reputation and equilibrium is shirk-work.

The shirk-work beliefs imply that the reputational dynamics diverge. Consider an equilibrium where the firm works if its reputation is below $x^* \in (0, 1)$, and works if it is above x^*. A firm that starts with reputation above x^* converges to reputation $x = 1$, absent a breakdown. If the firm is hit by such a breakdown while its product quality is still low, it gets stuck in a shirk-hole with reputation $x = 0$. A firm with reputation below x^* initially shirks and may have either rising or falling reputation, depending on parameters. In either case, its reputation will either end up at $x = 0$ or $x = 1$.

To allow for positive investment in equilibrium we impose the following assumption:

$$\frac{\lambda}{r + \lambda + \mu}\mu(1 - c)/r > c \tag{5.5}$$

Proposition 2 *Under perfect bad news learning:*

(a) Every equilibrium is shirk-work.

(b) If $x^* \in (0, 1)$ then reputational dynamics diverge to 0 or 1.

(c) Assume (5.5) holds and $\lambda \geq \mu$. There is a non-empty interval $[a, b]$ such that every cutoff $x^* \in [a, b]$ defines an equilibrium.

Proof. Part (a). Reputation x_t follows (5.3), so an increase in x_0 raises x_t at each point in time. Lemma 2 says that $V_L(x)$ is strictly increasing in x, so equation (5.4) implies that $\Delta(x_0)$ is increasing in x_0. Part (b) follows from (a).

Part (c). If $\lambda \geq \mu$, the dynamics are divergent at x^*: if $x_0 = x^* - \epsilon$, then $\lim x_t = 0$; if $x_0 = x^* + \epsilon$, then $\lim x_t = 1$. Thus, to define value functions and Δ at the cutoff x^* we need to specify whether or not x^* works. Denote by $\Delta^-_{x^*}(x)$ (resp. $\Delta^+_{x^*}(x)$) the value of quality at x when x^* is believed to be shirking (resp. working). For $x^* \in (0,1)$ it follows that $\Delta^-_{x^*}(x^*) = \lim_{x \nearrow x^*} \lambda \Delta_{x^*}(x)$ and $\Delta^+_{x^*}(x^*) = \lim_{x \searrow x^*} \lambda \Delta_{x^*}(x)$. Lemma 2 says that $V_L(x_t)$ is strictly increasing in x_t, so (5.4) implies that

$$\Delta^-_{x^*}(x^*) < \Delta^+_{x^*}(x^*). \tag{5.6}$$

A cutoff $x^* \in (0,1]$ then defines a shirk-work equilibrium if and only if[11]

$$\lambda \Delta^-_{x^*}(x^*) \leq c \leq \lambda \Delta^+_{x^*}(x^*). \tag{5.7}$$

Equation (5.4) implies that $\Delta^+_{x^*}(x^*)$ and $\Delta^-_{x^*}(x^*)$ are increasing and continuous in x^*. It remains to be shown that the interval of x^* satisfying (5.7) is not empty. The lower bound is always satisfied because a firm with no reputation that is believed to be shirking is stuck at 0 forever, and thus $\lambda \Delta^-_0(0) = 0$. For the upper bound, $\lambda \Delta^+_1(1) = \frac{\lambda}{r+\lambda+\mu}\mu(1-c)/r$ because $V_L(1) = \frac{r+\lambda}{r+\lambda+\mu}(1-c)/r$, $V_L(0) = 0$, and $\lambda \Delta^+_1(1) = \frac{\lambda\mu}{r+\lambda}\left(V_L(1) - V_L(0)\right)$. Under assumption (5.5), equation (5.7) therefore defines a non-empty interval of cutoffs, $[a, b]$. □

Suppose $\lambda \geq \mu$, so that whenever the firm is known to be shirking its reputation drifts down (see Figure 5). In this case, the region below x^* is a shirk-hole: when a firm's reputation is below the cutoff, it is certain to see its reputation decrease because of the adverse beliefs. Such a firm always shirks, eventually giving rise to a low quality product and a product breakdown destroying whatever is left of its reputation. When a firm's reputation is above the cutoff, favorable market beliefs contribute to an increasing reputation and the firm invests to insure itself against a product breakdown. At the cutoff, the firm works if it is believed to be working and shirks if it is believed to be shirking.

Theorem 2(c) shows that there is an interval of equilibrium cutoffs satisfying (5.7). The multiplicity is driven by a discontinuity in the value function at x^*, caused by the divergent reputational dynamics. Intuitively, the market's beliefs become self-fulfilling. If the market believes the firm is shirking, its reputation falls and investment incentives are low. Conversely, if the market believes the firm is working, its reputation rises and incentivizes the firm to invest in order to protect its

[11] For $x^* = 0$, condition (5.7) is necessary and sufficient for full work to be an equilibrium. However, working for $x > 0$ but shirking at $x = 0$ can be in equilibrium if (5.7) is not satisfied, but $c > \lambda \Delta^+_0(0)$. The necessary equilibrium condition is that $c \leq \lambda \Delta^-_0(x)$ for all $x > 0$, or equivalently that $c \leq \lambda \Delta^+_x(x)$ for all $x > 0$, which is weaker than (5.7). The reason for this is that $\Delta^+_x(x)$ is discontinuous in the limit $\lim_{x \to 0} \lambda \Delta^+_x(x) > \Delta^+_0(0)$, because quality is more valuable when there is a shirk-hole at $x = 0$ than it is under full work.

appreciating reputation.[12]

When $\lambda < \mu$ the dynamics have additional interesting features: Define $\widehat{x} = 1 - \frac{\lambda}{\mu} \in (0,1)$ to be the stationary point in the dynamics when the firm is believed to be shirking. There are two types of equilibria:

1. *Trapped equilibria.* When $\widehat{x} < x^*$, a firm with reputation $x \in (0, x^*)$ finds its reputation converging to $\widehat{x}$, and remains stuck in a shirk-hole. At some point it suffers a breakdown and remains at $x = 0$ thereafter. Since the dynamics are divergent at x^*, the value function is discontinuous and there is a continuum of such equilibria.

2. *Permeable equilibria.* When $\widehat{x} > x^*$, a firm with reputation $x \in (0, x^*)$ finds its reputation increasing. If x_t passes x^* before a breakdown hits, the firm starts to work and its reputation may converge to one. Since the value functions are continuous at a permeable cutoff x^*, there is at most one permeable equilibrium.

5.3 Perfect Good and Bad News

If the product can both enjoy breakthroughs revealing high quality with intensity μ_g, and suffer breakdowns revealing low quality with intensity $-\mu_b$, the reputational dividend is given by:

$$\begin{aligned} D_\theta(x) &= \mu_g (V_\theta(1) - V_\theta(x)) - \mu_b (V_\theta(x) - V_\theta(0)) \\ &= -(\mu_b + \mu_g) V_\theta(x) + \mu_g V_\theta(1) + \mu_b V_\theta(0). \end{aligned} \tag{5.8}$$

When good news is more frequent than bad news, $\mu_g > -\mu_b$, the analysis is similar to the perfect good news case of Section 5.1: Reputational dividends and value of quality are decreasing in reputation and any equilibrium must be work-shirk. However, full work can be an equilibrium now, because the perfect bad news signal incentivizes investment even for high reputations.

When bad news is more frequent than good news, $-\mu_b > \mu_g$, the analysis is similar to the perfect bad news case of Section 5.2: Reputational dividends and value of quality are increasing in reputation, and any equilibrium must be shirk-work. However, when costs are low, equilibrium must be full work, because the perfect good news signal guarantees that (HOPE) is satisfied and rules out shirking at the bottom.

[12] The shirk-work equilibrium under perfect bad news learning seems to be at odds with the unique work-shirk equilibrium under "almost perfect" bad news covered in Theorem 2. However, this limit is continuous since full work is an equilibrium under perfect bad news learning and low costs (see equation (D.14) in Appendix D.2). This full work profile is approximated by work-shirk equilibria when learning is via bad news signals y with high but finite ratios $\mu_{L,y}/\mu_{H,y}$. With such signals the reputational increment $d_y x$ and the reputational dividend remain high close to $x = 1$. Thus, the work-shirk equilibria with $x^* \approx 1$ for close to perfect learning approximate the full work equilibrium, with $x^* = 1$, under perfect learning.

6 Quality Choice

We now connect our analysis to the standard repeated games models, where quality is chosen in every period, by taking the quality obsolescence rate λ to infinity. As discussed in the introduction, we find that an increase in λ can be detrimental to incentives despite the benefit of front-loading the returns to investment. In particular, we find that the work-shirk equilibria of Theorem 2 disappear as $\lambda \to \infty$.

Theorem 3 *For any learning process, there exists λ^* such that for all $\lambda > \lambda^*$:*

(a) Pure shirking is an equilibrium.

(b) There is no work-shirk equilibrium with cutoff $x^* \in (0, 1]$.

Proof. See Appendix C. $\square$

Intuitively, when we fix a learning process dZ_t and choose λ sufficiently high, the reputational dynamics dx_t in (2.2) are dominated by equilibrium beliefs $\lambda(\widetilde{\eta} - x)$ and any effect of learning from actual quality is quickly lost. If beliefs $\widetilde{\eta}$ are work-shirk the firm's future reputation and revenue will be close to the cutoff x^*, irrespective of its investment, quality, or induced consumer utility.

More precisely, we can apply Theorem 1 to write the benefits of investment as the product of an average reputational dividend and an average time at which the dividend accrues:

$$\lambda\Delta(x) = \lambda \int_0^\infty e^{-(r+\lambda)t}\mathbb{E}[D(x_t)]dt \approx \underbrace{\frac{\lambda}{r+\lambda}}_{\leq 1} \underbrace{D(x_{future})}_{\to 0}$$

The first term $\frac{\lambda}{r+\lambda}$ captures the front-loading effect discussed in the introduction: As λ increases so does the probability that current investment affects quality. At the same time it increases the rate at which this quality becomes obsolete, and the effective discount rate $r+\lambda$. Aggregating these two effects, an increase in λ front-loads returns and increases $\lambda/(r+\lambda)$. Under complete information, this is the only effect when increasing λ. Under incomplete information, there is a second, negative effect that outweighs the benefits of front-loading. The reputational dividend of quality vanishes as a function of λ. To see this, we decompose the value of a firm, say that is shirking and has low quality, into current profits and continuation value $V_L(x) = x_t dt + (1 - rdt) V_L(x - dx)$. Rearranging terms, we find that marginal value $V_L'(x) = \frac{1}{dx/dt}(x_t - rV_L(x))$ is decreasing in $dx/dt \approx \lambda$. As the marginal value of reputation vanishes, so does the reputational dividend of quality, which is composed of jump-terms $\mu_y(V_\theta(x_y) - V_\theta(x))$ derived from the Poisson component of learning, and a continuous term $\mu_B^2 x(1-x)V_\theta'(x)$ derived from the Brownian component of learning.[13]

For the perfect Poisson learning processes of Section 5, we find:

Proposition 3 *There exists λ^* such that for all $\lambda > \lambda^*$:*

(a) Under perfect good news learning, full shirking is the only equilibrium.

(b) Under perfect bad news learning, any cutoff $x^* \in (0,1]$ defines a shirk-work equilibrium, if condition (5.5) is satisfied.

Proof. Part (a). By Proposition 1 any equilibrium under good news is work-shirk, but the only such equilibrium for high λ entails full shirking by Theorem 3. More explicitly, with perfect good news λ decreases incentives by decreasing the value of perfect reputation. Starting at $x_0 = 1$ reputation decreases exponentially, $x_t \leq e^{-\lambda t}$, so

$$V_L(1) = \int_0^\infty e^{-rt} x_t dt \leq \frac{1}{r+\lambda}$$

which disappears as $\lambda \to \infty$. Then $\Delta_0(0)$ disappears as well:

$$\lambda \Delta_0(0) = \lambda \int_0^\infty e^{-(r+\lambda)s} \mu \mathbb{E}_{\theta^s = H}\left[V_L(1) - V_L(x_s)\right] ds \leq \frac{\lambda\mu}{r+\lambda} V_L(1) \leq \frac{\lambda}{r+\lambda}\frac{\mu}{r+\lambda}.$$

This completes the proof, because incentives are greatest for $x = x^* = 0$ by equation (D.4).

Part (b). Pick any $x^* > 0$. First, suppose $x_0 > x^*$ and observe that $x_t \geq 1 - e^{-\lambda t}$. Equation (D.9) in Appendix D.2 gives an explicit formula for the value of quality. Taking limits,

$$\lim_{\lambda\to\infty} \lambda \Delta_{x^*}(x) \geq \lim_{\lambda\to\infty} \frac{\lambda\mu}{\lambda+\mu} \int_{t=0}^\infty e^{-rt}(1 - e^{-\lambda t} - c)(1 - e^{-(\lambda+\mu)t})\, dt = \mu(1-c)/r$$

where the final equality uses the fact that the integral converges to $(1-c)/r$. Assumption (5.5) implies that $\mu(1-c)/r > c$. Hence for sufficiently large λ, working is optimal for all $x > x^*$ and any x^*.

[13] Theorem 3 does not show that full shirking is the unique equilibrium in the Brownian case. In this sense, this result is weaker than the one in Sannikov and Skrzypacz (2007). However the following informal argument suggests that the stronger result is also true in our setting.

Consider the incentives in a shirk-work-shirk equilibrium with cutoffs $\underline{x} < \overline{x}$ for high values of λ. A firm with reputation $\overline{x}$ faces the risk of slipping through the work-region into the shirk hole. Let γ_θ be the intensity with which this happens to a firm with quality θ. The dividend of quality is now a function of the difference in these intensities, $\gamma_L - \gamma_H$. For large values of λ, the size of the work-region $\overline{x} - \underline{x}$ must be small to bound this difference from below. The analysis in Sannikov and Skrzypacz (2007) suggests that this increases γ_H. Even a firm with high quality slips into the shirk hole arbitrarily fast, and $V(\overline{x})$ and $\Delta(\overline{x})$ converge to 0.

Next suppose $x_0 < x^*$ and observe that $x_t \leq e^{-(\lambda-\mu)t}$. Equation (D.12) in Appendix D.2 derives an explicit formula for the value of quality. Taking limits,

$$\lim_{\lambda\to\infty} \lambda\Delta_{x^*}(x) \leq \lim_{\lambda\to\infty} \frac{\lambda\mu}{\lambda-\mu}\int_{t=0}^{\infty} e^{-rt}e^{-(\lambda-\mu)t}dt = 0$$

Hence for sufficiently large λ, shirking is optimal for all $x < x^*$ and any x^*. $\square$

The key difference that sustains investment in part (b) is that reputational dynamics are divergent at a shirk-work cutoff x^*, and value functions are discontinuous as discussed in Section 5.2. Marginal value then vanishes, the value functions approach step functions, and the reputational dividend equals the discontinuity $\mu\left(V_\theta\left(1\right) - V_\theta\left(0\right)\right)$ at x^*; this exceeds c under condition (5.5) as shown in Appendix D.2. High values of λ amplify the multiplicity of equilibria found in the bad news case. Intuitively, a firm that starts below the cutoff finds its reputation falling to 0 instantly and gives up, while a firm above the cutoff finds its reputation rising to 1 instantly and works to stay there.

Theorem 3 has a surprising consequence: Providing more information about the firm's quality may be detrimental to equilibrium investment. More specifically, consider a shirk-work equilibrium under perfect bad news learning. Suppose we improve the learning process by introducing additional perfect good news signals. If the arrival rate of the good news exceeds the arrival rate of the bad news, Section 5.3 shows that equilibria must be work-shirk. If additionally λ is high enough, full shirking is the only equilibrium by Lemma 3.

Ordering signal structures by sufficiency, equilibria by investment, and equilibrium sets by the set order we can summarize:

Corollary 2 *More information can lead to less investment.*

The economics of the counter-example resembles the problem of renegotiation-proofness in a repeated prisoners' dilemma: Under perfect bad news learning a firm with a high reputation works because a breakdown permanently destroys its reputation. Additional good news signals grant the firm a second chance after a breakdown, and undermine incentives to work hard in the first place.

7 Conclusion

This paper studies the moral hazard problem of a firm that produces experience goods and controls quality through its investment choice. Investment is incentivized by consumers' learning about product quality which feeds into the firm's reputation and future revenue.

The key feature distinguishing our paper from classical models of reputation and repeated games is that product quality is a function of past investments rather than current effort. This capital-theoretic model of persistent quality is realistic: The current state of General Motors is a function of its past hiring policies, investment decisions and reorganisations, all of which are endogenous and have lasting effects on quality. The model also yields new economic insights: When the market learns quality via breakthroughs of high quality products, a low-reputation firm has strong incentives to build its reputation, while a high reputation firm has weak incentives to maintain its reputation and chooses to run it down. When the market learns quality via breakdowns of low quality products, a low-reputation firm has weak incentives to invest into quality and reputation, while a high reputation firm keeps investing to protect its reputation.

There are many ways to extend this model to capture additional important aspects of firm reputation. One can allow high-quality firms to have lower investment costs than low-quality firms. For example with perfect good news, equilibria are still work-shirk, with two cutoffs $x_L^* < x_H^*$ for low and high quality firms, implying that the firm's reputation ultimately cycles over $[x_L^*, 1]$. In another variation of the model, the firm goes out of business when its reputation falls below an exogenous threshold, $\underline{x}$. The value functions then satisfy $V_L(\underline{x}) = V_H(\underline{x}) = 0$; with a Brownian learning component the firm shirks when its reputation gets close to $\underline{x}$ as its life expectancy gets short.

Beyond firm reputation, we hope that our model will prove useful in other fields. In corporate finance, where default is a bad news signal for a borrower, the shirk-work equilibria generate endogenous credit-traps. In political economy, where a scandal is bad news about a politician, the divergent dynamics imply that a politician who is caught will cheat even more, whereas a lucky politician will become more honest. And in personnel economics, our model predicts that in "superstar markets", where agents are judged by their successes, performance tends to be mean-reverting.

A General Results

A.1 Log-likelihood Ratio Transformation

For most of the technical proofs in the appendix, we represent reputation by the log-likelihood ratio $\ell(x) = \log(x/(1-x)) \in \mathbb{R} \cup \{-\infty, \infty\}$ of state H rather than the posterior x. The relevant transformation functions are:

$$\ell(x) = \log \frac{x}{1-x} \qquad x(\ell) = \frac{e^\ell}{1+e^\ell} \qquad \frac{dx}{d\ell} = \frac{e^\ell}{(1+e^\ell)^2} = x(1-x)$$

When we work in ℓ-space we write $\widehat{V}_\theta(\ell) := V_\theta(x(\ell))$ for value functions, $\widehat{D}_\theta(\ell) := D_\theta(x(\ell))$ for the reputational dividend, $\widehat{\Delta}(\ell) := \Delta(x(\ell))$ for the value of quality, and $\widehat{\eta}(\ell) := \eta(x(\ell))$ for investment. The advantage of this transformation is that reputational updating, equation (2.2), is more tractable in ℓ-space.

Market Learning: Bayesian updating from signals dZ_t is linear in ℓ-space:

$$\ell_{t+dt} = \ell_t + \log \frac{\Pr(dZ_t|H)}{\Pr(dZ_t|L)}. \tag{A.1}$$

More specifically, the Brownian component of learning $\theta\mu_B dt + dW$ imposes a Brownian motion on reputational dynamics:

$$d_\theta \ell = \mu_B^2 (\theta - 1/2) + \mu_B dW$$

The Poisson component of learning, with event y arriving at intensity $\mu_{y,\theta}$, affects reputational dynamics via

$$d\ell = \begin{cases} \delta_y & \text{in case of arrival } y, \\ -\sum_y \mu_y dt & \text{absent an arrival,} \end{cases}$$

where the jump equals $\delta_y = \log(\mu_{y,H}/\mu_{y,L})$.

Equilibrium Beliefs: This part of reputational updating is more complex in ℓ-space:

$$\frac{d\ell}{dt} = \frac{d\ell}{dx}\lambda(\widetilde{\eta} - x) = \lambda \frac{(1+e^\ell)^2}{e^\ell}\left(\widetilde{\eta} - \frac{e^\ell}{1+e^\ell}\right) = \begin{cases} \lambda\left(1+e^{-\ell}\right) & \text{for } \widetilde{\eta} = 1, \\ -\lambda\left(1+e^{\ell}\right) & \text{for } \widetilde{\eta} = 0. \end{cases} \tag{A.2}$$

Asymptotic Dynamics: Reputational dynamics are generally determined by equations (A.1) and (A.2). In a work-shirk profile with cutoff $\ell^* \gg 0$, dynamics of high reputations $0 \ll \ell \leq \ell^*$

are approximately independent of ℓ with

$$d_\theta \ell = \begin{cases} \left(\lambda\left(1+e^{-\ell}\right)+\mu_B^2\left(\theta-1/2\right)-\sum_y \mu_y\right)dt & \text{almost constant drift,} \\ +\mu_B dW & \text{constant wiggle,} \\ +\sum_{y\in Y}\delta_y & \text{constant jumps,} \end{cases} \tag{A.3}$$

while the boundary ℓ^* is approximately reflecting with $d\ell \approx -\lambda e^{\ell} \approx -\infty$.

Reputational Dividend: $\widehat{D}_\theta(\ell)$ is additive across the Brownian and Poisson component of market learning:

$$\widehat{D}_\theta(\ell) = \frac{\mu_B^2}{2}\widehat{V}_\theta'(\ell) + \sum_y \mu_y \left(\widehat{V}_\theta(\ell+\delta_y) - \widehat{V}_\theta(\ell)\right). \tag{A.4}$$

Value of Quality: The representation of $\widehat{\Delta}(\ell)$ as sum of future reputational dividends is the same in x-space and ℓ-space:

$$\widehat{\Delta}(\ell) = \int_0^\infty e^{-(r+\lambda)t}\mathbb{E}_{\ell_0=\ell,\theta^t=L}[\widehat{D}_H(\ell_t)]dt.$$

A.2 Value of Reputation: Reprise

We now extend Lemma 2 (positive marginal value of reputation in equilibrium) in two directions. Lemma 4 extends the result to non-equilibrium work-shirk profiles and provides an explicit formula for the marginal value of reputation (equation A.5). Lemma 5 provides a lower bound to incremental value that is uniform across equilibrium investment profiles.

We first prove an auxiliary lemma that uses flexible accounting in representing value functions as NPV of future profits. Future profits depend on current quality and investment via the reputational evolution. Lemma 3 shows one way of controlling this evolution, by replacing the actual investment function η by an arbitrary other, not necessarily Markovian, investment function $\overline{\eta}\left(Z^t\right)$.

Lemma 3 *For any investment and belief function $\langle \widehat{\eta}, \widetilde{\eta}\rangle$, and any (non-Markovian) alternative investment function $\overline{\eta}\left(Z^t\right)$ the firm's value equals:*

$$\widehat{V}_\theta(\ell_0) = \mathbb{E}_{\overline{\eta},\widetilde{\eta}}\left[\int_0^\infty e^{-rt}\left(\left(x\left(\ell_t\right)-c\overline{\eta}\left(Z^t\right)\right)+\left(\widehat{\eta}(\ell_t)-\overline{\eta}\left(Z^t\right)\right)\left(\lambda\widehat{\Delta}(\ell_t)-c\right)\right)dt\right]$$

Proof. Fix θ_0, ℓ_0 and $\widetilde{\eta}$. Consider first a "one shot deviation" from $\widehat{\eta}$, i.e. an alternative investment function $\overline{\eta}$ that differs from $\widehat{\eta}$ only for $t \in [0, dt]$, say $\overline{\eta} = 1$ while $\widehat{\eta} = 0$. A firm that invests according to $\widehat{\eta}$ but whose quality θ_{dt} is governed by $\overline{\eta}$ gains $\lambda\widehat{\Delta}(\ell_0)dt$. Thus, the firm's actual value is the value under the more favorable process $\overline{\eta}$, minus the fair value of the quality subsidy:

$$\widehat{V}_\theta(\ell_0) = \mathbb{E}_{\overline{\eta},\widetilde{\eta}}\left[\int_0^\infty e^{-rt}\left(x\left(\ell_t\right) - c\widehat{\eta}(\ell_t)\right)dt\right] - \lambda\widehat{\Delta}(\ell_0)dt.$$

For "multi-period" deviations, we accumulate a term $\left(\widehat{\eta}(\ell_t) - \overline{\eta}\left(Z^t\right)\right)\lambda\widehat{\Delta}(\ell_t)dt$ whenever $\widehat{\eta}(\ell_t) \neq \overline{\eta}\left(Z^t\right)$. Thus, in general we have

$$\begin{aligned}
\widehat{V}_\theta(\ell_0) &= \mathbb{E}_{\overline{\eta},\widetilde{\eta}}\left[\int_0^\infty e^{-rt}\Big(\left(x\left(\ell_t\right) - c\widehat{\eta}(\ell_t)\right) + \left(\widehat{\eta}(\ell_t) - \overline{\eta}\left(Z^t\right)\right)\lambda\widehat{\Delta}(\ell_t)\Big)dt\right] \\
&= \mathbb{E}_{\overline{\eta},\widetilde{\eta}}\left[\int_0^\infty e^{-rt}\Big(\left(x\left(\ell_t\right) - c\overline{\eta}\left(Z^t\right)\right) + \left(\widehat{\eta}(\ell_t) - \overline{\eta}\left(Z^t\right)\right)\Big(\lambda\widehat{\Delta}(\ell_t) - c\Big)\Big)dt\right]
\end{aligned}$$

as required. □

To state and prove the next lemma, we write the firm's reputation at time t as a function $\ell_t = \ell_t\left(\ell_0, Z^t, \widetilde{\eta}\right)$ of its initial reputation ℓ_0, realized utilities Z^t and the Markovian beliefs $\widetilde{\eta}$.

Lemma 4 *Fix a work-shirk profile with cutoff ℓ^*:*

(a) Reputational increments are decreasing:

$$\frac{\partial\ell_t}{\partial\ell}(\ell, Z^t, \widetilde{\eta}) < 1.$$

(b) Value $\widehat{V}_\theta\left(\ell\right)$ is continuous in reputation ℓ.

(c) If the cutoff ℓ^* weakly prefers to shirk, i.e. $\lambda\widehat{\Delta}_{\ell^*}\left(\ell^*\right) \leq c$, the marginal value of reputation is strictly positive:

$$\widehat{V}_\theta'\left(\ell\right) > 0 \textit{ for all } \ell \in \mathbb{R}.$$

(d) If the cutoff ℓ^* is indifferent, i.e. $\lambda\widehat{\Delta}_{\ell^*}\left(\ell^*\right) = c$, the marginal value of reputation equals:

$$\widehat{V}_\theta'\left(\ell\right) = \int_{t=0}^\infty e^{-rt}\mathbb{E}_{\theta_0=\theta}\left[\frac{e^{\ell_t}}{\left(1+e^{\ell_t}\right)^2}\frac{\partial\ell_t}{\partial\ell}(\ell, Z^t, \widetilde{\eta})\right]dt > 0 \textit{ for all } \ell \in \mathbb{R}. \quad \text{(A.5)}$$

(e) The value of quality at ℓ^* is strictly positive:

$$\widehat{\Delta}_{\ell^*}\left(\ell^*\right) > 0.$$

Proof. For part (a), consider the reputational trajectories ℓ_t, ℓ_t' originating at $\ell < \ell'$. Equations (A.1) and (A.2) imply that $\ell_t' - \ell_t$ is decreasing in t for all work-shirk profiles: Market learning (A.1) shifts ℓ_t and ℓ_t' by the same amount $\log\left(\Pr(dZ_t|H)/\Pr(dZ_t|L)\right)$, while equilibrium beliefs (A.2) shrink $\ell_t' - \ell_t$ at rate $\lambda\left(1+e^{-\ell_t'}\right) - \lambda\left(1+e^{-\ell_t}\right) \approx -\lambda e^{-\ell_t}\left(1-e^{-(\ell_t'-\ell_t)}\right) < 0$ in the work-region and similarly in the shirk-region.

Turning to points (b) to (e), let $\overline{\eta}\left(Z^t\right) = \widehat{\eta}(\ell_t(\ell', Z^t, \widetilde{\eta}))$ be the investment strategy of a firm with initial reputation ℓ who mimics ℓ', expressed in a non-Markovian way directly as a function of market signals Z^t. We decompose the incremental value of reputation as follows:

$$\widehat{V}_{\theta,\widehat{\eta}}\left(\ell'\right) - \widehat{V}_{\theta,\widehat{\eta}}\left(\ell\right) = [\widehat{V}_{\theta,\widehat{\eta}}(\ell') - \widehat{V}_{\theta,\overline{\eta}}(\ell)] + [\widehat{V}_{\theta,\overline{\eta}}(\ell) - \widehat{V}_{\theta,\widehat{\eta}}(\ell)] \tag{A.6}$$

The first term in (A.6) is the reputational advantage of starting with reputation $\ell' > \ell$, when the firm starting at reputation ℓ mimics the investment of the firm starting at ℓ'. It is determined by the derivative of future reputation with respect to current reputation:

$$\widehat{V}_{\theta,\widehat{\eta}}(\ell') - \widehat{V}_{\theta,\overline{\eta}}(\ell) = \int e^{-rt}\mathbb{E}_{\theta_0=\theta,\overline{\eta}}\left[\left(x\left(\ell_t(\ell', Z^t, \widetilde{\eta})\right) - x\left(\ell_t\left(\ell, Z^t, \widetilde{\eta}\right)\right)\right)\right] dt$$

This term is always positive. Taking the limit as $\ell' \to \ell$ and applying the chain rule gives rise to equation (A.5).

The second term in (A.6) is the net value of shirking whenever $\overline{\eta}\left(Z^t\right) = 0$ and $\widehat{\eta}(\ell_t) = 1$. We will now show that it is of order $o\left(d\ell\right)$.

$$\begin{aligned}
\widehat{V}_{\theta,\overline{\eta}}(\ell) - \widehat{V}_{\theta,\widehat{\eta}}(\ell) &= \mathbb{E}_{\ell_0=\ell,\theta_0=\theta,\overline{\eta},\widetilde{\eta}}\left[\int_0^\infty e^{-rt}\left(\widehat{\eta}(\ell_t) - \overline{\eta}\left(Z^t\right)\right)\left(c - \lambda\widehat{\Delta}(\ell_t)\right) dt\right] \\
&= \mathbb{E}_{\theta_0=\theta,\overline{\eta}}\left[\int_0^\infty e^{-rt}\left(\widehat{\eta}(\ell_t(\ell, Z^t, \widetilde{\eta})) - \widehat{\eta}(\ell_t(\ell', Z^t, \widetilde{\eta}))\right)\left(c - \lambda\widehat{\Delta}(\ell_t(\ell, Z^t, \widetilde{\eta}))\right) dt\right] \\
&= \mathbb{E}_{\theta_0=\theta,\overline{\eta}}\left[\int_{\ell_t(\ell, Z^t, \widetilde{\eta}) < \ell^* < \ell_t(\ell', Z^t, \widetilde{\eta})} e^{-rt}\left(c - \lambda\widehat{\Delta}(\ell_t(\ell, Z^t, \widetilde{\eta}))\right) dt\right] \\
&\leq \max_{\widetilde{\ell}\in[\ell^*, \ell^*+(\ell'-\ell)]}\left\{\left|c - \lambda\widehat{\Delta}\left(\widetilde{\ell}\right)\right|\right\}\left(\ell' - \ell\right)/2\lambda
\end{aligned}$$

The first equality applies Lemma 3. The second applies the definition of $\overline{\eta}\left(Z^t\right)$. The third uses that the investment functions disagree if and only if the trajectories are on opposite sides of the cutoff, i.e. $\ell_t(\ell, Z^t, \widetilde{\eta}) < \ell^* < \ell_t(\ell + d\ell, Z^t, \widetilde{\eta})$. The final inequality uses that by (a), $\ell_t(\ell + d\ell, Z^t, \widetilde{\eta}) - \ell_t(\ell, Z^t, \widetilde{\eta})$ is decreasing in t and by (A.2) it is decreasing at rate $-\lambda\left(2 + e^{-\ell_t(\ell)} + e^{\ell_t(\ell')}\right)$ whenever $\ell_t(\ell, Z^t, \widetilde{\eta}) < \ell^* < \ell_t(\ell', Z^t, \widetilde{\eta})$. This proves (b).

For (c) the integrand is positive when $\lambda\widehat{\Delta}\left(\ell^*\right) \leq c$: Reducing investment on the margin is profitable if cost exceeds benefit.

For (d), we have $\lambda\widehat{\Delta}\left(\ell^*\right) = c$ and $\widehat{\Delta}$ is continuous by (b). Then the upper bound of (A.5) is of order $o\left(\ell' - \ell\right)$ because now $\max_{\ell\in[\ell^*, \ell^*+(\ell'-\ell)]}\left\{\left|c - \lambda\widehat{\Delta}\left(\ell\right)\right|\right\}$ is of order $o\left(1\right)$.

To prove (e) assume to the contrary that the cutoff is non-positive $\widehat{\Delta}\left(\ell^*\right) \leq 0$. A fortiori

$\widehat{\Delta}(\ell^*) \leq c$ and by part (c) the marginal value of reputation is positive $\widehat{V}'_H(\ell) > 0$, as is the reputational dividend $\widehat{D}_H(\ell)$. But then, also the value of quality would be positive $\widehat{\Delta}(\ell^*) > 0$. $\square$

Lemma 4(d) has the flavor of the envelope theorem: when the firm's first-order condition holds at the cutoff, then a change in initial reputation only affects its payoff through the reputational evolution. Intuitively, a firm with a lower initial reputation works more, leading to a gain of $\widehat{\Delta}(\ell)$ when a technology shock hits. This gain is exactly offset by the extra cost born by the firm. The marginal value of reputation $\widehat{V}'_\theta(\ell)$ is thus determined solely by the "durability" of the reputational increment $\ell'_t - \ell_t$.

In important cases we can truncate the integral (A.5) at time T when the reputational evolution hits ℓ^*. When $\ell_t < \ell^* < \ell'_t$ the increment $\ell'_t - \ell_t$ decreases at rate $-\lambda\left(2 + e^{-\ell_t} + e^{\ell'_t}\right)$. If the cutoff ℓ^* is "reflecting" because either $\ell^* \approx \pm\infty$ or $\lambda \gg 0$, the reputational increment $\ell'_t - \ell_t$ approximately disappears at T and we can restrict the integral in (A.5) to $t \leq T$.

We now come to the second extension of Lemma 2, which we will use to prove essential uniqueness of equilibrium (Lemma 10).

Lemma 5 *There exists a continuous function $\alpha(\ell, \ell') > 0$, that is a lower bound of the incremental value of reputation $\alpha(\ell, \ell') < \widehat{V}_\theta(\ell') - \widehat{V}_\theta(\ell)$ for $\ell < \ell'$, uniformly across cost c and equilibrium investment $\widehat{\eta}$.*

Proof. Let $\ell' > \ell$. In equilibrium the firm with the higher reputation ℓ' cannot benefit by imitating the investment of the firm with the lower reputation via $\overline{\eta}\left(Z^t\right) = \widehat{\eta}(\ell_t(\ell, Z^t, \widetilde{\eta}))$. Therefore,

$$\begin{aligned}
\widehat{V}_\theta(\ell') - \widehat{V}_\theta(\ell) &\geq \int e^{-rt}\mathbb{E}_{\overline{\eta}}\left[x\left(\ell_t\left(\ell', Z^t, \widetilde{\eta}\right)\right) - x\left(\ell_t\left(\ell, Z^t, \widetilde{\eta}\right)\right)\right] dt \\
&\geq k_1 \frac{e^\ell}{\left(1 + e^\ell\right)^2} \int e^{-rt}\mathbb{E}_{\overline{\eta}}\left[\ell_t\left(\ell', Z^t, \widetilde{\eta}\right) - \ell_t\left(\ell, Z^t, \widetilde{\eta}\right)\right] dt \\
&\geq k_1 \frac{e^\ell}{\left(1 + e^\ell\right)^2} \int e^{-rt} \max\left\{0; \ell' - \ell - \lambda\left(2 + e^{\ell'_t} + e^{-\ell_t}\right)t\right\} dt \\
&\geq k_2 \frac{e^\ell}{\left(1 + e^\ell\right)^2} \frac{\left(\ell' - \ell\right)^2}{2\lambda\left(2 + e^\ell + e^{-\ell}\right)}
\end{aligned}$$

The first line bounds $\widehat{V}_\theta(\ell')$ below by the value when investing according to $\overline{\eta}$. The second line essentially applies the chain rule to factor out $dx/d\ell = e^\ell / \left(1 + e^\ell\right)^2$, and the constant k_1 accounts for the possibility that $dx/d\ell$ needs to be evaluated at ℓ_t instead of ℓ, and possibly $dx(\ell_t)/d\ell < dx(\ell)/d\ell$. The second line takes the most pessimistic stance on the evolution of $\ell'_t - \ell_t$ by assuming that $\eta(\ell_t) = 1$ and $\eta(\ell'_t) = 0$ in which case $d(\ell'_t - \ell_t)/dt = -\lambda(2 + e^{\ell'_t} + e^{-\ell_t})t$. The third line evaluates the integral over an essentially linear function, and k_2 accounts for $e^{-rt} < 1$ and the possibility that $\lambda(2 + e^{\ell'_t} + e^{-\ell_t}) < \lambda(2 + e^\ell + e^{-\ell})$. $\square$

B Proof of Theorem 2

We prove Theorem 2 in ℓ-space, introduced in Appendix A.1. We show that for sufficiently small c there exists a cutoff ℓ^* such that:

(a) The cutoff is indifferent: $\lambda\widehat{\Delta}_{\ell^*}(\ell^*) = c$ (Section B.1, Lemma 6),

(b) Low reputations work: $\lambda\widehat{\Delta}_{\ell^*}(\ell) > c$ for $\ell < \ell^*$ (Section B.2, Lemma 8),

(c) High reputations shirk: $\lambda\widehat{\Delta}_{\ell^*}(\ell) < c$ for $\ell > \ell^*$ (Section B.3, Lemma 9).

Section B.4 shows the essential uniqueness of the work-shirk equilibrium.

B.1 Indifference of Cutoff

We now show that for small costs there exists a high cutoff ℓ^* that satisfies the indifference condition. Since $\widehat{\Delta}$ and $\widehat{V}$ depend on c, we subscript them with c where useful.

Lemma 6 *For every $\ell \in \mathbb{R}$ there exists $c(\ell) > 0$ such that for all $c^* < c(\ell)$ there exists $\ell^* > \ell$ such that $c^* = \lambda\widehat{\Delta}_{\ell^*,c^*}(\ell^*)$.*

Proof. Fix $\ell \in \mathbb{R}$ and consider $\widehat{\Delta}_{\ell,c}(\ell)$ as a function of $c \in [0, \lambda/(r+\lambda)]$. By Lemma 4(e) we have $\widehat{\Delta}_{\ell,c}(\ell) > 0$ for all c. Since $\widehat{\Delta}_{\ell,c}(\ell)$ is continuous in c, it takes on its minimum $\widehat{\Delta}_{\ell,c'}(\ell) > 0$ at some c'.

Let $c(\ell) = \lambda\widehat{\Delta}_{\ell,c'}(\ell)$ and fix $c^* \in (0, c(\ell))$. Using the definitions of c' and c^*,

$$\lambda\widehat{\Delta}_{\ell,c^*}(\ell) \geq \lambda\widehat{\Delta}_{\ell,c'}(\ell) > c^*,$$

so the firm prefers to work. On the other hand

$$c^* > \lim_{\ell' \to \infty} \lambda\widehat{\Delta}_{\ell',c^*}(\ell') = 0,$$

so by continuity of $\widehat{\Delta}_{\ell',c^*}(\ell')$ as a function of $\ell' \in [\ell, \infty]$, there exists $\ell^* \in (\ell, \infty)$ with $c^* = \lambda\widehat{\Delta}_{\ell^*,c^*}(\ell^*)$. □

The daunting array of quantifiers in the statement of this lemma guarantees that we can assume ℓ^* with $c^* = \lambda\widehat{\Delta}_{\ell^*,c^*}(\ell^*)$ as large as necessary in the upcoming arguments.

B.2 Low Reputations Work

We first use equation (A.5) to prove an auxiliary result about the marginal value of reputation:

Lemma 7 *Fix any $\alpha > 0$, $M > 0$ and ℓ_V sufficiently large. Suppose $\ell^* > \ell_V$ is sufficiently high and $c = \lambda\widehat{\Delta}_{\ell^*,c}(\ell^*)$.*

(a) $\widehat{V}'_\theta(\ell)$ is decreasing on $[\ell_V, \ell^*]$.

(b) $\widehat{V}'_\theta(\ell)$ "diminishes" to the right of ℓ^* :

$$\widehat{V}'_\theta(\ell^* - \gamma) > M\widehat{V}'_\theta(\ell^* + \beta) \text{ for all } \gamma \in [\alpha, \ell^* - \ell_V] \text{ and all } \beta > 0.$$

Intuitively, incremental reputation above ℓ^* is less "durable" because it disappears when ℓ_t hits ℓ^* and reputational updating $\frac{d\ell}{dt}$ decelerates from $-\lambda\left(1 + e^{\ell}\right) \approx -\infty$ to $\lambda\left(1 + e^{-\ell}\right) \approx \lambda$. Formally, for $\ell < \ell^*$ let $T(\ell) = \min\{t|\ell_t \geq \ell^*, \ell_0 = \ell\}$ be the *cutoff time*: This is the first time that the reputational dynamics starting at ℓ reach, or exceed, the cutoff ℓ^*. For $\ell > \ell^*$ let $T(\ell) = \min\{t|\ell_t \leq \ell^*, \ell_0 = \ell\}$.

Proof. (a): For high enough ℓ^* and $\ell_0 \in (\ell_V, \ell^*)$, the reputational dynamics ℓ_t are approximately governed by a Brownian motion with constant drift and jumps (A.3). As long as $\ell_t < \ell^*$ we have $\frac{\partial \ell_t}{\partial \ell_0} \approx 1$. When the trajectory hits ℓ^* at time $T(\ell_0)$, then $\frac{\partial \ell_t}{\partial \ell_0} \approx 0$ for $t > T(\ell_0)$ since ℓ^* is reflecting. Using equation (A.5)

$$\widehat{V}'_\theta(\ell_0) = \int_0^\infty e^{-rt}\mathbb{E}_{\ell_0}\left[\frac{e^{\ell_t}}{(1+e^{\ell_t})^2}\frac{\partial \ell_t}{\partial \ell_0}\right]dt \approx \mathbb{E}_{\ell_0}\left[\int_0^{T(\ell_0)} e^{-rt}\frac{e^{\ell_t}}{(1+e^{\ell_t})^2}\frac{\partial \ell_t}{\partial \ell_0}dt\right]. \tag{B.1}$$

Since $e^{\ell_t}/(1+e^{\ell_t})^2$ is strictly decreasing for $\ell_t > 0$, and $T(\ell_0)$ is decreasing in ℓ_0, equation (B.1) is strictly decreasing in ℓ on $[\ell_V, \ell^*]$.

(b): Let $\gamma \in [\alpha, \ell^* - \ell_V]$. When ℓ^* is sufficiently high, the reputational dynamics are given by (A.3). The drift is finite, and the expected cutoff time $\mathbb{E}\left[T(\ell^* - \gamma)\right]$ is bounded below independently of ℓ^*. Thus, $\mathbb{E}\left[\frac{d\ell_t}{d\ell_0}(\ell')\right]$ for $\ell' \in [\ell_V, \ell^* - \alpha]$ is bounded away from 0 as $\ell^* \to \infty$.

Next, consider $\ell^* + \beta$. The expected time until cutoff $\mathbb{E}\left[T(\ell^* + \beta)\right]$ converges to 0 as $\ell^* \to \infty$, uniformly in δ. This is easier to see for the posterior x_t than for the log-likelihood-ratio ℓ, as $\mathbb{E}\left[\frac{d_\theta x}{dt}\right] = -\lambda x$ is bounded away from 0 while $1 - x^*$ converges to 0. Thus, $\mathbb{E}\left[\frac{d\ell_t}{d\ell_0}(\ell^* + \beta)\right]$ for $\beta > 0$ converges to 0 as $\ell^* \to \infty$.

For large values of ℓ^* we can ignore in (A.5) all terms with $t > T(\ell_0)$. Since $e^{\ell}/(1+e^{\ell})^2$ is decreasing in $\ell > 0$ we get bounds $e^{\ell_t(\ell^*-\gamma)}/(1+e^{\ell_t(\ell^*-\gamma)})^2 \geq e^{\ell^*}/(1+e^{\ell^*})^2 \geq e^{\ell_t(\ell^*+\beta)}/\left(1+e^{\ell_t(\ell^*+\beta)}\right)^2$ and equation (A.5) implies

$$\frac{\widehat{V}'_\theta(\ell^* - \gamma)}{\widehat{V}'_\theta(\ell^* + \beta)} \geq \frac{\frac{e^{\ell^*}}{(1+e^{\ell^*})^2}\mathbb{E}\left[\int_0^{T(\ell^*-\gamma)} e^{-rt}\frac{d\ell_t}{d\ell_0}(\ell^* - \gamma)dt\right]}{\frac{e^{\ell^*}}{(1+e^{\ell^*})^2}\mathbb{E}\left[\int_0^{T(\ell^*+\beta)} e^{-rt}\frac{d\ell_t}{d\ell_0}(\ell^* + \beta)dt\right]} \geq const.\frac{\mathbb{E}\left[T(\ell^* - \gamma)\right]}{\mathbb{E}\left[T(\ell^* + \beta)\right]}$$

Therefore, $\widehat{V}_\theta'(\ell^*-\gamma)/\widehat{V}_\theta'(\ell^*+\beta)$ diverges as $\ell^*\to\infty$, uniformly over all $\gamma\in[\alpha,\ell^*-\ell_V]$ and $\beta>0$. $\square$

Lemma 8 shows that firms with low reputations work. For reputations $\ell\in[\ell_\Delta,\ell^*]$ for some ℓ_Δ defined below, the optimality of working follows directly by showing that $\widehat{\Delta}(\ell)$ is decreasing on $[\ell_\Delta,\ell^*]$. For reputations $\ell<\ell_\Delta$ the result follows from the closeness of $\widehat{\Delta}_{\ell^*}(\cdot)$ and $\widehat{\Delta}_\infty(\cdot)$.

Lemma 8 *Assume ℓ^* is large, costs c are small and $\lambda\widehat{\Delta}_{\ell^*}(\ell^*)=c$. Then $\lambda\widehat{\Delta}_{\ell^*}(\ell)>c$ for all $\ell<\ell^*$.*

Proof. *Claim 1.* For any $\alpha>0$, there exists ℓ_D sufficiently large such that $\widehat{D}_H(\cdot)$ is strictly decreasing on $[\ell_D,\ell^*-\alpha]$ for any $\ell^*>\ell_D$.

Proof. By (A.4) $\widehat{D}_H(\cdot)$ is composed of $\widehat{V}_H'$-terms and jump terms. The former are taken care off by Lemma 7(a). For the jump terms, pick α,ℓ_V and $M=1$ as in Lemma 7, and choose $\ell_D=\ell_V+\max_y\{-\delta_y\}$. We need to show $\widehat{D}_H(\ell)>\widehat{D}_H(\ell')$ for all $\ell<\ell'$ in $[\ell_D,\ell^*-\alpha]$ and all $y\in Y$.

First consider good news events $y\in Y^+$ with $\delta_y,\mu_y>0$ and assume wlog that $\ell+\delta_y>\ell'$. Then

$$\begin{aligned}\widehat{D}_H(\ell)-\widehat{D}_H(\ell') &= \mu_y\left(\int_\ell^{\ell+\delta_y}\widehat{V}_H'\left(\widetilde{\ell}\right)d\widetilde{\ell}-\int_{\ell'}^{\ell'+\delta_y}\widehat{V}_H'\left(\widetilde{\ell}\right)d\widetilde{\ell}\right)\\ &= \mu_y\int_\ell^{\ell'}\left(\widehat{V}_H'\left(\widetilde{\ell}\right)-\widehat{V}_H'\left(\widetilde{\ell}+\delta_y\right)\right)d\widetilde{\ell}\\ &> 0\end{aligned}$$

where the inequality follows pointwise for every $\widetilde{\ell}$ from Lemma 7(a) if $\widetilde{\ell}+\delta_y<\ell^*$ and from Lemma 7(b) if $\widetilde{\ell}+\delta_y>\ell^*$.

Now consider bad news events $y\in Y^-$ with $\delta_y,\mu_y<0$. Then

$$\begin{aligned}\widehat{D}_H(\ell)-\widehat{D}_H(\ell') &= -\mu_y\left(\int_{\ell+\delta_y}^{\ell}\widehat{V}_H'\left(\widetilde{\ell}\right)d\widetilde{\ell}-\int_{\ell'+\delta_y}^{\ell'}\widehat{V}_H'\left(\widetilde{\ell}\right)d\widetilde{\ell}\right)\\ &= -\mu_y\int_{\ell+\delta_y}^{\ell}\left(\widehat{V}_H'\left(\widetilde{\ell}\right)-\widehat{V}_H'\left(\widetilde{\ell}+(\ell'-\ell)\right)\right)d\widetilde{\ell}\\ &> 0\end{aligned}$$

where the inequality follows pointwise for every $\widetilde{\ell}$ by Lemma 7(a) because $\ell+\delta_y\geq\ell_V$.

Claim 2. For any $\varepsilon > 0$, there exists $\alpha > 0$, ℓ_D arbitrarily high, and $\ell_\Delta > \ell_D$ sufficiently high such that for any $\ell^* > \ell_\Delta$ and $\ell' \in (\ell_\Delta, \ell^*)$, we have

$$(r+\lambda)\int e^{-(r+\lambda)t}\Pr(\ell'_t \in [\ell_D, \ell^* - \alpha])dt \geq 1 - \varepsilon.$$

Proof. This is because the reputational dynamics $d\ell_t$ in $[\underline{\ell}, \ell^*]$ are approximately governed by a Brownian motion with constant drift and jumps (A.3) reflected at ℓ^*.

Claim 3. There exists ℓ_Δ sufficiently large such that $\widehat{\Delta}(\cdot)$ is strictly decreasing on $[\ell_\Delta, \ell^*]$ for any $\ell^* > \ell_\Delta$.

Proof. Pick ℓ_D as in Claim 1 and $\ell_\Delta > \ell_D$ as in Claim 2. Claim 1 states that $\widehat{D}_H(\cdot)$ is strictly decreasing on $[\ell_D, \ell^* - \alpha]$ and claim 2 states that $\ell_t \in [\ell_D, \ell^* - \alpha]$ with probability close to 1. The claim follows because the reputational evolution ℓ_t is monotone, and the value of quality $\widehat{\Delta}(\ell)$ is the integral over the dividends $\widehat{D}_H(\ell_t)$.

Claim 4. Assume that $\lambda\widehat{\Delta}_{\ell^*}(\ell^*) = c$ and fix any $\bar{\ell}$. Then $\widehat{\Delta}_{\ell^*}(\cdot)$ converges to $\widehat{\Delta}_\infty(\cdot)$ uniformly on $[-\infty, \bar{\ell}]$ as $\ell^* \to \infty$.

Proof. As $\ell^* \to \infty$, $\widehat{\Delta}_{\ell^*}(\ell)$ converges pointwise to $\widehat{\Delta}_\infty(\ell)$ for all ℓ. Let $\ell^* \gg \bar{\ell}$. For any $\ell < \bar{\ell}$, equations (A.5) and (A.4) imply that, $\widehat{V}'_{\theta,\ell^*}(\ell)$ and $\widehat{D}_{\theta,\ell^*}(\ell)$, and thus $\widehat{\Delta}_{\ell^*}(\ell)$, depend on ℓ^* only on trajectories ℓ_t that reach ℓ^*. The future discounted probability of these trajectories converges to 0 as $\ell^* \to \infty$, so the convergence is uniform for $\ell < \bar{\ell}$.

Proof of Lemma. Choose $0 \ll \bar{\ell} \ll \ell^*$. Claim 3 implies that $\widehat{\Delta}_{\ell^*}(\ell)$ is strictly decreasing in ℓ for $\ell \in [\bar{\ell}, \ell^*)$. Since $\lambda\widehat{\Delta}_{\ell^*}(\ell^*) = c$, we have

$$\lambda\widehat{\Delta}_{\ell^*}(\ell) > c \quad \text{ for } \ell \in [\bar{\ell}, \ell^*).$$

The function $\widehat{\Delta}_\infty(\cdot)$ is bounded away from 0 on $[-\infty, \bar{\ell}]$. Hence Claim 4 implies that $\widehat{\Delta}_{\ell^*}(\ell)$ is bounded away from zero. For small costs c, we get

$$\lambda\widehat{\Delta}_{\ell^*}(\ell) > c \quad \text{ for } \ell \in [-\infty, \bar{\ell}],$$

as required. □

B.3 High Reputations Shirk

Lemma 9 *Suppose ℓ^* is large and $\lambda\widehat{\Delta}(\ell^*) = c$. Then $\lambda\widehat{\Delta}(\ell') < c$ for all $\ell' > \ell^*$.*

Proof. The idea of the proof is to develop $\widehat{\Delta}(\ell')$ into dividends $\widehat{D}(\ell'_t)$ and a continuation value $e^{-(r+\lambda)T}\widehat{\Delta}(\ell'_T)$ for a "small" stopping time T, and show that the dividends $\widehat{D}(\ell'_t)$ are small compared to dividends of $\widehat{\Delta}(\ell^*)$, while the continuation value $e^{-(r+\lambda)T}\widehat{\Delta}(\ell'_T)$ is no bigger than the respective term of $\widehat{\Delta}(\ell^*)$. There are two possible comparisons:

(a) Develop $\widehat{\Delta}(\ell^*)$ until T and compare dividends and appreciation of continuation values separately, i.e. show $\widehat{D}(\ell'_t) < \widehat{D}(\ell^*_t)$ and $\widehat{\Delta}(\ell'_T) - \widehat{\Delta}(\ell') \leq \widehat{\Delta}(\ell^*_T) - \widehat{\Delta}(\ell^*)$.

(b) Let $T = T(\ell')$ be the time when ℓ'_t first hits ℓ^* (so that $\widehat{\Delta}(\ell'_T) = \widehat{\Delta}(\ell^*)$) and compare the dividend $\widehat{D}(\ell'_t)$ to the annuity value of $\widehat{\Delta}(\ell^*)$.

Below we show $\widehat{D}(\ell'_t) < \widehat{D}(\ell^*_t)$ for $t \approx 0$ and use comparison (a), if there is some Poisson event $y \in Y^-$ signalling bad news, i.e. $\mu_y = \mu_{H,y} - \mu_{L,y} < 0$. On the other hand we can show $\widehat{D}(\ell'_t) < (r+\lambda)\widehat{\Delta}(\ell^*)$ and use comparison (b), if there is no such bad news Poisson event y.

Case (a): Assume that there is a bad news Poisson event $y \in Y$, with $\mu_y < 0$. Assume by contradiction, that $\ell' > \ell^*$ maximizes $\widehat{\Delta}(\ell)$ on $[\ell^*, \infty]$. We develop $\widehat{\Delta}(\ell')$ and $\widehat{\Delta}(\ell^*)$ until T and subtract $e^{-(r+\lambda)T}\widehat{\Delta}(\ell')$ to express the rental value of $\widehat{\Delta}(\ell')$ as sum of dividends and appreciation:

$$
\begin{aligned}
(r+\lambda)\, T\widehat{\Delta}(\ell') &= \int_0^T e^{-(r+\lambda)t}\mathbb{E}_{\theta^t=L}\left[\widehat{D}_H(\ell'_t)\right] dt + \mathbb{E}\left[\widehat{\Delta}(\ell'_T) - \widehat{\Delta}(\ell')\right], \\
(r+\lambda)\, T\widehat{\Delta}(\ell^*) &= \int_0^T e^{-(r+\lambda)t}\mathbb{E}_{\theta^t=L}\left[\widehat{D}_H(\ell^*_t)\right] dt + \mathbb{E}\left[\widehat{\Delta}(\ell^*_T) - \widehat{\Delta}(\ell^*)\right],
\end{aligned}
$$

Let $T > 0$ be sufficiently small[14] with $\ell'_t > \ell^*$ for $t \leq T$ and $\ell^*_T < \ell^*$. Then the appreciation $\mathbb{E}\left[\widehat{\Delta}(\ell^*_T) - \widehat{\Delta}(\ell^*)\right]$ is positive by Claims 2 and 3 in the proof of Lemma 8, while $\mathbb{E}\left[\widehat{\Delta}(\ell'_T) - \widehat{\Delta}(\ell')\right]$ is negative by choice of ℓ'.

We now argue that $\widehat{D}_H(\ell^*_t) > \widehat{D}_H(\ell'_t)$. We first compare the jump terms for bad news events $y \in Y^-$ with $\mu_y, \delta_y < 0$. Because of the reflecting dynamics we can assume that $\ell^*_t < \ell^*$. If $\ell^*_t < \ell'_t + \delta_y$ we can conclude by Lemma 7(b); otherwise we have $\ell^*_t + \delta_y < \ell'_t + \delta_y < \ell^*_t < \ell^* < \ell'_t$, and just like in the proof of claim 1 in Lemma 8 we obtain

$$
\mu_y\left(\widehat{V}_H(\ell^*_t + \delta_y) - \widehat{V}_H(\ell^*_t)\right) - \mu_y\left(\widehat{V}_H(\ell'_t + \delta_y) - \widehat{V}_H(\ell'_t)\right) = -\mu_y \int_{\ell^*_t}^{\ell'_t} \left(\widehat{V}'_H(\ell + \delta_y) - \widehat{V}'_H(\ell)\right) d\ell. \tag{B.2}
$$

The RHS is strictly positive because $\widehat{V}'_H(\ell + \delta_y)$ exceeds $\widehat{V}'_H(\ell)$ by Lemma 7(a) when $\ell < \ell^*$ and by Lemma 7(b) when $\ell > \ell^*$.

The (B.2)-term dominates the good news and Brownian component of $\widehat{D}_H(\ell'_t)$, i.e. $\mu_y \int_{\ell'_t}^{\ell'_t+\delta_y} \widehat{V}'_H\left(\widetilde{\ell}\right) d\widetilde{\ell}$ for $y \in Y^+$ and $\widehat{V}'_H(\ell'_t)\mu_B^2/2$. This is because $\widehat{V}'_H(\ell^*_t + \delta_y)/\widehat{V}'_H(\ell'_t)$ diverges as $\ell^* \to \infty$ by Lemma 7(b).

[14] Technically, T is a stopping time, but for notational simplicity we treat it as a real number here.

Case (b): Assume that there is no bad news Poisson event, i.e. $\mu_y \geq 0$ for all $y \in Y$. We develop $\widehat{\Delta}(\ell')$ until cutoff time $T(\ell')$. As there are no downward jumps in ℓ'_t, we have $\ell'_t > \ell^*$ for $t < T$.

$$\begin{aligned}\widehat{\Delta}(\ell') - \widehat{\Delta}(\ell^*) &= \mathbb{E}\left[\int_0^{T(\ell')} e^{-(r+\lambda)t}\widehat{D}_H\left(\ell'_t\right)dt + e^{-(r+\lambda)T(\ell')}\widehat{\Delta}\left(\ell^*\right)\right] - \widehat{\Delta}(\ell^*) \\ &= \mathbb{E}\left[\int_0^{T(\ell')}\left(e^{-(r+\lambda)t}\widehat{D}_H\left(\ell'_t\right) - (r+\lambda)\widehat{\Delta}\left(\ell^*\right)\right)dt\right] \end{aligned} \tag{B.3}$$

To show that the integrand is negative, we develop $(r+\lambda)\widehat{\Delta}(\ell^*)$ into future reputational dividends $\widehat{D}_H(\ell^*_t)$, that will on average exceed $\widehat{D}_H(\ell'_t)$:

$$\begin{aligned}(r+\lambda)\widehat{\Delta}\left(\ell^*\right) &= (r+\lambda)\int_0^\infty e^{-(r+\lambda)t}\mathbb{E}\left[\widehat{D}_H\left(\ell^*_t\right)\right]dt \\ &\geq (r+\lambda)\int_0^\infty e^{-(r+\lambda)t}\Pr\left(\ell^*_t \in [\ell_D, \ell^* - \alpha]\right)\inf_{\ell\in[\ell_D,\ell^*-\alpha]}\left\{\widehat{D}_H(\ell)\right\}dt \\ &\geq (r+\lambda)\int_0^\infty e^{-(r+\lambda)t}(1-\varepsilon)M\sup_{\ell>\ell^*}\left\{\widehat{D}_H(\ell)\right\}dt \\ &\geq \sup_{t\leq T(\ell')}\left\{\widehat{D}_H\left(\ell'_t\right)\right\}\end{aligned}$$

The third line uses that for $\alpha > 0$ sufficiently small, and ℓ^* sufficiently large we get $\Pr\left(\ell^*_t \in [\ell_D, \ell^* - \alpha]\right) = 1-\varepsilon$ by Claim 2 in the proof of Lemma 8, while by choosing ℓ^* large enough, we get $\inf_{\ell\in[\ell_D,\ell^*-\alpha]}\left\{\widehat{D}_H(\ell)\right\} > M\sup_{\ell>\ell^*}\left\{\widehat{D}_H(\ell)\right\}$ for any M by Lemma 7(b). $\square$

B.4 Essential Uniqueness

We prove Theorem 2(c) in two steps. Lemma 10 shows that for low costs, intermediate reputations prefer to work in any tentative equilibrium investment profile. Lemma 11 shows that if market learning ensures (HOPE) and intermediate reputations work, so do low reputations. Intuitively, a firm with reputation just below a tentative shirk-work cutoff hopes to achieve an intermediate reputation in the future and thus prefers to work.

Lemma 10 *On any bounded interval $\left[-\overline{\ell}, \overline{\ell}\right]$ the value of quality $\widehat{\Delta}(\ell)$ is bounded away from zero, uniformly across costs c and equilibrium investment profiles $\widehat{\eta}$.*

Proof. To show this, we prove existence of a continuous, strictly positive, lower bound of $\widehat{\Delta}(\ell)$.

If the Poisson component of market learning is non-trivial, i.e. there is $y \in Y$ with $\delta_y \neq 0$, then the reputational dividend is uniformly bounded below by Lemma 5, and so is $\widehat{\Delta}(\ell)$.

If market learning is pure Brownian and the value of quality equals $\widehat{\Delta}(\ell) = \int e^{-(r+\lambda)t}\mathbb{E}\left[\mu_B^2 \widehat{V}_\theta'(\ell_t)/2\right] dt$, this argument still holds. While Lemma 5 does not prove $\widehat{V}_\theta'(\ell_t) > 0$ for every ℓ_t, the volatility of the Brownian motion distributes ℓ_t over some interval $[\ell, \ell']$ and the expectation $\mathbb{E}\left[\frac{\mu_B^2}{2}\widehat{V}_\theta'(\ell_t)\right]$ is then bounded below by *const.* $\left(\widehat{V}_\theta(\ell') - \widehat{V}_\theta(\ell)\right)$. □

Lemma 11 *Fix $\overline{\ell} > 0$. If market learning ensures (HOPE) and costs are sufficiently low, then a firm with reputation below $\overline{\ell}$ works in equilibrium.*

Proof. From Lemma 10 we know that the firm prefers to work for $\ell \in \left[-\overline{\ell}, \overline{\ell}\right]$. By contradiction, assume that there is equilibrium shirking in $\left[-\infty, -\overline{\ell}\right]$ and let the highest shirk-work cutoff be $\ell_* < -\overline{\ell}$. We develop $\widehat{\Delta}(\ell_* - \varepsilon)$ until the first time T when $\ell_t = -\overline{\ell}$:

$$\widehat{\Delta}(\ell_* - \varepsilon) \geq \mathbb{E}[e^{-(r+\lambda)T}]\widehat{\Delta}(\ell_T). \tag{B.4}$$

As $\widehat{\Delta}(\ell_T) = \widehat{\Delta}(-\overline{\ell})$ is bounded below by $\beta\left(-\overline{\ell}\right)$ we just need to show that $\mathbb{E}[e^{-(r+\lambda)T}] > 0$, independently of investment $\widehat{\eta}$ and costs c: By the assumption that market learning ensures (HOPE), and by choosing $-\overline{\ell}$, and thus ℓ_*, low enough, the firm's initial reputation $\ell_* - \varepsilon$ will rise above ℓ_* with positive probability. Once $\ell_t > \ell_*$, equilibrium beliefs $\widetilde{\eta} = 1$ will push reputation $\ell_{t'}$ to $-\overline{\ell}$ in finite time with positive probability.

Thus the right-hand-side of (B.4) is strictly positive, and by choosing c small enough we get $\lambda\widehat{\Delta}(\ell_* - \varepsilon) > c$, and the desired contradiction. □

C Quality Choice

Proof of Lemma 3. We will show that the dividend $\widehat{V}'_\theta(\ell)$ approaches 0 uniformly over all reputations and all work-shirk investment profiles. That is, $\lim_{\lambda\to\infty}\sup_{\ell^*\in[-\infty,\infty],\ell\in\mathbb{R}}\widehat{V}'_{\theta,\ell^*}(\ell)=0$, where $\ell^*=\pm\infty$ captures the full shirk (resp. work) profile.

To do so, fix ϵ and let $\ell^{**}>0$ solve $e^{\ell^{**}}/\left(1+e^{\ell^{**}}\right)^2=\epsilon$. First, consider a cutoff ℓ^* in the tail, i.e. $|\ell^*|>\ell^{**}$.

$$\begin{aligned}
\lim_{\lambda\to\infty}\sup_{|\ell^*|>\ell^{**},\ell_0\in\mathbb{R}}\widehat{V}'_{\theta,\ell^*}(\ell_0) &= \lim_{\lambda\to\infty}\sup_{|\ell^*|>\ell^{**},\ell_0}\mathbb{E}\left[\int e^{-rt}\frac{e^{\ell_t}}{\left(1+e^{\ell_t}\right)^2}\frac{\partial\ell_t}{\partial\ell_0}dt\right]\\
&\leq \lim_{\lambda\to\infty}\sup_{|\ell^*|>\ell^{**},\ell_0}\mathbb{E}\left[\int e^{-rt}\left(\frac{\Pr(|\ell_t|<\ell^{**})}{4}+\frac{\Pr(|\ell_t|\geq\ell^{**})e^{\ell^{**}}}{\left(1+e^{\ell^{**}}\right)^2}\right)dt\right]\\
&\leq \varepsilon.
\end{aligned}$$

The first line applies Lemma 4 (d). The second line uses $e^{\ell}/(1+e^{\ell})^2\leq 1/4$ and Lemma 4 (a). The first term under the integral vanishes because $\lim_{\lambda\to\infty}\sup_{\ell_0}\Pr\left(|\ell_t|<\ell^{**}\right)=0$ for all $t>0$. The second term is bounded by ε.

Next, suppose that $|\ell^*|\leq\ell^{**}$. Suppose the process ℓ_t hits ℓ^* at time T. Using Lemma 4(d) we obtain:

$$\begin{aligned}
\lim_{\lambda\to\infty}\sup_{|\ell^*|<\ell^{**},\ell_0\in\mathbb{R}}\widehat{V}'_{\theta}(\ell_0) &= \lim_{\lambda\to\infty}\sup_{|\ell^*|<\ell^{**},\ell_0}\mathbb{E}\left[\int e^{-rt}\frac{e^{\ell_t}}{\left(1+e^{\ell_t}\right)^2}\frac{\partial\ell_t}{\partial\ell_0}dt\right]\\
&\leq \lim_{\lambda\to\infty}\sup_{|\ell^*|<\ell^{**},\ell_0}\frac{1}{4}\mathbb{E}\left[\int_{t=0}^{T}e^{-rt}dt+\int_{t=T}^{\infty}e^{-rt}\frac{\partial\ell_t}{\partial\ell_0}dt\right]\\
&= 0.
\end{aligned}$$

The first two lines are as above. The first integral vanishes because $\sup_{\ell_0,|\ell^*|<\ell^{**}}\mathbb{E}\left[T\right]\to 0$ as $\lambda\to\infty$. The second integral vanishes because reputational increments disappear at an absorbing boundary, i.e. $E[\partial\ell_t/\partial\ell_0(\ell^*)|t\geq T]\to 0$ as $\lambda\to\infty$.

Thus the marginal value of reputation $\widehat{V}'(\ell)$ uniformly approaches 0 as $\lambda\to\infty$. □

D Perfect Poisson Learning

In this appendix we solve the perfect learning specifications of Section 5 explicitly by calculating value functions in closed form. This approach highlights the analytic tractability of these learning specificications and delivers a more explicit understanding of value functions and the value of quality. Some of the derived expressions are also used in the proofs of Section 5.

We assume throughout that $\lambda \geq \mu$, so that the drift of the firm's reputation is determined by market beliefs.

D.1 Perfect Good News

Shirk-region, above the cutoff $x \geq x^*$. Suppose $x_0 = 1$ and let x_t solve the dynamics (5.1), absent a breakthrough. The firm weakly prefers to shirk; we assume it always works, so x_t is strictly decreasing until it stops at $x_t = x^*$. With a low quality product, reputation is deterministic and firm value is given by:

$$V_L(x_s) = \int_{t=0}^{\infty} e^{-rt} x_{t+s} dt \tag{D.1}$$

With a high quality product dynamics are more complicated, because the reputation jumps to 1 at a μ-shock and quality disappears at a λ-shock:

$$\begin{aligned} V_H(x_s) &= \int_{t=0}^{\infty} e^{-(r+\lambda+\mu)t}[x_{t+s} + \lambda V_L(x_{t+s}) + \mu V_H(1)]dt \\ &= \int_{t=0}^{\infty} x_{t+s} e^{-rt} \left[\frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu} e^{-(\mu+\lambda)t} \right] dt + \frac{\mu}{r+\lambda+\mu} V_H(1), \end{aligned} \tag{D.2}$$

where we rewrote the $\lambda V_L(x_{t+s})$-term by changing the order of integration:

$$\int_{t=0}^{\infty} e^{-(r+\lambda+\mu)t} \lambda V_L(x_{t+s}) dt = \frac{\lambda}{\lambda+\mu} \int_{t=0}^{\infty} x_{t+s} e^{-rt} [1 - e^{-(\mu+\lambda)t}] dt.$$

We evaluate (D.2) at $x_s = 1$, and rearrange

$$V_H(1) = \frac{r+\lambda+\mu}{r+\lambda} \int_{t=0}^{\infty} x_t e^{-rt} \left[\frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu} e^{-(\mu+\lambda)t} \right] dt.$$

The value of quality is the difference between the value functions (D.2) and (D.1)[15]

$$\Delta(x_s) = \frac{\mu}{r+\lambda} \int_{t=0}^{\infty} x_t e^{-rt} \left[\frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu} e^{-(\mu+\lambda)t} \right] dt - \frac{\mu}{\lambda+\mu} \int_{t=0}^{\infty} x_{t+s} e^{-rt} \left[1 - e^{-(\mu+\lambda)t} \right] dt. \tag{D.3}$$

[15] Equivalently, one could compute the reputational dividend from the value functions and substitute it into the dividend formula for the value of quality (5.2).

When $x_s = x^*$, we get

$$\Delta(x^*) = \frac{\mu}{r+\lambda} \int_{t=0}^{\infty} (x_t - x^*) e^{-rt} \left[\frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu} e^{-(\mu+\lambda)t} \right] dt. \tag{D.4}$$

Quality at x^* is valuable because of the possibility that reputation jumps from x^* to $x = 1$. The terms in brackets capture the possibilities of λ and μ shocks while x_t descends from 1 to x^*.

Work-region, below the cutoff $x \leq x^*$. Next, suppose $\tilde{x}_0 = 0$ and let $\tilde{x}_t$ solve the dynamics (5.1), absent a breakthrough. The firm weakly prefers to work; we assume it always works, so $\tilde{x}_t$ is strictly increasing until it stops at $\tilde{x}_t = x^*$. With a high quality product, the firm's reputation drifts up until $\tilde{x}_t = x^*$, or a μ-shock hits:

$$V_H(\tilde{x}_s) = \int_{t=0}^{\infty} e^{-(r+\mu)t} [(\tilde{x}_{t+s} - c) + \mu V_H(1)] dt. \tag{D.5}$$

With a low quality product, the firm's reputation drifts up until $\tilde{x}_t = x^*$, or a λ-shock hits:

$$\begin{aligned} V_L(\tilde{x}_s) &= \int_{t=0}^{\infty} e^{-(r+\lambda)t} [(\tilde{x}_{t+s} - c) + \lambda V_H(\tilde{x}_{t+s})] dt \\ &= \int_{t=0}^{\infty} (\tilde{x}_{t+s} - c) \left[\frac{\lambda}{\lambda-\mu} e^{-(r+\mu)t} - \frac{\mu}{\lambda-\mu} e^{-(r+\lambda)t} \right] dt + \frac{\lambda}{r+\lambda} \frac{\mu}{r+\mu} V_H(1), \end{aligned} \tag{D.6}$$

where we rewrote the $\lambda V_H(\tilde{x}_{t+s})$-term by changing the order of integration:

$$\int_{t=0}^{\infty} e^{-(r+\lambda)t} \lambda V_H(\tilde{x}_{t+s}) dt = \frac{\lambda}{\lambda-\mu} \int_{t=0}^{\infty} (\tilde{x}_{t+s} - c) e^{-rt} (e^{-\mu t} - e^{-\lambda t}) dt + \frac{\lambda}{r+\lambda} \frac{\mu}{r+\mu} V_H(1).$$

The value of quality is the difference between the value functions (D.5) and (D.6):

$$\Delta(\tilde{x}_s) = \frac{r}{r+\lambda} \frac{\mu}{r+\mu} V_H(1) - \frac{\mu}{\lambda-\mu} \int_{t=0}^{\infty} (\tilde{x}_{t+s} - c) e^{-rt} (e^{-\mu t} - e^{-\lambda t}) dt$$

The first term captures the value of the high quality firm's breakthroughs while the second term captures the opportunity cost.

D.2 Perfect Bad News

Work-region, above the cutoff $x > x^*$. First assume that $x^* > 0$, so that $V_L(0) = 0$. Consider starting at x_0, just above x^*, and let x_t solve the dynamics (5.3), absent a breakdown. Then, x_t is strictly increasing and converges to 1. With a high quality product, reputation is deterministic and firm value equals

$$V_H(x_s) = \int_{t=0}^{\infty} e^{-rt} (x_{t+s} - c)] \, dt. \tag{D.7}$$

With a low quality product, dynamics are more complicated because the reputation jumps to 0 at a μ-shock and quality improves at a λ-shock,

$$\begin{aligned} V_L(x_s) &= \int_{t=0}^{\infty} e^{-(r+\lambda+\mu)t}[(x_{t+s}-c)+\lambda V_H(x_{t+s})+\mu V_L(0)]\,dt. \\ &= \int_{t=0}^{\infty} e^{-rt}(x_{t+s}-c)\left[\frac{\lambda}{\lambda+\mu}+\frac{\mu}{\lambda+\mu}e^{-(\lambda+\mu)t}\right]\,dt, \end{aligned} \tag{D.8}$$

where we rewrote the $\lambda V_H(x_{t+s})$-term by changing the order of integration:

$$\int_{t=0}^{\infty} e^{-(r+\lambda+\mu)t}\lambda V_H(x_{t+s})\,dt = \frac{\lambda}{\mu+\lambda}\int_{t=0}^{\infty} e^{-rt}(x_{t+s}-c)[1-e^{-(\lambda+\mu)t}]\,dt.$$

The value of quality is the difference between the value functions (D.7) and (D.8),

$$\Delta(x_s) = \frac{\mu}{\lambda+\mu}\int_{t=0}^{\infty} e^{-rt}(x_{t+s}-c)(1-e^{-(\lambda+\mu)t})\,dt. \tag{D.9}$$

The cost of low quality is the loss of reputation when the μ-shock hits before the λ-shock.

Shirk-region, below the cutoff $x < x^*$. Next, consider starting at $\tilde{x}_0$ just below x^*, and let $\tilde{x}_t$ solve the dynamics (5.3), absent a breakdown. Then, $\tilde{x}_t$ is strictly decreasing and converges to 0. With a low quality product, reputation is deterministic and firm value equals

$$V_L(\tilde{x}_s) = \int_{t=0}^{\infty} e^{-(r+\mu)t}\tilde{x}_{t+s}\,dt. \tag{D.10}$$

With a high quality product, quality disappears at a λ-shock and the firm's value function is

$$\begin{aligned} V_H(\tilde{x}_s) &= \int_{t=0}^{\infty} e^{-(r+\lambda)t}[\tilde{x}_{t+s}+\lambda V_L(\tilde{x}_{t+s})]\,dt. \\ &= \int_{t=0}^{\infty} \tilde{x}_{t+s}\left[\frac{\lambda}{\lambda-\mu}e^{-(r+\mu)t}-\frac{\mu}{\lambda-\mu}e^{-(r+\lambda)t}\right]\,dt, \end{aligned} \tag{D.11}$$

where we rewrote the $\lambda V_L(\tilde{x}_{t+s})$-term by changing the order of integration:

$$\int_{t=0}^{\infty} e^{-(r+\lambda)t}\lambda V_L(\tilde{x}_{t+s}) = \frac{\lambda}{\lambda-\mu}\int_{t=0}^{\infty} e^{-rt}\tilde{x}_{t+s}(e^{-\mu t}-e^{-\lambda t})dt.$$

The value of quality is the difference of the value functions (D.11) and (D.10):

$$\Delta(x_s) = \frac{\mu}{\lambda-\mu}\int_{t=0}^{\infty} e^{-rt}x_{t+s}(e^{-\mu t}-e^{-\lambda t})\,ds. \tag{D.12}$$

Again, the cost of low quality is the loss of reputation when the μ-shock hits before the λ-shock.

Full work. Suppose $x^* = 0$ so the firm always works. Consider $x_0 = 0$ and let x_t denote the dynamics (5.3), absent a breakdown. The value function for the high quality firm is given by (D.7). The value function of the low quality firm becomes

$$V_L(x_s) = \int_{t=0}^{\infty} e^{-rt}(x_{t+s} - c)\left[\frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu}e^{-(\lambda+\mu)t}\right] dt + \frac{\mu}{r+\lambda+\mu}V_L(0). \tag{D.13}$$

Setting $s = 0$, we obtain

$$V_L(0) = \frac{r+\lambda+\mu}{r+\lambda}\int_{t=0}^{\infty} e^{-rt}(x_{t+s} - c)\left[\frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu}e^{-(\lambda+\mu)t}\right] dt.$$

The value of quality is the difference between (D.7) and (D.13):

$$\Delta(x_s) = \frac{\mu}{\lambda+\mu}\int_{t=0}^{\infty} x_{t+s}e^{-rt}\left[1 - e^{-(\lambda+\mu)t}\right] dt - \frac{\mu}{r+\lambda}\int_{t=0}^{\infty} x_t e^{-rt}\left[\frac{\lambda}{\lambda+\mu} + \frac{\mu}{\lambda+\mu}e^{-(\mu+\lambda)t}\right] dt.$$

where cost terms cancel since both high- and low quality firms always work. This equation parallels equation (D.3) in the good-news case. When $s = 0$, this becomes

$$\Delta(0) = \frac{\mu}{\lambda+\mu}\int_{t=0}^{\infty} e^{-rt}x_t\left[\frac{r}{r+\lambda} - \frac{r+\mu+\lambda}{r+\lambda}e^{-(\lambda+\mu)t}\right] dt. \tag{D.14}$$

Here, the value of quality realizes when a μ-shock hits before the first λ-shock.

References

Abreu, D., P. Milgrom, and D. Pearce (1991): "Information and Timing in Repeated Partnerships Information and Timing in Repeated Partnerships," *Econometrica*, 59(6), 1713–1733.

Bar-Isaac, H., and S. Tadelis (2008): "Seller Reputation," *Foundations and Trends in Microeconomics*, 4(4), 273–351.

Benabou, R., and G. Laroque (1992): "Using Priviledged Information to Manipulate Markets: Insiders, Gurus and Credibility," *Quarterly Journal of Economics*, 107(3), 921–958.

Cabral, L., and A. Hortaçsu (2009): "The Dynamics of Seller Reputation: Theory and Evidence from eBay," forthcoming.

Cohen, W. M., and D. A. Levinthal (1990): "Absorptive Capacity: A New Perspective on Learning and Innovation," *Administrative Science Quarterly*, 35(1), 128–152.

Cripps, M. W., G. Mailath, and L. Samuelson (2004): "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, 72(2), 407–432.

Diamond, D. W. (1989): "Reputation Acquisition in Debt Markets," *Journal of Political Economy*, 97(4), 828–862.

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

Fernandes, A., and C. Phelan (2000): "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory*, 91(2), 223–247.

Fudenberg, D., D. Kreps, and E. Maskin (1990): "Repeated Games with Long-Run and Short-Run Players," *Review of Economic Studies*, 57(4), 555–573.

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

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

Jarque, A. (2008): "Repeated Moral Hazard with Effort Persistence," Working paper, Federal Reserve Bank of Richmond.

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

Keller, G., and S. Rady (1999): "Optimal Experimentation in a Changing Environment," *Review of Economic Studies*, 66(3), 475–507.

Liu, Q. (2009): "Information Acquisition and Reputation Dynamics," Working paper, University of Pennsylvania.

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

Sannikov, Y., and A. Skrzypacz (2007): "Impossibility of Collusion under Imperfect Monitoring with Flexible Production," *American Economic Review*, 97(5), 1794–1823.

Tadelis, S. (1999): "What's in a Name? Reputation as a Tradeable Asset What's in a Name? Reputation as a Tradeable Asset," *Journal of Political Economy*, 89(3), 548–563.