Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 16-014

"Promoting a Reputation for Quality"

by

Daniel Hauser

http://ssrn.com/abstract=2845375

Promoting a Reputation for Quality*

Daniel N Hauser†

September 29, 2016

Abstract

I consider a model in which a firm invests in both product quality and in a costly signaling technology, and the firm's reputation is the market's belief that its quality is high. The firm influences the rate at which consumers receive information about quality: the firm can either promote, which increases the arrival rate of signals when quality is high, or censor, which decreases the arrival rate of signals when quality is low. I study how the firm's incentives to build quality and signal depend on its reputation and current quality. The firm's ability to promote or censor plays a key role in the structure of equilibria. Promotion and investment in quality are complements: the firm has stronger incentives to build quality when the promotion level is high. Costly promotion can, however, reduce the firm's incentive to build quality; this effect persists even as the cost of building quality approaches zero. Censorship and investment in quality are substitutes. The ability to censor can destroy a firm's incentives to invest in quality, because it can reduce information about poor quality products.

KEYWORDS: Reputation, Advertising, Promotion, Censorship, Dynamic Games
JEL: C73, D82, D83, D84

*Thanks to Aislinn Bohren, George Mailath, and Steven Matthews for their advice and support. I'm also grateful for suggestions from Simon Board, David Dillenberger, Mallesh Pai, Alessandro Pavan, Andrew Postlewaite, Philipp Strack, Rakesh Vohra, Yuichi Yamamoto and audiences at the Midwest Economic Theory Conference, North American Summer Meetings, GAMES 2016, the Stony Brook Game Theory Festval, the University of Pennsylvania Theory Lunch, Theory Reading Group, and Empirical Micro Reading Group.

†Email: dahauser@sas.upenn.edu; University of Pennsylvania

1 Introduction

In many markets, such as smartphones and television series, the quality of the next version of the product is not observed until after purchase. When the quality is persistent, the firm's past products provide consumers with important information about future quality. Therefore, a firm has an incentive to build a reputation for producing high quality products, which incentivizes the firm to invest in quality. Concurrently, a firm informs consumers about the quality of its product to establish a reputation. The firm can promote high quality by publicizing good reviews or permitting customers to try out the product, and hide poor quality by concealing bad reviews, failed safety tests or other negative information about the product. The firm's ability to control information about its product greatly impacts a firm's ability to build and maintain a reputation.

In the reputation literature, information is treated as exogenous. Liu (2011), Liu and Skrzypacz (2014), Hu (2015) and Board and Meyer-ter-Vehn (2013) show how the information consumers see influences a firm's incentives to build and maintain reputation. However, firms often have a great deal of control over the information consumers are able to observe. It seems natural for the information a firm reveals to be endogenous. In this paper, I allow firms to strategically choose the information consumers observe about quality and show that this plays a crucial role in the incentive to build and maintain a reputation. In equilibrium, the firm forgoes investing in quality at high reputations to instead promote. The firm's ability to censor information can completely destroy its incentive to build quality. The firm may benefit from committing to an information structure that releases information independently of the firm's reputation. This provides a novel explanation for why firms may want to hire an outside firm to handle advertising or monitor the firm for ethical violations.

In the model, a long lived firm sells a product to short lived consumers. The product's quality changes at a Poisson rate λ, and is fixed between arrivals, as in Board and Meyer-ter-Vehn (2013). Quality is not observed by consumers; instead they observe a Poisson process with an arrival rate that depends on both the quality of the product and the firm's actions. Whenever an arrival occurs, consumers observe the quality of the firm's product. The firm can promote its product by increasing the intensity of this news process or censor information by decreasing the intensity of this process. Consumer beliefs about past investment and past quality play a crucial role in the firms current payoff. A consumer is more willing to buy a product today if she thinks it was a high quality product yesterday, so the firm has incentives to both invest in quality and in controlling the information consumers observe in order to maintain its reputation.

I first consider the good news case, where the firm selects the arrival rate of news that the product is high quality. In this case, the firm chooses to generate signals only when it is producing a high quality product. This creates incentives for the firm to invest in quality, since the firm knows that it will then be able to promote the product and sell it at a higher price. The ability to promote can create periods of time where reputation cycles, but the firm is solely investing in promotion and not in quality. But, the ability to control information can also damage incentives and hurt the firm's payoffs. Investing in quality and investing in promotion are complements, in that a firm only benefits from investing in quality if it also promotes. The need to promote makes investing more costly, because even if the firm creates a high quality product, it still has to promote it in order to benefit from the quality breakthrough. The firm is not only negatively impacted by the cost of promotion, but is also damaged by lower consumer beliefs. Endogenous promotion decreases a firm's incentives to invest in quality, which in turn leads to lower consumer beliefs and lower payoffs for the firm. Investment in quality is under-incentivized. Even as investing in quality becomes arbitrarily cheap the firm still shirks when it has a high reputation.

I focus on Markov Perfect Equilibrium (MPE) in the firm's reputation. Every MPE of the good news game is characterized by three regions. When consumer beliefs are high, the firm chooses not invest in quality or promotion. At intermediate beliefs, the firm promotes if it is selling a high quality product but it does not choose to invest in quality. The firm will never invest in quality when it wouldn't promote if it had a high quality product, because there is no way for consumers to detect or punish low levels of investment if they do not expect a firm with a high quality product to be promoting. Finally, at low beliefs, the firm promotes when it has a high quality product, and the firm invests in quality independent of current quality. This equilibrium structure is distinct from the equilibrium in Board and Meyer-ter-Vehn (2013) and Marinovic et al. (2015), who consider a similar model. In this intermediate region, the firm invests in promotion to restore its high reputation. This insight applies to many economic settings, such as when a video game company releases a special edition of a popular game; a studio launching a second promotional campaign when a movie is released on DVD; or a tech company releasing a new variant of a popular product. In the good news case, the firm utilizes its ability to control promotion and the quality of its product to create and maintain its reputation.

Next, I consider the bad news case, where bad news about the product arrives at a fixed rate, and the firm can suppress this information. Censorship and investment are substitutes. A firm can reduce bad news by either investing in quality or censoring the news. This can cause incentives to disintegrate. The firm is often willing to forgo investment in quality, because it knows censorship is a viable alternative. Oil companies or cigarette manufacturers are prime examples, who instead of investing in better, safer products, they have invested significantly in suppressing and undermining research that proves their products are harmful.

In the bad news case, any MPE can be characterized by two reputation cutoffs. The firm invests in quality when its reputation is higher than a cutoff and doesn't invest otherwise. Similarly, the firm censors bad news when its reputation is sufficiently high, and doesn't otherwise. When censorship is cheaper than investment, the firm opts not to invest in quality. Instead, it opts to censor bad news when its reputation is high. Unless investment in quality is sufficiently less expensive than censorship, the firm will always substitute censorship for investment in quality, and therefore never invests in the quality of its product. In these equilibria, the firm's reputation is transient. Consumers have no reason to believe the firm is investing in its product, and therefore are skeptical of the product. This creates a situation where in the firm would prefer to commit to never censor, and, in fact, may benefit from making it more costly to censor.

Literature Review

There is an extensive literature on reputations, starting with Kreps et al. (1982), and continuing in papers like Mailath and Samuelson (2001), Cripps et al. (2004), and Faingold and Sannikov (2011). These papers model reputation as the consumer's belief about the firm's type, which can either be a strategic type or a commitment type. These reputational concerns allow the firm to achieve payoffs that exceed the payoffs of the game where the consumers know the firm's type. Cripps et al. (2004) makes the observation that in many cases reputation effects vanish over time, and consumers can eventually determine the firm's type.

Cripps et al.'s (2004) observation has led to a large body of literature on when a firm can have a persistent reputation or when a firm's reputation cycles. Liu (2011), Liu and Skrzypacz (2014), Hu (2015) and Dilmé (2014) consider models similar to the standard reputation model with commitment types, and show that with some modifications to the environment reputations no longer vanish. Dilmé (2014) adds switching costs which add some persistence to the firm's action. Liu (2011) considers an environment where it is costly for consumers to observe the history. In Liu and Skrzypacz (2014) and Hu (2015), consumers only observe part of the history. This new information structure enhances reputation building by making it more difficult, or impossible, for consumers to eventually detect strategic types. These papers illustrate the importance of the information consumers see, as different information structures can have very different affects on how a firm builds and maintains a reputation.

There is also a literature on designing in a information structure to encourage reputation building. Dellarocas (2005) considers an environment where a long lived seller selling to short lived buyers has a moral hazard problem, and the designer designs a signal to help resolve this problem. This persistent signal improves the sellers incentives to exert effort, since they will be rewarded in future periods for high values of the signal, and punished for low values. Ekmekci (2011) considers designing a rating system for a product choice game, where the agent could either be a strategic type or a commitment type. Hörner and Lambert (2015) considers a similar problem in a career concerns environment and in addition allows the designer to choose multiple signals. In these papers, a precise choice of information structure can reduce the moral hazard problem. Finally, Pei (2015) considers how a firm can control the information the market sees about a worker in order to keep the worker's outside option low. Similar to this paper, he finds that the ability to hide information reduces the firm's ability to induce high levels of effort.

A second set of papers resolves the issue of vanishing reputation by making the firm type endogenous and persistent. Bohren (2012) models reputation as a publicly observable state variable driven by a Brownian motion that depends on the players action, which allows for reputation cycles. Board and Meyer-ter-Vehn (2013) models a firm investing in its products quality. Quality shocks arrive at a Poisson rate, and the quality of a product after a shock depends on the effort the firm takes. They then characterize how a firm builds and maintains a reputation for selling a high quality product. This reputation does not necessarily vanish, because the type is constantly evolving, and it incentivizes costly effort. Since effort today can determine future quality, the firm wants to exert effort because it will potentially allow them to benefit from a higher reputation in the future. Moreover, Board and Meyer-ter-Vehn (2013) considers a variety of information structures, and finds that incentives, and how persistent vs vanishing reputation depend crucially on the information consumers observe. Board and Meyer-ter-Vehn (2015) and Halac and Prat (2014) consider similar models.

The paper that most relates to this one is Marinovic et al. (2015). This paper considers a similar environment, where quality of the product evolves as in Board and Meyer-ter-Vehn (2013), but the firm can now deterministically reveal its quality at any instant of time, for a cost. They find that every Markov Perfect Equilibrium has many counterintuitive features. In their model, making it cheaper for the firm to reveal its type always lowers payoffs, and the firm almost never benefits from a higher reputation. They resolve this issue by showing that by using time as a state variable, they can significantly enrich the state space, and support higher payoffs through a more elaborate system of punishments and rewards. My paper demonstrates that to some extent, the counter-intuitiveness of the Markov Perfect Equilibria reflect their modeling assumptions, and break down to some extent if the firm's advertisements convince consumers that the firm is selling a high quality product at random times. Unlike in their model, the firm can invest (and will) invest in promotion and not quality for a while, and its ability to promote allows the firm to extend the period it collects reputational dividends, which cannot happen in Markov Perfect Equilibria in the Marinovic et al. (2015) model.

2 Model

The Firm

Time is continuous, $t \in [0, \infty)$. There is a single long lived firm with stochastic quality $\theta_t \in \{L, H\}$. The firm has discount rate r. At each instant of time, the firm chooses an action, which consists of a level of effort $a_t \in [0, 1]$ and a level of promotion $\pi_t \in [0, \bar{\pi}]$, for costs ca_t and $k\pi_t$, where $c > 0$, $k > 0$ and $\bar{\pi} > 0$. The upper bound on promotion, $\bar{\pi}$, is a fixed and measures the maximum arrival rate of news generated by the firm. A firm's strategy is a stochastic process $(a_t, \pi_t)_{t=0}^{\infty}$ that determines the effort choice and level of promotion at each instant of time given the history the firm has observed. These strategies are predictable processes with respect to the σ-algebra generated by the quality Poisson process and the news Poisson process. Quality evolves via Poisson shocks, as in Board and Meyer-ter-Vehn (2013). Specifically, there is a Poisson process with intensity $\lambda > 0$. Whenever there is an arrival of this process, θ_t becomes H with probability a_t and L with probability $1 - a_t$, and is fixed between arrivals. The firm observes θ_t.

Consumers

Consumers do not observe θ_t, a_t or π_t directly. Consumers observe a news process, which provides a noisy signal that they use to form beliefs about θ_t. I consider two news processes, good news and bad news. The news process has intensity π_t in the good news case, and intensity $(\bar{\pi} - \pi_t)1_{\theta_t=L}$ in the bad news case, where $\bar{\pi}$ is the rate bad news would arrive at absent any censorship. When a signal arrives at time t, the firm observes θ_t. Let $x_t = Pr(\theta_t = H|\mathscr{F}_t^s)$, where $\mathscr{F}_t^s$ is the σ-algebra generated by the news process, and the probability measure is the measure induced by the consumers' beliefs about the firm's strategy.

Payoffs

The firm receives a flow payoff of x_t. This can be motivated either as the willingness to pay of consumer with utility $1_{\theta_t=H}$ or as consumers who are willing to pay 1 arriving at some rate proportional to the public belief. The firm maximizes

$$\max_{\hat{a},\hat{\pi}} E_{\hat{a},\hat{\pi}} \left(\int_0^\infty e^{-rt}[x_t - c\hat{a}_t - k\hat{\pi}_t]dt \right)$$

where the expectation is taken with respect to the actual probability measure induced by the firm's chosen effort and promotion levels, while x_t is determined by what consumers believe about the firm's effort and promotion choices.

Solution Concept

I characterize Markov Perfect Equilibrium in consumers' beliefs and current quality.

Definition 1 *A pure strategy Markov Perfect Equilibrium (MPE) consists of an effort strategy $a : [0,1] \times \{H, L\} \to [0,1]$, which maps public beliefs to effort, and a promotion strategy $\pi : [0,1] \times \{H, L\} \to [0, \bar{\pi}]$ such that (a, π) maximize payoffs given x_t and x_t is formed via Bayes Rule.*

An additional admissibility restriction must be placed on beliefs in order to ensure the belief process has a unique solution in the MPE. These restrictions are identical to the restrictions placed in Board and Meyer-ter-Vehn (2013). If beliefs follow the law of motion $\dot{x}_t = g(x)$, then the drift $g(x)$ must satisfy one of the following conditions:

1. $g(x) = 0$
2. $g(x) > 0$ and $g(x)$ is right continuous at x,
3. $g(x) < 0$ and $g(x)$ is left continuous at x,

at any cutoff and beliefs can be partitioned into a finite set of intervals such that both the effort and promotion choices are Lipschitz continuous on the interior of all these intervals.

These restrictions are placed on consumers beliefs about the firm's strategies, not on the firm's strategies themselves. Admissibility ensures that $\dot{x} = g(x)$ has a solution and when there are multiple solutions, I select the one consistent with the discrete time approximation (for details, see Board and Meyer-ter-Vehn (2013) or Klein and Rady (2011)). As in Board and Meyer-ter-Vehn (2013), this is a relatively mild assumption that ensures beliefs are defined everywhere and are right continuous when viewed as a function of t.

Discussion of the Model

There are many ways to interpret the quality process. For instance, arrivals can be viewed as the firms ability to incorporate new discoveries, or as the firm purchasing a rate of technological improvement or abating deterioration of the firms technology. While the technology process is stylized, it satisfies many intuitive properties of persistent quality. Beliefs drift down when the firm is believed to be shirking and up when the firm is believed to be working absent any promotion, and the firm's incentives to exert effort depend both on present and future quality.

The news process is similarly stylized. The assumption that arrivals of news can only occur in when the firm has a high quality product in the good news case, or when the firm has a low quality product in the bad news case is a strong assumption. This can be viewed as a requirement that a successful promotional campaign contain evidence that the firm's product is actually good, and that a negative news story needs evidence that the firm's product is actually bad. The Poisson structure captures large jumps in beliefs about quality. For example, an actor winning an academy award after a production company's promotional campaign, or a news story about a popular brand of tires exploding causes beliefs to jump significantly. This process abstracts away from infinitesimal information consumers receive constantly about a product, for instance, a consumer continuously updating beliefs about Apple while they are using their iPod, which is information that seems more difficult for the firm to control.

This information process and $\bar{\pi}$, the arrival rate of news, are important objects in this model and the analysis. I view $\bar{\pi}$ as a measure of how difficult it is to successfully create an ad campaign (or some other sort of signal) that convinces consumers that the firm indeed is selling a high quality product. This upper bound encapsulates both the difficulty of creating a promotional campaign, how long it takes for the campaign to permeate the public consciousness and the difficulty of actually persuading consumers, so it is in some ways a measure of persuasiveness.

3 The Good News Case

In the good news case, the signal process has intensity $\pi_t 1_{\theta_t = H}$. Arrivals can only occur when the firm has a high quality product, in which case, beliefs jump to 1. Between arrivals, beliefs follow the law of motion

$$\dot{x}_t = \underbrace{\lambda(\tilde{a}_t - x_t)}_{\text{Quality Breakthroughs}} - \underbrace{\tilde{\pi}_t x_t(1 - x_t)}_{\text{Absence of signals}},$$

where $\tilde{a}_t$ is the believed amount the firm is investing in quality and $\tilde{\pi}_t$ is the believed level of promotion.

In between arrivals, the drift of consumer's beliefs can be expressed as the sum of two terms. The firm's effort choice determines to the first term $(\tilde{a}_t - x_t)\lambda$. Depending on how much effort the firm is believed to be exerting, this term is either positive or negative depending on whether it is more likely for a low quality firm to switch to a high quality firm or for a high quality firm to switch to a low quality firm. The second term, $-\tilde{\pi}_t x_t(1 - x_t)$ is determined by the firm's choice of promotion. This term is negative if consumers believe the firm is promoting, since it is more likely that no news arrives if the firm is selling a low quality product, and 0 if the firm is not promoting.

The firm's value function is

$$V(x, \theta) = \max_{a,\pi} E\left(\int_0^\infty e^{-rt}[x_t - ca_t - k\pi_t]dt \middle| \theta_0 = \theta, x_0 = x \right).$$

This can be rewritten as

$$\begin{aligned} V(x_0, \theta) = \max_{a,\pi} \int_0^\infty e^{-\int_0^t (r+\pi(x_s,\theta)+\lambda)ds}[x_t + a(x_t, \theta)(\lambda(V(x_t, H) - V(x_t, L)) - c) + \lambda V(x_t, L) \\ + \pi(x_t, \theta)1_{\theta=H}(V(1, \theta) - k)]dt, \end{aligned}$$

where x_t is the solution to $\dot{x}_t = \lambda(\tilde{a}_t - x_t) - \tilde{\pi}_t x_t(1 - x_t)$. As in Board and Meyer-ter-Vehn (2013), this can be rewritten as

$$V(x_t, \theta) = \max_{a,\pi} \int_0^\infty x_t + a(x_t, \theta)(\lambda D(x_t) - c) + \pi(x_t, \theta)(\Delta(x_t) - k) + \lambda V(x_t, L) - \lambda V(x_t, \theta) - rV(x_t, \theta)dt,$$

where $D(x) = V(x, H) - V(x, L)$ and $\Delta(x) = V(1, H) - V(x, H)$. This implies the following lemma

Lemma 1 *(**Sequential Rationality**) (a, π) are a MPE if and only if $a(x, \omega)$ solves*

$$\max_{a\in[0,1]} \lambda D(x)a - ca,$$

$\pi(x, H)$ solves

$$\max_{\pi\in[0,\bar{\pi}]} \Delta(x)\pi - k\pi,$$

and $\pi(x_t, L) = 0$ for all x.

The firm's incentive to exert effort is driven by $D(x)$, the change in the firm's expected payoffs if a quality change arrives at that instant, while the firm's incentive to promote is driven by $\Delta(x)$, the change in payoffs if a beliefs jumped at that instant. An important feature of this model is that the effort choice is independent of the firm's type, which greatly simplifies the analysis.

$D(x)$ and $\Delta(x)$ have a tight relationship, since the increase in a firms payoff from having high quality is due to the potential reputational benefits the firm receives in the future once it is selling a high quality product. This is captured in the following proposition.

Lemma 2 *The difference in expected payoffs between a firm selling a high quality product and a firm selling a low quality product can be written as*

$$D(x_0) = \int_0^\infty e^{-(r+\lambda)t}\pi(x_t, H)[\Delta(x_t) - k]dt.$$

Moreover, payoffs satisfy the following properties for any beliefs

1. $V(\cdot, H)$ *and* $V(\cdot, L)$ *are strictly increasing.*
2. $V(x_t, H) \geq V(x_t, L)$ *for all* x_t.
3. $\Delta(x)$ *is decreasing.*
4. $D(x)$ *is decreasing.*

The value functions are increasing. Given any two initial conditions x_0 and x_0^*, if $x_0 > x_0^*$ then $x_t > x_t^*$ until the first arrival. The firm can then mimic the $*$-firm and induce the same probability measure over θ_t and arrivals at all points, while receiving a strictly higher flow payoff, so its payoff must be higher. Similarly, $V(x_t, H) \geq V(x_t, L)$ because a high quality firm can play the exact same strategy as the low quality firm to guarantee themselves the same payoff. This directly implies that $\Delta(x)$ is decreasing, and this, combined with the optimality of π and the previous observation about x_t and x_t^*, implies that $D(x)$ is decreasing.

The previous lemmas imply that optimal strategies are cutoff strategies. The firm shirks and doesn't promote at high beliefs and as beliefs drift down eventually start promoting and working. This greatly simplifies the existence argument and implies that in any equilibrium beliefs can be categorized into four different regions, the region where the firm is promoting or working, the region where the firm is promoting but not working, the region where the firm is working but not promoting and the region where the firm is working and promoting. The following proposition establishes the unique structure of any Markov Perfect Equilibrium.

Proposition 1 *A Markov Perfect Equilibrium exists. Moreover, every Markov Equilibrium is characterized by two cutoffs* $0 \leq x^* \leq x^{**} < 1$ *where optimal strategies satisfy*

$$a(x) = \begin{cases} 1 \text{ if } x < x^* \\ 0 \text{ if } x > x^* \end{cases},$$

$$\pi(x,H) = \begin{cases} \bar{\pi} \ \textit{if}\ x < x^{**} \\ 0 \ \textit{if}\ x > x^{**} \end{cases},$$

and $\pi(x, L) = 0$*. When* $\bar{\pi} \leq \lambda$*,* $x^* < x^{**}$ *in any MPE where* $x^{**} > 0$*.*

As beliefs decrease, the firm always begins promoting before it starts working. If there was a region where consumers believed the firm was working and not promoting, their beliefs would drift up, so it would never have any incentive to promote once beliefs were high enough. But if this was the case, then the firm would have no incentive to work, since the information consumers expect to see wouldn't change if the firm shirked. So, the firm must start promoting before it starts working. This can lead to long cycles where the firm builds and maintains reputation through promotion without exerting effort. After a quality breakthrough and a successful promotion, the firm collects reputational dividends and tries to use its ability to promote to extend the period of time when it collects these dividends and without investing in quality.

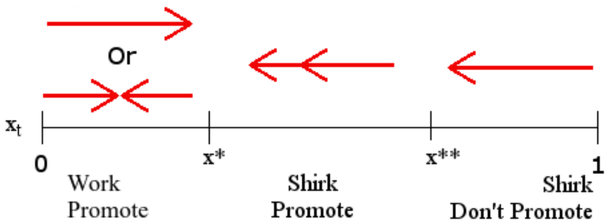

Figure 1: Structure of a MPE

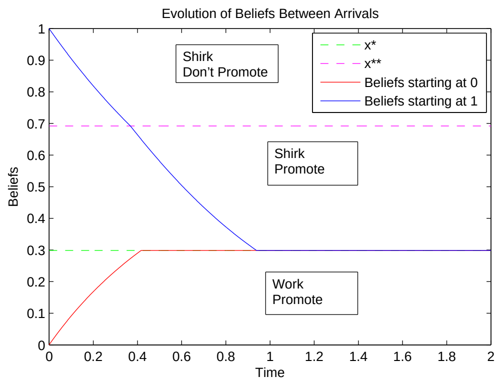

Figure 2: Belief Dynamics between arrivals

The result that the firm never exerts effort before it starts promoting is also present in Marinovic et al. (2015). But, unlike in Marinovic et al. (2015), this is not always bad for the firm, since beliefs no longer jump to 0 as soon as consumers believe the firm would be promoting but don't see any news. In that paper, promotion never has any value to the firm in equilibrium, the best possible case is that the marginal benefit from promotion is equal to its cost. As opposed to in the MPEs in Marinovic et al. (2015), the firm exerts effort when beliefs are non-zero and benefits from consumers giving them the benefit of the doubt when they don't see any news.

I consider two different types of commitment to investigate how endogenous promotion changes firm payoffs. I first examine how the game with endogenous promotion compares to the game with exogenous promotion from Board and Meyer-ter-Vehn (2013). In this game, the firm pays a lump sum at the beginning of the game in order to commit to promoting whenever they have a high quality product for no flow cost[1]. I then consider what strategy promotion strategy the firm would like commit in the game with endogenous promotion. Since the firm's promotion strategy was not observable by consumers, the firm is unable to fully internalize the impact it's choice of promotion has on the drift of beliefs and on the level of investment. By allowing the firm to commit to a promotion strategy, I can investigate whether the firm is over or under-investing in promotion.

The ability to promote creates a moral hazard problem. The firm needs to be promoting before it has any incentive to invest in its product. Consider the alternative game from Board and Meyer-ter-Vehn (2013), where the firm cannot control the news process, instead consumers learn by observing an exogenous news process with arrival rate $\bar{\pi}1_{\theta_t=H}$.

Board and Meyer-ter-Vehn (2013) show that in the game with exogenous promotion, all MPE have the following form

$$a(x,\theta) = \begin{cases} 1 \text{ if } x < x^*_{BM} \\ 0 \text{ if } x > x^*_{BM} \end{cases}$$

for some cutoff $x^*_{BM} \in [0,1]$. Comparing this cutoff to the cutoff in when the firm can endogenously promote allows me to look at how the ability to endogenously choose the level of promotion changes how reputational concerns incentivize the firm to invest in quality.

Proposition 2 *If $\lambda \geq \bar{\pi}$ then the equilibrium of the game with exogenous promotion is unique and $x^*_{BM} \geq x^*$ (and is strict as long as $x^*_{BM} \neq 0$).*

So, the firm's incentives to invest in quality are reduced when compared to the game studied in Board and Meyer-ter-Vehn (2013). While the firm still eventually invests in quality, it starts investing at lower reputations, so the incentive to exert effort to avoid a bad reputation is diminished. The ability to choose when to promote decreases the gap between the payoffs received by a firm selling a high quality product and a firm selling a low quality product, so the firm chooses to start working at a lower belief. This does not necessarily imply the firm is always investing less in quality, because beliefs move in a fundamentally different way in the two models. While the firm with exogenous promotion starts working at a higher belief, it expect beliefs to jump to 1 more, since arrivals can occur anywhere, and these arrivals delay when it has to begin investing in effort. On the other hand, when the firm has a low quality product, it will begin investing sooner in the game with exogenous promotion than in the game with endogenous promotion, because no news can arrive to delay investment.

The firm receives some benefit from exogenous promotion. For instance, firm selling a low quality product when beliefs are close to 1 receives higher flow payoffs for longer

[1]If promotion still had a flow cost, the firm may never want to invest in quality because it doesn't want to pay for promotion.

because consumers no longer make inference from not seeing any news, but there also is a cost. After news arrives, consumer's beliefs can drop to a lower point than it would in the game with exogenous promotion, which can lower the firm's payoffs. This suggests that a firm may benefit from committing to exogenous promotion, by hiring an outside advertising agency, in order to incentivize effort at higher beliefs.

When promotion is exogenous, consumers learn that that the firm is selling a high quality product more frequently.

Corollary 1 *Let $x_0 = 1$, and $\lambda \geq \bar{\pi}$. Let $\tau = \inf\{t > 0 : x_t = 1\}$, and $\tau_{BM} = \inf\{t > 0 : x_{BM,t} = 1\}$. Then $E(\tau) > E(\tau_{BM})$.*

The expected time it takes for reputation to return to 1 is smaller when the firm has committed to exogenous promotion. Reputation cycles more frequently when the firm has committed to promote. The increased incentive to invest leads to a higher shirk-work cutoff x^*, which in turn leads to consumers learning that the firm is selling a high quality product more frequently.

This suggests that the firm may benefit overall from committing to promote whenever they have a high quality product, independent of their reputation. If the firm could pay an advertising agency a lump sum in exchange for the ad agency committing to promote, they would benefit from the enhanced incentives for investment in quality.

In equilibrium, the firm is not fully internalizing the impact of their choice on promotion on the law of motion for beliefs and on their choice of investment. To explore this effect more, suppose the firm could commit to any cutoff x^{**}_{FB} ex-ante. With this commitment power, the firm internalizes its ability to control the drift of beliefs, and how its choice determines what investment strategies are credible.

Proposition 3 *Suppose $\lambda \geq \bar{\pi}$ and $\theta_0 = H$. If the firm could commit to promoting below any cutoff, they would commit to promoting at a cutoff $x^{**}_c \leq x^{**}$ for any x^{**} that is a promotion cutoff in a MPE of the original game. In the game where the firm has commited to the optimal promotion strategy, it chooses to invest in quality at a higher reputation, $x^*_c \geq x^*$ for any x^* that is a investment cutoff in the MPE of the original game.*

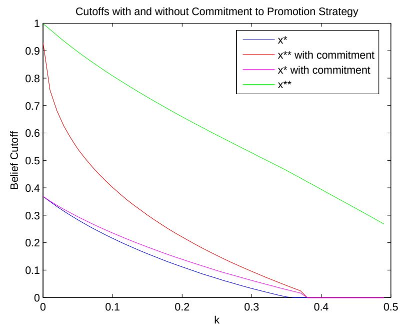

Figure 3: Optimal level of investment and promotion with and without commitment.

The firm commits to a higher level of promotion because they now are internalize the impact their promotion has on the drift of beliefs. This in turn, can increase the incentives to invest, by raising the payoff from a high reputation. The firm is under-investing in quality and over-investing in promotion. The firm would like to be promoting less, in order to prevent consumers from interpreting lack of news as a bad sign, but without commitment power, it cannot convince consumers that it wouldn't promote at a higher reputation.

In order to capture how the firm's ability to promote impacts incentives, it is interesting to consider the limiting case as $c \to 0$. This simplifies the firm's effort choice, and allows me to focus on the firm's promotion decision and how that decision impacts payoffs. In addition, while there may be multiple equilibrium in this model, these equilibrium all converge to the same limiting equilibrium as $c \to 0$, which allows me to characterize some comparative statics.

Proposition 4 *As $c \to 0$, if $\lambda \geq \bar{\pi}$,*

1. $|x^* - x^{**}| \to 0$
2. *There is a unique point $\bar{x}^{**}$ such that for any sequence of costs $(c_n)_{n=1}^{\infty}$ and a corresponding sequence of equilibria, characterized by cutoffs $(x_n^*, x_n^{**})_{n=1}^{\infty}$, $x_n^{**} \to \bar{x}^{**}$. moreover $\bar{x}^{**} < 1$.*
3. *In the limiting equilibrium, as k decreases payoffs increase.*

So, unlike in Marinovic et al. (2015), decreasing k doesn't always decrease firm payoffs. When costs are sufficiently low, the firm benefits from cheaper promotion, because it increases the firm's incentives to work. The gap between the work cutoff and the promotion cutoff vanishes, but the firm still only works after they've begun promoting, even when costs are arbitrarily low, since the firm can never credibly work unless consumers expect them to also promote when they have a high quality product. So even as effort becomes arbitrarily cheap, the firm still doesn't work everywhere. The firm still has to pay for promotion to make effort worthwhile, which indirectly makes effort costly.

In this limit, I can compare the payoffs from the game with exogenous promotion to the game with endogenous promotion.

Proposition 5 *As $c \to 0$, if $\lambda \geq \bar{\pi}$, $V_{BM}(x,\theta) > V(x,\theta) + E\left(\int_0^{\infty} e^{-rt} k\pi_t | x_0 = x, \theta_0 = \theta\right)$, where V_{BM} is the value function from the game where news arrives at exogenous rate $\bar{\pi} 1_{\theta_t = H}$.*

So the firm would be willing to pay more than its total expected expenditure on promotion on order to commit to promoting at rate $\bar{\pi}$ whenever it has a high quality product. In this limit, the firm is hurt by the ability to control promotion because it leads to worse consumer beliefs. The firm benefits from hiring an advertising agency, even if the agency doesn't have any sort of comparative advantage or superior technology.

The firm benefits simply because the ad agency can be used as a commitment device, convincing consumers that the firm will promote whenever it has a high quality product.

The assumption that $\lambda \geq \bar{\pi}$ makes taking the limit much easier. From a technical perspective, this assumption implies that at x^* the drift is always 0 at any equilibrium. In turn, this fact is enough to ensure the uniqueness of the limiting equilibrium for any sequence of equilibria as $c \to 0$. This assumption means that the quality of the firm's product changes faster than the firm can signal that it has a high quality product. This is a natural assumption for a variety of applications. For instance, an athlete's ability is relatively responsive to how much effort they exerted at practice that season, while our beliefs about his or her quality change relatively infrequently. Tech firms are constantly innovating, changing both hardware and software for a new phone, while our beliefs that the next phone is going to be good or bad seem to change relatively slowly.

The ability to promote also can lower the firm's payoffs by causing beliefs to drift down faster. If the firm could commit to never promote, beliefs would always follow the law of motion $\dot{x}_t = \lambda(\tilde{a}_t - x_t)$ instead of drifting down at rate $\dot{x}_t = \lambda(\tilde{a}_t - x_t) - \tilde{\pi}_t(x_t)(1 - x_t)$, but the firm also loses the ability to benefit from investing in quality, since consumers no longer see any news. In the limit, as $c \to 0$, I can characterize which of these effects is larger.

Corollary 2 *As $c \to 0$, if $\lambda \geq \bar{\pi}$, payoffs in the limiting equilibrium are greater than payoffs in the equilibrium where the firm has committed to never promote.*

The firm still does benefit from promotion, even though it makes effort costly. Since promotion allows consumers to distinguish between high and low quality firms, it gives firms the ability to build a reputation. Without this promotion technology, the firm would have no incentive to work, because shirking would be undetectable. These results contrasts with Marinovic et al. (2015), where in any Markov Perfect Equilibria, payoffs are decreasing as k decreases and the firm would rather commit to not promoting in any MPE. The deterministic nature of the firm's signaling technology in their model forces the firm to signal so frequently that the firm no longer benefit at all from building a reputation.

4 The Bad News Case

In this section I consider the case where consumers learning about the firm's quality through arrivals of bad news, and the firm tries to censor this bad news. Unlike promotion, censorship is a substitute for effort and can cause incentives to completely break down.

The difference between the good news and bad news model is the news process. Now the news process has arrival rate $\max(0, \bar{\pi}1_{\theta_t = L} - \pi_t)$ and $\pi_t \in [0, \bar{\pi}]$. Consumer beliefs now follow law of motion

$$\dot{x}_t = \underbrace{\lambda(\tilde{a}_t - x_t)}_{\text{Quality breakthroughs}} + \underbrace{(\bar{\pi} - \tilde{\pi}_t)x_t(1 - x_t)}_{\text{Absence of news}}.$$

Unlike in the good news case, the signaling term now causes beliefs to drift up. If consumers believe the firm is not censoring, not seeing any news is more likely to occur when the firm is selling a high quality product. When news arrives, consumers learn that $\theta_t = L$ and adjust their beliefs to 0.

Incentives to work and censor are now driven by the firm's fear of consumers figuring out that it is selling a low quality product. As before, the firm's value function can be written as

$$\begin{aligned} V(x, \theta) = \int_0^\infty e^{-(\int_0^t r+\lambda+\bar{\pi}-\pi_s ds)}[x_s &+ \bar{\pi}1_{\theta_t=L}(V(0,L) - V(x_t, L)) \\ &+ \pi_t(V(x_t, L) - V(0, L) - k)) \\ &+ a_t(\lambda(V(x_t, H) - V(x_t, L)) - c) \\ &+ \lambda(V(x_t, L) - V(x_t, \theta)) - rV(x_t, \theta)]dt. \end{aligned}$$

The firms incentives to invest in effort and promotion are similar to before. The incentive to invest is still driven by $D(x)$, but now the incentive to promote is driven by $\Delta(x) = V(x_t, L) - V(0, L)$.

Lemma 3 ***(Sequential Rationality)*** *The optimal $a(x_t, \theta)$ solves*

$$\max_{a \in [0,1]} \lambda D(x_t)a - ca$$

and the optimal $\pi(x_t, L)$ solves

$$\max_{\pi \in [0,\bar{\pi}]} \Delta(x_t)\pi - k\pi.$$

This leads to very different effort choices than the good news case. Now, the firm has incentive to work when beliefs are high, because it is worried about consumers finding out that it has a low quality product. Like before, $D(x)$ and $\Delta(x)$ are related.

Lemma 4 *The difference in expected payoffs between a firm selling a high quality product and a firm selling a low quality product can be written as*

$$D(x_0) = \int_0^\infty e^{-(r+\lambda)t}[(\bar{\pi} - \pi(x_t, L))(\Delta(x_t)) - k\pi(x_t, L)]dt.$$

Moreover, payoffs satisfy the following properties for any beliefs

1. *$V(\cdot, H)$ and $V(\cdot, L)$ are strictly increasing.*
2. *$V(x_t, H) \geq V(x_t, L)$ for all x_t.*

3. $\Delta(x)$ *is increasing.*

4. $D(x)$ *is increasing.*

$D(x)$ is increasing in x by similar logic as the good news case. $\Delta(x_t)$ is clearly increasing in x, and π_t maximizes pointwise $\Delta(x_t)\pi_t - k\pi_t$. This is equivalent to minimizing $(\bar{\pi} - \pi_t)\Delta(x_t) + k\pi_t$. Since $\Delta(x_t)$ is strictly increasing in x, so the minimum value of this increasing in x. Moreover, as before, beliefs that start out higher always stay higher, so for any $x > x'$, the integrand is always greater for x_t than for x'_t, so $D(x) > D(x')$. This means, as before, equilibrium can be characterized in terms of cutoffs x^* and x^{**}.

Proposition 6 *A MPE exists. Moreover, in any MPE, there exist cutoffs $x^* \in [0,1]$ and $x^{**} \in (0,1]$ such that equilibrium strategies take the form*

$$a(x,\theta) = \begin{cases} 1 \text{ if } x > x^* \\ 0 \text{ if } x < x^* \end{cases},$$

for $\theta \in \{L, H\}$,

$$\pi(x,L) = \begin{cases} \bar{\pi} \text{ if } x > x^{**} \\ 0 \text{ if } x < x^{**} \end{cases},$$

and $\pi(x,H) = 0$.

Consider any equilibrium where x^* and x^{**} are less than 1. Then there is a region near 1 where beliefs are drifting up, the bad quality firm is censoring and working and the firm never leaves this region. So, in this region $D(x_t) = \frac{k\bar{\pi}}{r+\lambda}$, the difference in the payoff for the high and low quality firms comes purely from the low quality firm having to pay the additional cost of censorship. This naturally implies the following stark result

Proposition 7 *When $(1 + \frac{r}{\lambda})c > k\bar{\pi}$, there is a unique MPE and in that MPE the firm never invests in quality.*

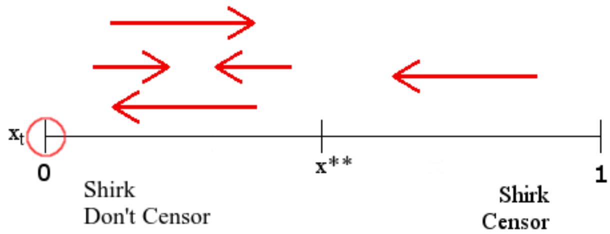

Figure 4: The structure of a full shirk equilibrium

The firm's ability to censor gives it quite a bit of control over how much it loses from quality switching from high to low. Since the firm can always censor bad news after quality switches, it never loses more than $k\bar{\pi}$ at each instant of time from being a low quality firm. So it would never exert effort when $k\bar{\pi}$ is low enough relative to the cost of effort. If censoring bad news is too easy (for instance, cheaper than exerting effort), than the firm would always choose to substitute effort for censorship.

This intuition that leads to the firm never invests in quality holds for much more general classes of equilibria than Markov Perfect. This argument that the firm never invests in quality straightforwardly generalizes to any equilibrium that uses calender time in addition to belief and firm type as the state variables, and in fact it holds for the class of all public sequential equilibria[2]

Corollary 3 *In any Public Sequential Equilibrium, the firm never invests in quality if* $(1+\frac{r}{\lambda})c > k\bar{\pi}$.

This contrasts dramatically from the result in Board and Meyer-ter-Vehn (2013), who find that without the ability to censor bad news, there are equilibria where the firm works. In fact, when $\lambda \geq \bar{\pi}$ and $\frac{\lambda\bar{\pi}}{(r+\lambda)(r+\bar{\pi})} > c$, there are constants $a < b$ such that all $x^* \in [a,b]$ are is a cutoff of an equilbria where the firm works above x^* if the firm was unable to censor. But, giving the firm the ability to censor for a low enough cost completely destroys the firms incentives to work.

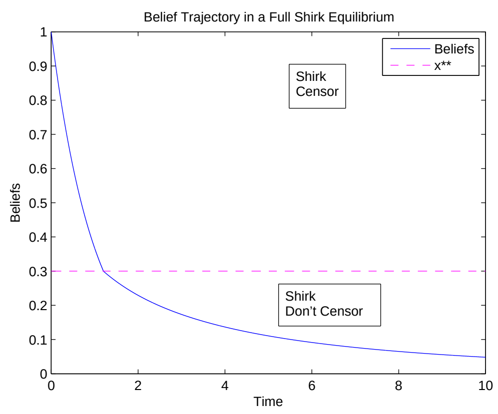

Figure 5: An example of the trajectory beliefs follow until news arrives in a full shirk equilibrium.

This full shirk equilibrium has many undesirable properties from the firm's perspective. While an equilibrium where the firm is working can have persistent reputation building, with probability 1 in a full shirk equilibrium the firm's reputation vanishes in finite time. A firm that could commit to not censoring would not suffer from these sorts of incentive problems. Board and Meyer-ter-Vehn (2013) identifies the possibility of equilibria where the firm works when the bad news technology is exogeneous, but introducing sufficiently cheap censorship technology completely destroys these incentives.

Corollary 4 *Suppose* $(1+\frac{r}{\lambda})c > k\bar{\pi}$. *Let* $(r_n)_{n=0}^{\infty}$ *be a sequence of discount rates such that* $r_n \to 0$. *Consider a sequence of equilibrium with discount rate* r_n, *and a sequence of equilibria of the game where the firm has committed to not censor. Then:*

[2] I define a public sequential equilibrium to be an equilibrium where the firm's strategy only depends on the public history and the firm's type at the current instant, and where beliefs update to 0 after news arrives. This is a richer class then equilibria that use time as a state variable, since they can in addition condition on the sequence of times when news arrived.

1. $\lim_{n\to\infty} r_n \left(V_n(x,\theta) + E\left(\int e^{-r_n t} k\pi_t dt | x_0 = x, \theta_0 = \theta\right)\right) \leq \liminf r_n V_{c,n}(x,\theta)$.

2. *This inequality is strict for any* $x > \min(\frac{\lambda}{\bar{\pi}} - 1, \limsup x^*_{c,n})$.

The firms willingness to pay for not censoring is always higher than the amount the firm expects to spend on censorship in the game with endogenous censorship as the firm becomes patient. As the firm becomes patient, all that matters is where beliefs end up. In the game with censorship, 0 is the unique absorbing state, beliefs always eventually converge to 0. On the other hand, in the game where the firm has committed, in any region where beliefs are drifting up, with positive probability the firm eventually starts selling a high quality product, never stops, and beliefs drift up towards 1. So the firm receives a positive payoff with positive probability forever.

When $k\bar{\pi} > \left(1 + \frac{r}{\lambda}\right) c$, there may exist equilibria where the firm invests in quality. The firm now finds investing desirable because once has a high quality product, it has either shut down bad news or no longer needs to censor bad news. There always exist equilibria of the form illustrated in the following figure.

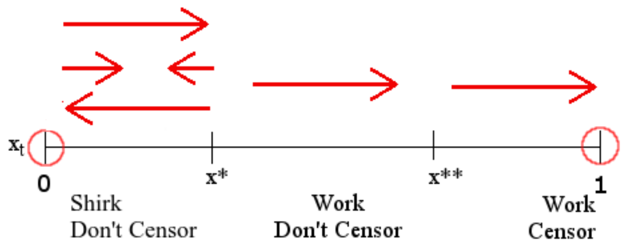

Figure 6: The structure of an equilibrium where the firm stops censoring before it stops working.

In this equilibrium, the firm works and censors when it has a high reputation, stops censoring at an intermediate reputation and stops working when its reputation is low enough. In equilibrium of this type, as long as $x^* > 0$ beliefs can converge to 0 or 1. At 1, the firm has the incentive to invest because if the firm chooses not to invest and their product becomes a low quality product, it will then have to censor news to maintain its high reputation. But, since censoring is sufficiently expensive, the firm always has the incentive to invest in quality in order to not have to censor bad news. At 0, if $x^* > 0$, the firm can never convince consumers that it is investing in quality, so it never benefits from investing in quality, and its reputation can never recover. For very low c, the firm can have the incentive to work everywhere, in which case beliefs converge to 1 with probability 1.

In this equilibrium, beliefs depend crucially on the both the initial condition and are path dependent. In the region where the firm is censoring and investing beliefs converge to 1 with probability 1. If beliefs start in the region near 0 where beliefs are drifting down, beliefs converge to 0 with probability 1. In every other region, with positive probability beliefs converge to either 1 or 0.

While the game without censorship from Board and Meyer-ter-Vehn (2013) has similar qualitative features, now a low quality firm with a high reputation can benefit from the ability to censor. The low quality firm can't be caught selling a low quality product if it's censoring, so it is better able to maintain a high reputation. On the other hand, a firm with a high quality product and a high reputation is only hurt by the ability to censor. Since the firm always has the incentive to invest in quality in the region where they have chosen to censor, a firm that starts with a high quality product never actually censors any bad news. On the other hand, the firm faces beliefs that drift up slower than they would if the firm could convince consumers that it was working and never censoring.

Equilibrium could also have the structure illustrated in the following figure.

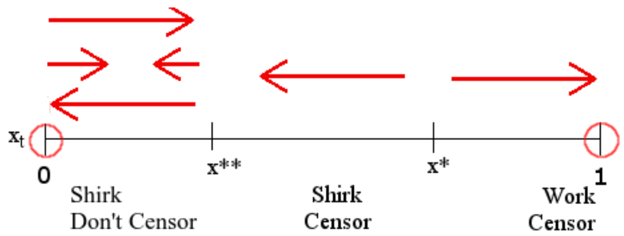

Figure 7: The structure of an equilibrium where the firm stops working before it stops censoring.

Equilibria with this structure may not always exist, but there are parameter values where equilibria of this form do exist. In this class of equilibria the limiting distribution of beliefs is entirely determined by the initial belief. Unless beliefs start in the region where the firm invests and censors, beliefs go to 0 with probability 1, and otherwise beliefs go to 1 with probability 1. As in the previous case, the ability to censor allows a low quality firm with a high reputation to preserve its reputation until it successfully develop a better product. On the other hand, it now allows a low quality firm with an intermediate reputation that never plans on investing in quality to prevent its reputation from dropping to 0 until it is already very low.

5 Discussion

At first glance, many of these results seem to be driven by the firms ability to set $a = 1$ or to fully shut down news in the censorship case. This is not entirely the case. These results for the most part carry over to the model where the highest level of effort is ϵ less than 1 or the where news can't be shut down completely, with some minor modifications. For instance, for very low k, if the maximum level of effort is ϵ less than 1, there is a small region of beliefs where the firm could be working and not promoting, but this region and any region where the firm is actually working and not promoting would vanish as $\epsilon \to 0$. Similarly, there could be a small region close to $x = 1$ where the firm is working in the bad news model, even if effort is sufficiently cheap, but this would also vanish as the distance between $\bar{\pi}$ and $\bar{\pi}$ goes to 0.

On the other hand, these properties may not extend to arbitrary public perfect equilibria that are not markovian in the public belief. As in Marinovic et al. (2015) and Halac and Prat (2014), there may be equilibria in richer state spaces that have different properties and rely on more complicated cycles of rewards and punishments through consumer's public beliefs. Equilibria that are markovian in the public belief capture features of how firm incentives for investment and signaling are driven by reputation concerns, but there certainly may be a larger class of equilibria that also have interesting features.

It would be interesting to consider an environment where the firm can both censor bad news and promote good news. I conjecture that depending on the parameter values, equilibria in this sort of environment would either resemble the good news game or the bad news game, depending on if incentives are driven more by the desire to avoid bad news or to speed up good news. A model where the firm can promote when it's selling a low quality product, either up to a lower $\bar{\pi}$ or for a higher cost may be worth investingating. In this model, the existence argument becomes significantly more complex, since what consumers believe after an arrival is now also endogenous, but I suspect, to some extent, the qualitative features of the equilibria would remain unchanged.

6 Conclusion

A firm's ability to control what information consumers see about a product is a crucial part of a firms ability to build and maintain a reputation. After a successful movie, an actor has incentives to invest in building and promoting their brand, instead of investing in future projects. Even as effort becomes arbitrarily cheap, a firm will still spend some time not working, because it knows there is no benefit from investment in quality unless it invests in promotion.

Censoring bad news can have much starker effects on incentives. While promotion and effort are compliments, I exert effort because I know I can promote my successes in the future, effort and censorship are substitutes. This leads to a cases where, when the cost of censorship is sufficiently lower, the firm never invests in its product, instead choosing to invest in hiding bad news about the product.

This is reminiscent of oil companies suppressing research about climate change or cigarette companies suppressing research about how unhealthy cigarettes are. While these companies could be investing in products that are safer and more effective, instead they've invested significant amounts of money in hiding bad news about their products. This has not only been bad for consumers, but seems to have also been bad for the companies, who now are trying to escape their negative reputations. Elaborate re-branding campaigns by Phillip Morris or British Petroleum seem to be designed entirely to escape the firms reputation for selling a dangerous, harmful product.

Controlling information is an important component of how firm a builds and maintains its reputations. The ability to promote and the ability to censor lead to interesting equilibrium dynamics and have real consequences on both firm strategies and firm payoffs. There are many other environments, including the more traditional reputation environments where players have beliefs over strategic types of the other players, where a similar analysis may be worth investigating. There are also many other ways a firm could control information, for instance, by hiding some of the history or choosing which summary statistics to display to consumers that may also be worthwhile to investigate.

References

Bar-Isaac, Heski and Steven Tadelis (2008), "Who wants a reputation." *Foundation and Trends in Microeconomics*.

Benabou, Roland and Guy Laroque (1992), "Using privileged information to manipulate markets: Insiders, gurus, and credibility." *The Quarterly Journal of Economics*, 921–958.

Board, Simon and Moritz Meyer-ter-Vehn (2013), "Reputation for quality." *Econometrica*, 81, 2381–2462.

Board, Simon and Moritz Meyer-ter-Vehn (2015), "A reputational theory of firm dynamics."

Bohren, Aislinn (2012), "Stochastic games in continuous time: Persistent actions in long-run relationships."

Bonatti, Alessandro and Johannes Hörner (2014), "Career concerns with exponential learning."

Che, Yeon-Koo and Johannes Hörner (2013), "Optimal design for social learning." Technical report.

Cripps, Martin W, George J Mailath, and Larry Samuelson (2004), "Imperfect monitoring and impermanent reputations." *Econometrica*, 72, 407–432.

Davis, Mark HA (1993), *Markov Models & Optimization*, volume 49. CRC Press.

Dellarocas, Chrysanthos (2005), "Reputation mechanism design in online trading environments with pure moral hazard." *Information Systems Research*, 16, 209–230.

Dilmé, Francesc (2014), "Reputation building through costly adjustment."

Ekmekci, Mehmet (2011), "Sustainable reputations with rating systems." *Journal of Economic Theory*, 146, 479–503.

Faingold, Eduardo and Yuliy Sannikov (2011), "Reputation in continuous-time games." *Econometrica*, 79, 773–876.

Fudenberg, Drew and Jean Tirole (1986), "A" signal-jamming" theory of predation." *The RAND Journal of Economics*, 366–376.

Halac, Marina and Andrea Prat (2014), "Managerial attention and worker engagement."

Holmström, Bengt (1999), "Managerial incentive problems: A dynamic perspective." *The Review of Economic Studies*, 66, 169–182.

Holmstrom, Bengt and Paul Milgrom (1991), "Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design." *Journal of Law, Economics, & Organization*, 24–52.

Hörner, Johannes and Nicolas Lambert (2015), "Motivational ratings."

Hu, Ju (2015), "Biased learning and permanent reputation."

Keller, Godfrey and Sven Rady (2010), "Strategic experimentation with poisson bandits." *Theoretical Economics*, 5, 275–311.

Klein, Nicolas (2015), "The importance of being honest." *Theoretical Economics.*

Klein, Nicolas and Tymofiy Mylovanov (2010), "Expert experimentation." Technical report, working paper, University of Munich and Pennsylvania State University.

Klein, Nicolas and Sven Rady (2011), "Negatively correlated bandits." *The Review of Economic Studies.*

Kreps, David M, Paul Milgrom, John Roberts, and Robert Wilson (1982), "Rational cooperation in the finitely repeated prisoners' dilemma." *Journal of Economic Theory*, 27, 245–252.

Liu, Qingmin (2011), "Information acquisition and reputation dynamics." *The Review of Economic Studies*, 78, 1400–1425.

Liu, Qingmin and Andrzej Skrzypacz (2014), "Limited records and reputation bubbles." *Journal of Economic Theory*, 151, 2–29.

Mailath, George J and Larry Samuelson (2001), "Who wants a good reputation?" *The Review of Economic Studies*, 68, 415–441.

Mailath, George J and Larry Samuelson (2006), *Repeated games and reputations*, volume 2. Oxford university press Oxford.

Marinovic, Iván, Andrzej Skrzypacz, and Felipe Varas (2015), "Dynamic certification and reputation for quality." *Available at SSRN 2697762.*

Pei, Harry Di (2015), "Reputation with strategic information disclosure."

Tadelis, Steven (1999), "What's in a name? reputation as a tradeable asset." *The American Economic Review*, 89, 548.

Appendix 1 - Good News Proofs

Proof of Lemmas 1 and 2

These proofs are very similar to the corresponding arguments in Board and Meyer-ter-Vehn (2013).

Lemma 1 Proof. These conditions follow from Lemma 5 of Board and Meyer-ter-Vehn (2013). They show that the function $\psi(t) = \int_t^\infty e^{-\int \rho(s)} \phi(t) dt$ is the unique bounded solution to

$$f(t) = \int_t^\infty \phi(s) - \rho(s) f(s) ds.$$

So the value function can be rewritten as

$$V(x, \theta) = \max_{a,\pi} \int_0^\infty x_t + a(x_t, \theta)(\lambda D(x_t) - c) + \pi(x_t, \theta) 1_{\theta=H}(\Delta(x_t) - k) + \lambda V(x_t, L) - \lambda V(x_t, \theta) - rV(x_t, \theta) dt,$$

where x_t solves $\dot{x}_t = (\tilde{a}_t - x_t)\lambda - \tilde{\pi}_t x_t (1 - x_t)$ with $x_0 = x$.

Therefore the optimal $a(x_t, \omega)$ solves

$$\max_{a \in [0,1]} \lambda D(x_t) a - ca$$

and the optimal $\pi(x_t, H)$ solves

$$\max_{\pi \in [0, \bar{\pi}]} \Delta(x_t)\pi - k\pi.$$

since these maximize the integrand pointwise. ∎

Lemma 2 Proof. Subtracting the two value functions gives

$$D(x) = \int_0^\infty \pi(x_t, H)(\Delta(x_t) - k) - (r + \lambda) D(x_t) dt.$$

Using Lemma 5 again,

$$D(x) = \int_0^\infty e^{-(r+\lambda)t} \pi(x_t, H)(\Delta(x_t) - k) dt.$$

Now I need to show a series of monotonicity properties.

Lemma 5 *If $x_0 > x_0^*$ then $x_t \geq x_t^*$ for all $t < \tau$, where τ is the first time news arrives.*

Proof. Both x_t and x_t^* are continuous in t. If $x_t = x_t^*$ at any t, then for any $s > t$, x_s and x_s^* both solve

$$\dot{x}_s = \lambda(a(x_s) - x_s) - \pi(x_s, H) x_s (1 - x_s).$$

with the same initial condition. Since I have selected the unique solution that is consistent with the discrete time approximation, x_s^* and x_s must be equal. So, if x_t and x_t^* cross, they must be the same from then on, so $x_t \geq x_t^*$. ■

This implies that the value functions are increasing in x. For any two initial conditions, $x > x'$, the firm facing a sequence of consumers with prior $x_0 = x$ could instead follow the strategy they would have followed if they faced a sequence of consumers with prior $x_0 = x'$. This would induce the same probability measure over signals, quality shocks, and quality. Moreover, by Lemma 5, the flow payoffs would be strictly higher than they were for the firm that starts at x'. Since equilibrium payoffs must be (weakly) greater than this, value functions must be increasing. Therefore, $V(x, \omega)$ is increasing in x.

The value of a signal $\Delta(x) = V(1, H) - V(x, H)$ is decreasing since $V(x, H)$ is increasing.

The high quality firm gets a higher payoff than the low quality firm since it can generate the exact same distribution over signals and quality as a firm that starts with low quality. So $V(x, H) \geq V(x, L)$.

Finally, $D(x)$ is decreasing. This difference can be written as

$$\int_0^\infty \pi(x_t, H)[\Delta(x_t) - k]dt,$$

where $\pi(x_t, H)$ maximizes

$$\pi(x_t, H)[\Delta(x_t) - k].$$

For larger x_t, $\Delta(x_t) - k$ is smaller. Therefore, if beliefs start at a larger x_0, the integrand is smaller pointwise. So $D(x)$ is decreasing. ■

Proposition 1 - Existence and Structure

The result that the equilibrium can be characterized in terms of two cutoffs follows directly from the monotonicity properties. Existence of equilibrium and $x^{**} \geq x^*$ are slightly more involved.

Lemma 6 *In any equilibrium, there exists points x^{**} and x^* such that*

$$a(x) = \begin{cases} 1 \ \textit{if}\ x < x^* \\ 0 \ \textit{if}\ x > x^* \end{cases},$$

$$\pi(x, H) = \begin{cases} \bar{\pi} \ \textit{if}\ x < x^{**} \\ 0 \ \textit{if}\ x > x^{**} \end{cases}.$$

Moreover, $x^{**} \geq x^*$.

Proof. Consider any equilibrium, let x^* be the highest public belief where the firm exerts effort[3] and let x^{**} be the highest public belief where the firm promotes.

Suppose, that $x^* > x^{**}$. Then consider any belief $x' \in (x^*, x^{**})$ where the firm exerts effort. There are a few possibilities that must be considered.

If beliefs are drifting up or are constant at this point, then they never cross x^*. So, if the firm followed this strategy there would be no arrivals of news and the firm would receive

$$\int e^{-rt}[x_t - ca_t]dt.$$

So, the firm prefers to never work, since it gets a payoff of $\int_0^\infty e^{-rt}x_t$ from never working.

Suppose at point in (x^*, x^{**}) are beliefs drifting up or constant. Then I have to consider the case where beliefs are drifting down at all points in (x^*, x^{**}). Since beliefs would be drifting up if $a(x') = 1$, and the firm is assume to be exerting effort, it must be that $a(x') \in (0, 1)$. This can only happen if $D(x) = \frac{c}{\lambda}$, for all $x' - \epsilon < x < x'$ and some $\epsilon > 0$, since the firm needs to choose a level of effort in $(0, x)$ for beliefs to drift down. This means that $D(x) = \frac{c}{\lambda}$ for all $x \in (x^*, x^{**})$. But

$$D(x') = \int_t^{t+\delta} e^{-(r+\lambda)}0\, dt + e^{-(r+\lambda)\delta}D(x)$$

for sufficiently small δ and some x, s.t.$x' - \epsilon < x < x'$. So, $D(x') < \frac{c}{\lambda}$, but the firm was assumed to be exerting effort at that point, which is a contradiction.

Therefore, $x^{**} \geq x^*$. Moreover, $D(x)$ cannot be constant below x^* because

$$\pi(x_t, H)(\Delta(x_t) - k) = \bar{\pi}(\Delta(x_t) - k),$$

so for any two beliefs x, x' below x^*, $D(x_0) > D(x_0')$ if $x_0' > x_0$ since the integrand at any point x_t' is strictly less than the integrand at x_t. ■

A MPE Exists

Theorem 1 *A Markov Perfect Equilibrium Exists.*

Proof. The proof proceeds as follows, I define an correspondence $T_\gamma : [0, 1]^2 \to 2^{[0,1]^2}$, I'll argue that this operator is UHC, non-empty and convex valued, conclude that this operator has a fixed point. Then, I argue that the limit of a sequence of these fixed points as $\gamma \to 0$ characterizes an equilibrium.

Let $\gamma > 0$. Let $U_{\gamma,H} : [0, 1]^3 \to \mathbb{R}$ as

$$U_{\gamma,\theta}(x_0, x^*, x^{**}) = \max_{a,\pi} E_{a,\pi}\left(\int e^{-rt}[x_{\gamma,t}(x^*, x^{**}) - ca_t - k\pi_t]dt|\theta_0 = \theta, x_0 = x\right)$$

[3] It follows from this proposition that it is also the lowest x where the firm shirks

where $x_{\gamma,t}$ behaves as follows. The firm is believed to be playing according to cutoffs x^* and x^{**}. So the firm works below x^* and promotes below x^{**}. Belief follow the modified the law of motion

$$\dot{x}_{\gamma,t} = \begin{cases} \min(-\gamma, \lambda(\tilde{a}_t - x_{\gamma,t}) - \tilde{\pi}_t x_{\gamma,t}(1 - x_{\gamma,t})) \text{ if } x_{\gamma,t} > x^{**} \\ \lambda(\tilde{a}_t - x_{\gamma,t}) - \tilde{\pi}_t x_{\gamma,t}(1 - x_{\gamma,t}) \text{ otherwise,} \end{cases}$$

and strategies are consistent with admissibility[4].

Let

$$\Gamma_\gamma(x^*, x^{**}) = \begin{cases} \arg\min_{\hat{x}^* \in [0,1]} \int_{\hat{x}^*}^1 \lambda(U_{\gamma,H}(s, x^*, x^{**}) - U_{\gamma,L}(s, x^*, x^{**})) - c\, ds \\ \arg\min_{\hat{x}^{**} \in [0,1]} \int_{\hat{x}^{**}}^1 U_{\gamma,H}(1, x^*, x^{**}) - U_{\gamma,H}(s, x^*, x^{**}) - k\, ds \end{cases}.$$

Since the value functions are bounded, the objective functions are continuous and well defined, so this operator is well defined and non-empty. By the monotonicity of $D_\gamma(s, x^*, x^{**}) = U_{\gamma,H}(s, x^*, x^{**}) - U_{\gamma,L}(s, x^*, x^{**})$ and $\Delta_\gamma(s, x^*, x^{**}) = U_{\gamma,H}(1, x^*, x^{**}) - U_{\gamma,H}(s, x^*, x^{**})$, this is convex valued.

Continuity in Initial Values

Lemma 7 *For given cutoffs (x^{**}, x^*) and for any $\epsilon > 0$, there exists a $\delta > 0$ such that for any $\hat{x}_0 \in [x_0 - \delta, x_0 + \delta]$ and any history then*

$$\int e^{-rt} |x_{\gamma,t} - \hat{x}_{\gamma,t}| dt \leq \epsilon.$$

Lemma 8 *$U_{\gamma,\theta}(\cdot, x^*, x^{**})$ is continuous in the first argument.*

Lemma 7 Proof. After the first arrival, beliefs coincide forever, so it is sufficient to show that $\hat{x}_t$ and x_t stay close together before an arrival. For some cases, it may be easier to consider the log likelihood ratio $l_t = \log(x_t/(1 - x_t))$. Fix $l_0 > \hat{l}_0 + \delta$, the argument is analogous for the other case.

1. l_t and $\hat{l}_t$ are both in the work/promote region.

   The reputational drift $\lambda(1 + e^{-l}) - \bar{\pi}$ is decreasing in the work/promote region, so beliefs are drifting closer together as long as they both lie in that region.

2. l_t is in the work/no promote region, $\hat{l}_t$ is in the work/promote region.

   l_t reaches the work/promote region in finite time, and is at most δ away from the edge of this region, so as by making δ small, the time these beliefs are both not in this region can be made arbitrarily small. Once the two beliefs enter this region, they must start drifting together by 1, so the most the two beliefs can drift apart is bounded. For sufficiently small δ, they won't spread apart by more than $r\epsilon$.

[4]specifically when the drift switches signs so that its positive below the cutoff and negative above the cutoff, it is 0 at the cutoff, at any cutoff where the drift doesn't change sign, the choice of strategy doesn't matter.

3. l_t is in the shirk/promote region or the shirk/no promote and $\hat{l}_t$ is in the work/promote region.

   Either $\min(x^*, x^{**})$ is convergent, in which case the drift above the cutoff is negative and the drift below the cutoff is positive, so beliefs are drifting together, or it is permeable, so the drift is negative on both sides. As before, since l_t reaches the cutoff in finite time, and then beliefs start drifting closer together again, δ can be chosen so that $|l_t - \hat{l}_t|$ separate by at most $r\epsilon$.

4. l_t and $\hat{l}_t$ are in the work/no promote region.

   x_t and $\hat{x}_t$ are both drifting at the same rate until they leave the region and enter the case from section 2, so they don't move any further apart until they enter the region in 2.

5. l_t is in the shirk/no promote region, $\hat{l}_t$ is in the work/no promote region. l_t is drifting down, so it reaches the work/no promote region in finite time, so δ can be chosen to make l_t and $\hat{l}_t$ stay arbitrarily close together.

6. l_t and $\hat{l}_t$ are in the shirk/promote region.

   In this region, x_t and $\hat{x}_t$ are determined by a standard differential equation of the form $x'(t) = f(x,t)$ and $x_0 = x$, $f(x,t)$ is a continuous function and is lipschitz in x, so by standard results from the differential equation literature, it is continuous in the initial conditions until $\hat{x}$ leaves the shirk/promote region. If the shirk/promote cutoff is interior, δ can be chosen so that beliefs are at most δ' apart when they hit that cutoff, in which case an argument from a previous case applies. Otherwise δ can be chosen so that beliefs are at most δ' apart when $\hat{x}_t$ hits 1/2, after which, the x_t and $\hat{x}_t$ only drift apart a small amount more before x_t hits 1/2 and they start drifting together since $\dot{x}_t - \dot{\hat{x}}_t = \lambda(\hat{x}_t - x_t) + \bar{\pi}(\hat{x}_t(1-\hat{x}_t) - x_t(1-x_t)) < 0$ below 1/2.

7. l_t is in the shirk/no promote region, $\hat{l}_t$ or the shirk/promote region. l_t must enter the shirk/promote region in finite time, after which we are in one of the previous cases.

8. l_t and $\hat{l}_t$ are in the shirk/no promote region.

   In this region $\dot{x}_t - \dot{\hat{x}}_t = \lambda(\hat{x}_t - x_t) < 0$, so beliefs are drifting close together until they enter a region from one of the other cases.

Finally, bounding the distance between the likelihood ratios is sufficient because, by the mean value theorem

$$|x_t - \hat{x}_t| \leq \max_x \frac{1}{(\frac{d}{dx}\log x/(1-x))}|l_t - \hat{l}_t| \leq |l_t - \hat{l}_t|.$$

■

Lemma 8 Proof. This follows directly from the previous claim. The optimal strategy is independent of the initial conditions, and x_t is continuous in the initial condition, so $U_{\gamma,\theta}(\cdot, x^*, x^{**})$ must be continuous. ■

Continuity in Cutoffs

Lemma 9 *$U_\theta^\gamma(x_0, x^*, x^{**})$ is continuous in x^* and x^{**}.*

Proof. It is sufficient to show that for cutoffs (x^*, x^{**}) there is a $\delta > 0$ s.t. $(\hat{x}^*, \hat{x}^{**}) \in [x^* - \delta, x^* + \delta] \times [x^{**} - \delta, x^* + \delta]$ such that

$$E_{a,\pi}\left(\int e^{-rt}|x_{\gamma,t} - \hat{x}_{\gamma,t}|dt\right) \leq \epsilon.$$

since

$$\begin{aligned} U_{\gamma,\theta}(x_0, x^*, x^{**}) &\geq E_{\hat{a},\hat{\pi}}\left(\int e^{-rt}[x_{\gamma,t} - ca_t - k\pi_t]dt|\theta_0 = \theta, x_0 = x\right) \\ &\geq E_{a,\pi}\left(\int e^{-rt}[\hat{x}_{\gamma,t} - c\hat{a}_t - k\hat{\pi}_t]dt|\theta_0 = \theta, x_0 = x\right) - \epsilon \geq U_{\gamma,\theta}(x_0, \hat{x}^*, \hat{x}^{**}) \end{aligned}$$

where $\hat{a}, \hat{\pi}$ is the optimal strategy if the cutoffs are $\hat{x}^*$ and $\hat{x}^{**}$. Since the problem is symmetric, this implies that $U_{\gamma,\theta}$ is continuous.

First I'll show continuity in x^*. Fix x^{**} and a strategy profile (a, π). I need to show that for any $\epsilon > 0$ there's a δ s.t. for $\hat{x}^* \in [x^* - \delta, x^* + \delta]$ then

$$\int e^{-rt}|x_{\gamma,t} - \hat{x}_{\gamma,t}|dt \leq \epsilon.$$

I call a cutoff X convergent if $\lim_{x\to X+} g(x) \geq 0$, $\lim_{x\to X-} g(x) \leq 0$, divergent if $\lim_{x\to X+} g(x) < 0$ and $\lim_{x\to X-} g(x) > 0$ and permeable if it is neither convergent or divergent. At a convergent cutoffs, beliefs that start in a neighborhood of the cutoff converge to it, and stay there until an arrival, at a divergent cutoff beliefs are pushed away from it, and at a permeable cutoff beliefs cross the cutoff and continue to drift in the same direction.

Let $x^g = \min(\lambda/\bar{\pi}, 1, x^{**})$. x^* is convergent if it is below x^g. First suppose that x^* is convergent. Let $\hat{x}^* \in [x^*, \min(x^* + r\epsilon, 1)]$. Then either $\hat{x}^*$ is convergent, or the convergent point for the $\hat{x}$ process is between $\hat{x^*}$ and x^*. The largest possible distance between the two processes is the distance between these two cutoffs, since beliefs coincide until they hit the cutoffs, and then get stuck at the two convergent points until an arrival. So $|\hat{x}_{\gamma,t} - x_{\gamma,t}| \leq r\epsilon$.

Now suppose x^* is divergent. Then beliefs drift down to a convergent point x^g. Consider a $\hat{x}^* \in [x^*, \min(x^* + \delta', 1)]$. $\hat{x}^*$ is also divergent. The two trajectories coincide until x hits $\hat{x}^*$. Then the $\hat{x}$ process drifts down slower. It takes at most $\frac{\delta'}{\gamma}$ for it to reach x^*, while it takes the x process at least $\frac{\delta'}{\min(\lambda,\gamma)}$. So $\hat{x}^*$ can be chosen so that the distance between these two processes can be made arbitrarily small until they've dropped below x^*. Then they decrease at the same speed (γ), until x^{**}. So, $\hat{x}^*$ can be chosen so that $|x_t - \hat{x}_t| < r\epsilon$ for all times t where either process is above x^{**}. Once the $\hat{x}$ process is at x^{**}, they both behave like the process from Lemma 7 until the next arrival, with two different initial conditions, $\hat{x}^*$ can be chosen so that $|x_t - \hat{x}_t| < r\epsilon$ for all x_t until an arrival, after which the two processes both coincide until they hit x^{**} again.

This argument can be reversed to show continuity in the other direction. The proof for continuity in the x^{**} cutoff is analogous. ■

So, by the dominated convergence theorem, the objective function is continuous. Define $T_{\gamma(x^*,x^{**})} = \Gamma_{\gamma(x^*-x^{**})}$. This is UHC, convex valued (by monotonicity of $\Delta(\cdot)$ and $D(\cdot)$), compact valued and non-empty. So it has a fixed point.

Consider a sequence of fixed points (x_n^{**}, x_n^*), where each point is a fixed point of $T_{\frac{1}{n}}$. This sequence has a convergent subsequence, since it is a sequence in a sequentially compact space $[0,1]^2$. Let (x^{**}, x^*) be the limit of this sequence.

If $x^{**} > x^*$, then an equilibrium exists. There exists a further subsequence of points where $x_{n_k}^{**} > x_{n_k}^*$ and the drift is the true drift, since when this holds, beliefs follow law of motion that is at most $-\lambda x$ above $x_{n_k}^{**}$, so there exist a sequence such that $|x_{n_k}^{**} - x^{**}| < \epsilon$, $x_{n_k}^{**} > 0$ and a n_0 such that $\frac{1}{n_0} < \lambda(x^{**} + \epsilon)$ so that for all points in the sequence the drift follows the actual law of motion. Any of these points are an equilibrium.

Similarly, if there is any point in the original sequence where $x_n^{**} \geq x_n^*$ for sufficiently large n, that is an equilibrium[5]. So, in any case that remains there must be a further subsequence where the work cutoff is larger than the promote cutoff at each point.

So, it remains to analyze the case where neither of these things hold. . For the remainder of this section, to avoid a lot of messy subscripts, (x_n^{**}, x_n^*) is a subsequence of the original sequence such that converges, and $x_n^* > x_n^{**}$ for all points in the sequence.

Case 1: $x^{**} < x^*$. If this is the case, there exists a large enough N such that for all $n \geq N$, $|x_n^{**} - x_n^*| \geq \frac{x^*-x^{**}}{2}$. Then

$$D^{\frac{1}{n}}(x_n^*) = \int_{T_n^{**}}^{\infty} e^{-(r+\lambda)t}\bar{\pi}[\Delta(x_s) - k]ds.$$

where T_n^{**} is the time it takes to get from x^* to x^{**}. This is bounded below by $\frac{1}{\frac{1}{4}\gamma}\frac{x^*-x^{**}}{2}$, which goes to ∞ as $\gamma \to 0$. Therefore, for sufficiently large n, $D^{\frac{1}{n}}(x_n^*) < c$, so the firm is promoting when it is unprofitable or $x^* = 0$, which is a contradiction.

Case 2: $x^{**} = x^*$.

This is also impossible, unless $x^{**} = x^* = 0$.As long as beliefs and the value functions are continuous in γ the limiting equilibrium effectively behaves like an equilibrium in a game where the drift at $x^{**} = x^*$ is 0 as beliefs approach x^* from above (since as γ goes to 0, the time it takes to get from x^* to x^{**} goes to ∞.

[5]The case where there are infinitely many such points, but they are all $(0,0)$ requires a slightly different, simpler argument, following directly everything being continuous in γ

Continuity in in γ

Lemma 10 *For any fixed γ, for any $\epsilon > 0$, there exists a δ s.t. for any $\hat{\gamma} \in [\gamma - \delta, \gamma + \delta]$ such that if x is the belief trajectory in the game with maximum drift $-\gamma$ and $\hat{x}$ is the corresponding belief for the game with maximum drift $-\gamma'$, the function*

$$E_{a,\pi}\left(\int_0^\infty e^{-rt}|x_t - \hat{x}_t|dt\right) \leq \epsilon.$$

Moreover, the time it takes for the two processes to hit both cutoffs is also continuous in γ.

Proof.

Fix an initial value x_0, cutoffs x^* and x^{**}, and strategy profile (a, π). Consider $\gamma, \hat{\gamma} > 0$. For any t, $x_t = \hat{x}_t$ until the first t where x is less than x^* and greater than x^{**}, and beliefs coincide until this point after any arrival. So, it is sufficient to bound the integral if $x_0 = x^*$ in the event where no signals arrive. Assume $\gamma > \hat{\gamma}$, the proof for the other direction is similar.

Let T and $\hat{T}$ be the time it takes to get to x^{**} from x^* for the two processes. Let $K = |x^{**} - x^*|$

$$\int_0^T e^{-rt}(x_t - \hat{x}_t)dt = \int_0^{\frac{1}{\gamma}} e^{-rt}(\hat{\gamma}t - \gamma t)(x^* - x^{**})dt \leq \delta K_2.$$

The second interval, from T to $\hat{T}$ is bounded by $\frac{1}{\hat{\gamma}} - \frac{1}{\gamma} \leq (\frac{1}{\gamma - \delta} - \frac{1}{\gamma})$ and since $\frac{1}{x}$ is continuous, there exists δ that make this arbitrarily small. Moreover, the total distance between x_{T^*} and $\hat{x}_T^*$ can be made arbitrarily small by choosing a small δ.

Finally, after time T^*, beliefs follow the same law of motion from then on, with both drift and promotion, but have different initial values. These are now two processes with the same laws of motions and different initial conditions until an arrival, so by lemma 7, there exists a δ' s.t. if $|x_{T^*} - \hat{x}_{T^*}| < \delta'$, then $\int_{T^*}^\infty e^{-r(t-T^*)}|x_t - \hat{x}_t|dt < \frac{\epsilon}{3}$.

Therefore, there exists a δ s.t. $\int e^{-rt}|x_t - \hat{x}_t|dt \leq \epsilon$.

∎

So, the value function $U_{\gamma,\theta}(x_0, x^*, x^{**}) = E(\int e^{-rt}[x_t^\gamma(x_0, x^*, x^{**}) - ca_t - k\pi_t]dt)$ is continuous in γ, x^* and x^{**}. So it converges given the sequence $(\gamma, x^*, x^{**})_n = (\frac{1}{n}, x_n^*, x_n^{**})$. Since x_n^* and x_n^{**} are always sequentially rational, they also are in the limit.

The drift converges to 0 when x^* was approached from above and consider $D(x_n^*)$. Moreover, because $x_n^* > x_n^{**}$ at all points in the sequence,

$$\lim D_n(x_n^*, x_n^{**}) \to \int_0^\infty e^{-(r+\lambda)}\pi(x^*, x^{**}, +)[\Delta(x_t, +) - k]\, dt \leq 0$$

, where + denotes these values as x_t approaches them from the right. So

$$\frac{c}{\lambda} = D(x^*) = \int e^{-rt}[\Delta(x^*+) - k]\,dt \leq 0.$$

which is a contradiction, since this implies that for sufficiently small n, $D(x_n^*) < \frac{c}{\lambda}$, so x_n^* was not sequentially rational in the perturbed game. Therefore, the drift at x^* drifts in the direction it should, so there must exist an n where $x^* < x^{**}$, which is an equilibrium of the true game. ■

Exogenous Promotion

Proof of Proposition 2: Let $x^* > 0$ be the work/shirk cutoff in an equilibrium of the game where the firm hasn't committed to working everywhere (in the case where this is equal to 0, the argument is trivial). Consider the corresponding commitment game (where all payoffs are considered without that fixed cost the firm pays to commit), with cutoff x^*. Let $D_{c,x^*}(x^*)$ be the corresponding difference between high and low quality payoffs at that point in the commitment game where the firm has cutoff x^*. By Board and Meyer-ter-Vehn (2013), the function $D_{c,x}(x)$ is continuous in x and $D_1^c(1) = 0$, so it suffices to show that $D_{c,x^*}(x^*) > D(x^*) = \frac{c}{\lambda}$.

Suppose not. Then $D_{c,x^*}(x^*) \leq D(x^*)$. This implies that the firm is prefers (weakly) not working to working at all beliefs $x \geq x^*$ in the commitment game.

In the non-commitment game and the commitment game, beliefs follow the same law of motion between x^{**} and x^*. Moreover, payoffs in the non-commitment game can be evaluated as if when news arrives, beliefs jump to x^* instead of 1, and the firm doesn't pay cost k. In other words

$$V(x_0, H) = \int e^{-\int_0^t (r+\lambda+\pi_s)ds}[x_t + \pi_t(V(x^*, H)) + \lambda(V(x_t, L))]dt,$$

and the firm can be treated as never working at any point between in $[x^*, 1]$ for both the comment firm and the non-commitment firm. This means that at every instant of time, the firm in the modified game where beliefs jump to x^* instead of 1 is receiving weakly lower flow payoffs than the firm who has committed to promoting everywhere, and the low quality firm is receiving exactly the same payoffs in either case, since it is never working and news never arrives. In other words

$$V_c(x^*, L) = V(x^*, L)$$
$$V_c(x^*, H) > V(x^*, H)$$

(the strict inequality comes from $x^{**} < 1$), which is a contradiction. So $D_{c,x^*}(x^*) > D(x^*)$ which means there is an equilibrium in the commitment game where the firm works at a higher belief by the intermediate value theorem.

Propositions 3 - Limit Behavior

Proposition 3 *As $c \to 0$, there is a unique limit. In the limit as $c \to 0$, equilibrium payoffs are increasing in k as long as $\lambda \geq \bar{\pi}$.*

Lemma 11 *For any sequence $c_n \to 0$ and for any sequence of equilibria, $|x_n^* - x_n^{**}| \to 0$.*

Suppose not. Then there exists an $\epsilon > 0$ and a subsequence s.t. $|x_{n_k}^* - x_{n_k}^{**}| > \epsilon$ for all k. I'm going to drop the k subscript for convenience. Consider $\Delta_n(x_n^*)$. This can be rewritten as

$$\begin{aligned}
\Delta_n(x_n^*) - k &= V_n(1, H) - k - V_n(x_n^*, H); \\
&= V_n(x_n^{**}, H) - V_n(x_n^*, H); \\
&> \int_{T^{**}}^{\infty} e^{-(r+\bar{\pi}+\lambda)t}[(x_{n,t} - x_n^*) + \lambda(V_n(x_{n,t}, L) - V_n(x_n^*, L))]dt; \\
&> \int_{T^{**}}^{\infty} e^{-(r+\bar{\pi}+\lambda)t}[(x_{n,t} - x_n^*)]dt.
\end{aligned}$$

The first inequality comes from considering the deviation where after beliefs hit x^{**}, the firm plays the strategy it would have played at x^* and the second comes from the monotonicity of the value function.

From the firm's perspective, the worst possible law of motion for $x_{n,t}$ is $-\lambda - \frac{1}{4}\bar{\pi}$ until $x_{n,t} = x^*$ and then $\dot{x}_{n,t} = 0$ after that. This gives

$$\Delta(x^*) - k > \int_0^{\frac{\epsilon}{\lambda + \frac{1}{4}\bar{\pi}}} e^{-(r+\bar{\pi}+\lambda)}(\epsilon - (\lambda + \frac{1}{4}\bar{\pi})u)du = B.$$

But, as $n \to \infty$, $c \to 0$. Moreover,

$$D(x_n^*) > \int e^{-(r+\lambda)t}[B]dt = \frac{1}{r+\lambda}B$$

But, $\lambda D(x_n^*) \leq c_n$ for all n by sequential rationality and $c_n \to 0$. This is a contradiction. So $|x_n^* - x_n^{**}|$ must converge.

Lemma 12 *The cutoff x_n^{**} converges to a unique cutoff. Cutoff determined by solution to:*

$$rk = \frac{r}{r+\lambda} + \frac{\lambda}{r+\lambda}e^{-(r+\lambda)T} - e^{-\lambda T},$$

where $x^{**} = e^{-\lambda T}$.

Suppose that the sequence of x_n^{**}'s did not converge to the cutoff x^{**}. Let T_n^* be the amount of time it takes for beliefs to go from 1 to x_n^*. Then there exists an $\epsilon > 0$ s.t.

$$|rk - \left(\frac{r}{r+\lambda}(1 - e^{-(r+\lambda)T_n^*}) + e^{-(r+\lambda)T_n^*} - e^{-\lambda T_n^*}\right)| \geq \epsilon$$

infinitely often. The value function can be rewritten as

$$
\begin{aligned}
V_n(1,H) = & \int_0^{T_n^{**}} e^{-(r+\lambda)t} x_{n,t} + \lambda V_n(x_{n,t},L)dt; \\
& + \int_{T_n^{**}}^{T_n^{*}} e^{-(r+\lambda)t-\bar{\pi}(t-T_n^{**})}[x_{n,t} + \bar{\pi}(V_n(1,H)-k) + \lambda V_n(x_{n,t},L)]dt \\
& + e^{-(r+\bar{\pi}+\lambda)T_n^{*}} V_n(x^*,H); \\
& = \int_0^{T_n^{**}} e^{-rt} x_{n,t}dt + \int_{T_n^{**}}^{\infty} e^{-ru} x_{n,u}(1-e^{-\lambda T_n^{*}}) \\
& + \int_{T_n^{**}}^{T_n^{*}} e^{-(r+\bar{\pi}+\lambda)t}[x_{n,t} + \bar{\pi}(V_n(1,H)-k) + \lambda V_n(x_{n,t},L)]dt \\
& + (e^{-(r+\lambda)T_n^{*}} - 1)[x_n^* + c].
\end{aligned}
$$

And using the indifference condition at the cutoff and the convergence of the two cutoffs,

$$
\begin{aligned}
V_n(1,H) - V_n(x^{**},H) = & \int_0^{T_n^{**}} e^{-rt} x_{n,t}dt + \int_{T_n^{**}}^{\infty} e^{-ru} x_{n,u}(1-e^{-\lambda T_n^{*}}) \\
& + \int_{T_n^{**}}^{T_n^{*}} e^{-(r+\bar{\pi}+\lambda)t}[x_{n,t} + \bar{\pi}(V_n(1,H)-k) + \lambda V_n(x_{n,t},L)]dt \\
& + (e^{-(r+\lambda)T_n^{*}} - 1)[x_n^* + c] - V_n(x^{**},H) \\
\Delta(x^{**}) = & \int_0^{T_n^{**}} e^{-rt} x_{n,t}dt + \int_{T_n^{**}}^{\infty} e^{-ru} x_{n,u}(1-e^{-\lambda T_n^{*}}) \\
& + \int_{T_n^{**}}^{T_n^{*}} e^{-(r+\bar{\pi}+\lambda)t}[x_{n,t} + \bar{\pi}(V_n(1,H)-k) + \lambda V_n(x_{n,t},L)]dt \\
& + (e^{-(r+\lambda)T_n^{*}} - 1)[x_n^* + c] \\
& - e^{(r+\bar{\pi}+\lambda)T_n^{**}}(\int_{T_n^{**}}^{T_n^{*}} e^{-(r+\bar{\pi}+\lambda)t}[x_{n,t} + \bar{\pi}(V_n(1,H)-k)] + \lambda V_n(x_{n,t},L)dt \\
& + e^{-(r+\lambda-\bar{\pi})(T_n^{*}-T_n^{**})} V_n(x^*,H)) \\
rk = & \int_0^{T_n^{*}} re^{-(r)t} x_{n,t}dt + (1-e^{-\lambda T_n^{*}}) \int_{T_n^{*}}^{\infty} re^{-rt} x_{n,t}dt + e^{-(r+\lambda)T_n^{*}} x_n^* - x_n^* + \delta(T_n^*, T_n^{**}) \\
rk = & \frac{r}{r+\lambda}(1-e^{-(r+\lambda)T_n^{*}}) + (1-e^{-\lambda T_n^{*}})e^{-(r+\lambda)T_n^{*}} + e^{-(r+\lambda)T_n^{*}} x_n^* - x_n^* + \delta(T_n^*, T_n^{**}) \\
rk = & \frac{r}{r+\lambda}(1-e^{-(r+\lambda)T_n^{*}}) + e^{-(r+\lambda)T_n^{*}} - e^{-\lambda T_n^{*}} + \delta(T_n^*, T_n^{**})
\end{aligned}
$$

Where the second equation comes from expanding $V_n(x_t, L)$ out, and using integration by parts, and the integrals are finally evaluated by using the fact that before time T^*, beliefs are equal to $e^{-\lambda t}$.The term $\delta(T_n^*, T_n^{**})$ collects all the terms that are integrals from T_n^{**} to T_n^*, and goes to 0 as $T_n^* - T_n^{**} \to 0$. Since there exists an N such that for all $n \geq N$ $\delta(T_n^*, T_n^{**} < \epsilon$, this is a contradiction.

It remains to consider what happens if x_n^{**} is 0 infinitely often. This can only happen if

$$k > \Delta_n(0) = \int_0^\infty e^{-(r+\lambda)t} dt.$$

Suppose the sequence (x_n^{**}) had another subsequence that was never 0. Then, by the logic from above, in the limit of this subsequence

$$rk = \frac{r}{r+\lambda}(1 - e^{-(r+\lambda)T^*}) + e^{-(r+\lambda)T^*} - e^{-\lambda T^*}.$$

But, the left hand side of this is bounded above by $\frac{r}{r+\lambda}$, and rk is bounded below by $\frac{r}{r+\lambda}$, so this cannot hold. Therefore, the sequence must not be non-zero infinitely often. So $x_n^{**} \to 0$, and the proposition is satisfied.

The properties of the limit equilibrium then follow directly.

Commitment in the limit

Proof of Proposition 4:

As $c \to 0$, $V_{BM}(1, H) = V_{BM}(1, L) = \frac{1}{r} > V(1, H)$. $x_{BM,t} \geq x_t$ until the first arrival (since $\lambda \geq \bar{\pi}$, and the firm that has committed is believed to be exerting effort everywhere), and after the first arrival $x_{BM,t} = 1$ forever, while x_t still drifts down. Therefore, even with costs added back in, the committed firm receives higher flow payoffs everywhere, and strictly higher payoffs after the first arrival, so

$$V_{BM}(x, \theta) > V(x, \theta) + E\left(\int_0^\infty e^{-rt} k\pi_t dt | x_0 = x, \theta_0 = \theta\right).$$

Commitment to a Strategy

Proof of Proposition 5:

Lemma 13 *Let $U_c(x, \theta)$ be the value function for the commitment game. If $U_c(1, H) > U(1, H)$, then $x^{**} > x_c^{**}$.*

Proof. Note, that in this new game, the sequential rationality condition for the investment decision still holds.

Let x^* and x^{**} be cutoffs of the equilibrium that induces the highest cutoff for investing in quality. Suppose that the firm deviates to a promotion cutoff above x^{**}. Let U_{dev} denote the value function for this deviation. Then $U_{dev}(1, H) < U(1, H)$ as is $U_{dev}(x^*, H) < U(x^*, H)$, because the firm is playing suboptimally.

If beliefs follow the law of motion induced by this deviating cutoff, payoffs decrease further, because the measure over arrivals is the same, but beliefs drift down faster.

So, in order for committing to this to be optimal, it must induce a higher level of investment in quality.

If it did not, then deviating to follow the strategy in the new game would be a profitable deviation in the original game, since absent arrival beliefs are weakly higher at every point, and the firm can always deviate to induce the same measure over arrivals. This means there must be some $x_c^* > x^*$ such that

$$D_c(x_c^*) = \int e^{-(r+\lambda)t}\bar{\pi}[U_c(1,H) - \frac{x_c^* - c + \pi[U_c(1,H) - k]}{r+\bar{\pi}}] > c/\lambda$$

This did not hold before beliefs updated, since in equilibrium at x_c^* $D(x_c^*) < c/\lambda$ and the firm's devaition decreases $U_{dev}(x_c^*, H)$ and doesn't change $U_{dev}(x_c^*, L)$.

The $x_c^*/(r+\bar{\pi})$ term was $\int_0^\infty e^{-(r+\bar{\pi})t}x_t < x_c^*/(r+\bar{\pi})$. The term $U_c(1,H) < U_{dev}(1,H)$ but by a smaller factor, $e^{-rT}\int_0^\infty e^{-(r+\pi)t}x_t dt$. Therefore, $D_c(x_c^*) < D_{dev}(x_c^*)$, the $D_{dev}(x)$ the firm receives from working at x_c^* if beliefs followed the law of motion that didn't take into account the new cutoff. But, as argued before,

$$D_{dev}(x^*) = \int e^{-(r+\lambda)t}[U_{dev}(1,H) - \frac{x^* - c + \pi[U_{dev}(1,H) - k]}{r+\bar{\pi}}] < D(x^*) \leq c/lambda,$$

since the only thing that's changed is $U_p(1,H)$, which has decreased, so it is not optimal to exert effort at x^*. This, it can't be optimal to invest at x_c^*, but $D_{dev}(x)$ is decreasing, and it was optimal to invest at a higher belief x_c^*, so this is a contradiction. ∎

Lemma 14 *For any initial condition, $x_0 = x \in [0,1]$, $\theta_0 = H$ or $x_0 = x$ and $\theta_0 = L$ where x is a point where the firm was investing before and after commitment, in the commitment game $U_c(1,H) \geq U(1,H)$.*

Proof. If x_0 is such that before and after the firm commits the firm is investing at x_0, then the only way the firm's payoffs can increase is if $U_c(1,H) > U(1,H)$. So the initial conditions that matter are initial conditions where the firm is not investing in the commitment game. Suppose that $U(1,H) > U_c(1,H)$, then the only way payoffs can go up at x_0 is if either $x^{**} > x_c^{**}$ or $x_c^* > x^*$. If neither of these held, then the firm in equilibrium would be able to mimic the strategy played with commitment, and get a strictly higher payoff than the commitment payoff, which must be higher than its equilibrium payoff.

If $x_c^* > x^*$[6] then

$$\begin{aligned}
c/\lambda &= \int_0^\infty \bar{\pi} e^{-(r+\lambda)t}[U_c(1,H) - U_c(x_c^*, H)]dt \\
&= \int_0^\infty \bar{\pi} e^{-(r+\lambda)t}[\frac{(r+\lambda)}{r+\lambda+\bar{\pi}} U_c(1,H) - \frac{1}{r+\lambda+\bar{\pi}}[x_c^*(1+\lambda/r) - k\bar{\pi}]] \\
&< \int_0^\infty \bar{\pi} e^{-(r+\lambda)t}[\frac{(r+\lambda)}{r+\lambda+\bar{\pi}} U(1,H) - \frac{1}{r+\lambda+\bar{\pi}}[x^*(1+\lambda/r) - k\bar{\pi}]] \\
&= \int_0^\infty \bar{\pi} e^{-(r+\lambda)t}[U(1,H) - U(x^*, H)]dt \\
&= c/\lambda,
\end{aligned}$$

which is a contradiction.

If $x^{**} > x_c^{**}$ and $x^* \geq x_c^*$ then the firm only benefits from the new cutoff if $x_0 > x_c^{**}$. If $\theta_0 = H$, then $U_c(1,H)$ can be written as $\int_0^t e^{-\int_0^s (r+\lambda)d\xi}[x_{c,s} + \lambda U_c(x_s, L)]ds + e^{-\int_0^t (r+\lambda)ds} U(x,H)$. So $U_c(1,H)$ is lower if $U_c(x_s, L)$ is lower for some x_s. At x_c^{**}, $U(x_c^{**}, \theta) > U_c(x_c^{**}, \theta)$ since the investment cutoff is lower and $U_c(1,H)$ is lower. Moreover, the difference between $U(x^* *_c, H)$ and $U(x_c^{**}, L)$ is greater than the difference between $U_c(x^* *_c, H)$ and $U_c(x_c^{**}, L)$, because $U(1,H)$ is larger than $U_c(1,H)$. This also holds between x_c^{**} and x^{**}, because one firm is promoting and the other isn't, and the point with this larger difference is reached faster for the firm without commitment, so $U_c(x^{**}, L) > U(x^{**}, L)$, which must also hold at all points above x^{**}, which is a contradiction. ■

Lemma 15 *If $U_c(1,H) > U(1,H)$, then $x_c^* \geq x^*$*

Proof. Suppose not, then[7]

$$\begin{aligned}
c/\lambda &= \int_0^\infty \bar{\pi} e^{-(r+\lambda)t}[U_c(1,H) - U_c(x_c^*, H)]dt \\
&= \int_0^\infty \bar{\pi} e^{-(r+\lambda)t}[\frac{(r+\lambda)}{r+\lambda+\bar{\pi}} U_c(1,H) - \frac{1}{r+\lambda+\bar{\pi}}[x_c^*(1+\lambda/r) - k\bar{\pi}]] \\
&> \int_0^\infty \bar{\pi} e^{-(r+\lambda)t}[\frac{(r+\lambda)}{r+\lambda+\bar{\pi}} U(1,H) - \frac{1}{r+\lambda+\bar{\pi}}[x^*(1+\lambda/r) - k\bar{\pi}]] \\
&= \int_0^\infty \bar{\pi} e^{-(r+\lambda)t}[U(1,H) - U(x^*, H)]dt \\
&= c/\lambda,
\end{aligned}$$

■

[6] if $x^* = 0$, then $U(1,H) > U_c(1,H)$ implies that this condition cannot hold, so I only need to consider the other case.

[7] If $x^* = 0$ (which $x_c^* = 0$ implies), then this holds vacuously, so I am considering the other case.

Properties of Equilibrium

Proof of Lemma 4 and Proposition 6:

As in the good news case, the value functions can be rewritten as

$$V(x,\theta) = \int_0^\infty e^{-\int_0^t (r+\lambda+\bar{\pi}-\pi_s)ds}[x_t - (\bar{\pi}-\pi_t)1_{\theta=L}(\Delta(x_t)) + \lambda a_t D(x_t) + \lambda(V(x_t,L) - V(x_t,\theta)) - k\pi_t - ca_t]ds,$$

which directly gives the sequential rationality conditions. The monotonicity conditions follow from the same logic as in the good news case. Moreover,

$$D(x) = \int_0^\infty e^{-(r+\lambda)t}(k\pi_t + (\bar{\pi} - \pi_t)\Delta(x_t))dt.$$

By the sequential rationality condition, this is bounded above by $\frac{k\bar{\pi}}{r+\lambda}$, so if $c \geq \frac{\lambda}{r+\lambda}k\bar{\pi}$, then the firm never works.

Existence

Proof of Proposition 6:

Case 1: $c > \frac{\lambda}{r+\lambda}k\bar{\pi}$.

Now, it can never be optimal for the firm to exert effort. It is thus sufficient to find the cutoff where the firm prefers to stop censoring, because for any beliefs, the optimal choice of this cutoff will automatically preclude any effort exerted by the firm. Consider the modified problem where x_t follows law of motion $\dot{x}_t = -\lambda x_t + \bar{\pi} x_t(1-x_t)$ if this is negative and $\dot{x}_t = 0$ otherwise and the firm never censors. Denote the value function for this problem as V_{nc}. Then the solution to

$$\max_{x^{**}} \int_0^{x^{**}} (V_{nc}(x,L) - k)dx.$$

exists, and $V_{nc}(x,L) \geq V(x,L)$ for all beliefs below x^{**}, and $V_{nc}(x^{**},L) = V(x^{**})$. $V_{nc}(\cdot,L)$ is continuous, so $V(x^{**},L) \geq k$ above x^{**} (or $x^{**} = 1$ and since $V(x,L)$ is increasing in x, this holds for all $x > x^{**}$ so this is sequentially rational.

Case 2: $c \leq \frac{\lambda}{r+\lambda}k\bar{\pi}$.

I will show the existence of an equilibrium where the firm stops censoring at a higher belief than when it stops working. Consider the following problem.

$$\max_{x^{**}} \int_{x^{**}}^1 \bar{\pi}(V_{wc}(x,L) - k)dx.$$

where V_{wc} is the value function corresponding to the problem where the firm is working and censoring everywhere except at 0 where it is doing nothing and is believed to be doing that. This is well defined, strictly increasing, and continuous except at 0. This maximization problem has a solution, x^{**}.

Define $D_{x^*}(x^*+) = \lim_{\epsilon\to 0} D^{x^*}(x^*+\epsilon)$ and $D_{x^*}(x^*-) = \lim_{\epsilon\to 0} D^{x^*}(x^*-\epsilon)$. As shown in Board and Meyer-ter-Vehn (2013), function $D_{x^*}(x^*\pm)$ satisfies

1. $D_{x^*}(x^*+)$ is continuous and strictly increasing on $[0, x^{**})$ and $\lim_{x^*\to x^{**}} D_{x^*}(x^*+) = \lim_{x^*\to x^{**}} D_0(x^*+)$.
2. $\lim_{x^*\to x^{**}} D_{x^*}(x^*-) = \lim_{x^*\to x^{**}} D_{x^{**}}(x^{**}-)$, and if $\lambda \geq \bar{\pi}$ then $D_{x^*}(x^*-)$ is continuous and strictly increasing with $\lim_{x^*\to 0} D^{x^*}(x^*-) = 0$.
3. At any cutoff x^*, if beliefs are drifting up above x^* and down below x^*, then $D_{x^*}(x^*+) > D_{x^*}(x^*-)$.

Property 1 follows from the following observation. When the firm is believed to be working and not censoring, beliefs are drifting up, and never drop below x^* unless news arrives and they go to 0. So, $D_{x^*}(x)$ is the same as $D_0(x)$, the difference between the value of a high and low quality product of a firm who works everywhere but at $x = 0$, for all $x^{**} > x > x^*$. So $D_{x^*}(x^*+) = D_0(x^*+)$, and $D_0(x)$ is strictly increasing and continuous by the previous lemma, and $D_{x^*}(x^{**}+) = D_0(x^{**}+)$.

Property 2 follows from the same logic as property 1. For any cutoff $x^{**} > x^* \geq \max(1 - \lambda/\bar{\pi}, 0)$, beliefs always stay below x^* once they start below x^*, so the firm's payoffs are the same as the payoffs they'd receive if they started shirking at x^{**}, and since $D_{x^{**}}(x)$ is continuous and increasing below x^{**} and above $\max(1 - \lambda/\bar{\pi}, 0)$, so is $D_{x^*}(x^*-)$.

Property 3 follows directly from the monotonicity of $D_{x^*}(x)$. Since this is increasing, it must always be that $D_{x^*}(x^*+) \geq D_{x^*}(x^*-)$ with this holding strictly at a divergent cutoff, where the function must be discontinuous.

A cutoff x^* is consistent with an equilibrium if $D_{x^*}(x^*-) \leq c/\lambda \leq D_{x^*}(x^*+)$ or if $D_0(0+) \geq c/\lambda$. The continuity and monotonicity of $D_{x^*}(x^*+)$ implies that an x^* exists that satisfies the sequential rationality condition, since $D_{x^*}(x^*+)$ must either cross c or always lie above or below c on $[0, x^{**})$, and above x^{**}, for any cutoff $x^* \in [0, x^{**})$ it must be that $D_{x^*}(x) = \frac{1}{r+\lambda} k\bar{\pi} > c/\lambda$, so, the firm always wants to work there.

Finally, since the drift above x^{**} is positive x_t never drops below x^{**} once it has gone above it, so $V(x^{**}, L) = V_{wc}(x^{**}, L)$, where V is the value function when the firm is playing x^{**} and x^* as cutoffs, so x^{**} is still sequentially rational.

Commitment

Proof of Corollary: Consider a sequence $r_n \to 0$. In this sequence of equilibrium

$$\lim_{n\to\infty} r_n(V_n(x,\theta) - E(\int e^{-r_n t} k\pi_t)) = \lim_{n\to\infty} E\left(r_n \int_0^\infty e^{r_n} x_t dt\right) = 0,$$

since

$$r_n \int_0^\infty e^{r_n} x_t dt = r_n \int_0^T e^{-r_n t} x_t dt + \int_T^\infty e^{-(r_n+\mu t} x_t dt = 0,$$

where T is the time it takes to reach the region where the firm doesn't censor. On the other hand, if the firm has committed to not censor, for any x where $x > \min(\frac{\lambda}{\pi}-1, \limsup x^*_{c,n})$, there exists an N such that for any $n \geq N$, beliefs are drifting up in all equilibria in the sequence at that point. Finally

$$r_n V_n(x,H) = r_n \int e^{-r_n t}[x_t - c]dt > 0.$$

and

$$r_n V_n(x,L) = r_n \int e^{-(r_n+\lambda)t}[x_t - c + \lambda V_n(x,H)]dt \geq \frac{\lambda}{r_n+\lambda}(r_n V_n(x^*_{c,n},H)) > 0,$$

so the firm always benefits from commitment.