Experimenting with Career Concerns*

Marina Halac† Ilan Kremer‡

July 18, 2018

Abstract

A manager who learns privately about a project over time may want to delay quitting it if recognizing failure/lack of success hurts his reputation. In the banking industry, managers may want to roll over bad loans. How do distortions depend on expected project quality? What are the effects of releasing public information about quality? A key feature of banks is that they learn about project quality from bad news, i.e. a default. We show that in such an environment, distortions tend to increase with expected quality and imperfect information about quality. Results differ if managers instead learn from good news.

*We thank Simon Board, Patrick Bolton, Tri Vi Dang, Peter DeMarzo, Wouter Dessein, Marco Di Maggio, Brett Green, Zhiguo He, Matt Jackson, Navin Kartik, Avi Lichtig, Qingmin Liu, Matteo Maggiori, George Mailath, Konrad Mierendorff, Pauli Murto, Andrea Prat, Canice Prendergast, Andy Skrzypacz, Lars Stole, Phillip Strack, various seminar and conference audiences, and two anonymous referees for helpful comments. We also thank Johannes Hörner, Paul Milgrom, Dan Quint, and Curt Taylor for valuable discussions of the paper. Yifeng Guo provided excellent research assistance.

†Yale University and CEPR. Email: marina.halac@yale.edu.

‡The Hebrew University and University of Warwick. Email: ikremer@huji.ac.il.

When a manager learns privately about a project over time, the market cannot assess the full consequences of his behavior, and the manager may want to take suboptimal actions that make a better impression (e.g., Prendergast and Stole, 1996). In particular, a manager may want to delay quitting a project if recognizing failure/lack of success hurts his reputation. How do distortions depend on expected project quality? Are managers more likely to keep bad ventures during good times, when expected quality is higher? What are the effects of releasing public information about project quality?

Following the financial crises over the last 30 years, there has been a growing concern about banks' behavior during boom times. It is by now well documented that financial crises have often been preceded by credit booms (e.g., Schularick and Taylor, 2012). One reason for unhealthy credit growth is that banks may lower their standards and lend to low-quality borrowers who are unlikely to repay in a downturn. But there is also another reason for concern: during good times, banks may prefer not to force bankruptcy and instead roll over bad debt, providing life support to projects that, from an economic point of view, should be terminated. This practice of rolling over bad loans indeed appears to have played an important role in many crises, including Japan's in the early 90s,[1] and is at the center of current concerns about China. According to The Wall Street Journal (2013), "the reason China's bad-debt levels are so low boils down to the tendency of the country's banks to routinely extend or restructure loans to borrowers, or sell them, rather than admit they have gone bad and record a loss in their accounts."[2] These problems have prompted efforts in several countries to generate more information about banks' assets, for example by adopting stress testing and requiring public disclosure of test results (e.g., Hirtle and Lehnert, 2014).

In a seminal paper, Rajan (1994) finds evidence of banks rolling over bad debt in good times and proposes a simple static model of career concerns to explain bank managers' incentives. The pattern of distortions in a richer setting, however, is not obvious. On the one hand, as Rajan (1994) points out, a bank's reputation cost of recognizing bad debt is larger in good times, when the perceived quality of loans is higher, compared to bad times. On the other hand, the fraction of problematic loans is also smaller in good times, so the potential for distortions is lower. The effects of releasing information are also *a priori* unclear: perfect public

[1] Sekine, Kobayashi and Saita (2003) and Peek and Rosengren (2005) find evidence that banks in Japan continued to lend to severely impaired borrowers in order to avoid realizing losses on their own balance sheets. See also Caballero, Hoshi and Kashyap (2008).

[2] The article shows that while China had a low nonperforming loan ratio in 2012 compared to other countries, its nonperforming loans in billions of yuan were steadily increasing over 2011-2013. According to The Economist (2014), "a culture of bankruptcy should replace the lifelines and 'evergreening' of useless loans" for China not to "repeat Japan's malaise." Evaluating credit growth in China a year later, The Economist (2015) reports that "it is not hard to find examples of companies on life support that in other countries might have perished by now."

information about loan quality would eliminate any scope for distortions, but is imperfect information also beneficial? We develop a dynamic model of career concerns to examine the pattern of distortions and the value of information. Our dynamic framework reveals that the nature of information managers receive over time is important for understanding these issues. Specifically, a key factor in the banking industry is that managers learn about project quality from the arrival of "bad news": the structure of debt contracts implies that banks get more information when a borrower is in distress and defaults than when the loan is paid in full.[3] This contrasts, for example, with the case of an entrepreneur investing in a technological innovation, who learns from the good news that arrive when a breakthrough occurs.

In our model, a manager decides at each time whether to continue to invest in a given project or abandon it. The quality of the project can be either good or bad and is initially unknown to all parties. The manager cares not only about the payoffs from the project but also about the market's perception of the project's quality. The market's perception is based on the publicly observable actions of the manager, namely whether he continues or not with the project. The manager learns about the project's quality from privately observed lump sum payoffs that arrive at random times. We contrast two scenarios. Our main focus is on a bad news setting, where the manager learns from the arrival of a negative payoff that indicates that the project is bad (and thus expected to generate losses). Here "no news is good news": over time, in the absence of a negative payoff, the manager becomes more optimistic about the quality of the project. As a benchmark we examine a good news setting, where the manager learns from the arrival of a positive payoff that indicates that the project is good (and thus expected to generate positive profits). In this case "no news is bad news": the manager becomes more pessimistic as time passes and a positive payoff does not arrive.

We characterize the manager's decision of whether and when to abandon the project, and how in turn the market updates its belief over time. In both of the information environments we study, we solve for the (essentially) unique equilibrium in closed form.[4] We show that the manager's career concern generates an inefficiency: the manager runs the project for too long relative to the first-best solution that maximizes the expected payoff from the project. In the setting in which learning about project quality is through the arrival of bad news, the manager follows a pure strategy: he abandons the project if and only if bad news arrive before a given date t^*. As the manager continues with the project before reaching t^*, the market's belief that the project is good increases; beyond t^*, the reputation cost of quitting is so large that the manager prefers to continue with the project even when he knows it will generate losses. In

[3] As Townsend (1979) and Dang, Gorton and Holmström (2015a,b) argue, debt contracts enable banks to minimize the cost of monitoring by not acquiring information when there is no default.

[4] Section 1 defines our equilibrium concept. We use a refinement in the same spirit as Divinity.

the good news setting, instead, the manager uses a mixed strategy: as time passes without the arrival of good news and the manager becomes more pessimistic, he follows a random quitting policy, abandoning the project at a later time than the efficient (pure) stopping time. In both settings, distortions are increasing, and welfare is decreasing, in the importance of career concerns.

Our characterization yields two main results. First, we show that in the bad news setting, distortions are more pronounced when the expected quality of the project is relatively higher.[5] Distortions in this setting take the form of the manager keeping the project after learning that it is bad, namely when bad news first arrive after time t^*. When the prior probability of a good project increases, t^* decreases, meaning that the manager is even less likely to quit following bad news. We show that this higher tendency not to terminate bad projects more than compensates for the fact that a bad project is less likely, so the overall distortion increases when expected quality rises. This result contrasts with what we find in the good news setting. Distortions in that setting take the form of the manager keeping the project after enough time has passed and good news have not arrived, namely after he has reached the efficient stopping time without good news. When the prior probability of a good project increases, the manager is more likely to keep the project for a longer period of time; however, the efficient stopping time also increases, so good news are more likely to arrive prior to this time. We show that as a consequence, distortions can decrease with expected quality in the good news setting.

Our second main result concerns the effects of information. Suppose that it is possible to release a public signal at the beginning of time that makes the market and manager's common prior on the project more precise. This signal allows for a better assessment of the quality of projects and thus always weakly increases first-best welfare. The effects on equilibrium welfare depend on how the signal affects distortions. Naturally, if the signal is perfect, it eliminates distortions, as the manager's actions cannot influence the market's belief when the project's quality is known. We show however that in the bad news setting, the effects of information are non-monotonic: a sufficiently imperfect signal increases the distortion relative to first best and reduces welfare. Intuitively, a high signal realization increases the prior that the project is good and (by the result described above) increases the distortion relative to first best, whereas a low signal realization reduces the prior and thus reduces the distortion. We find that distortions are convex in the prior, and therefore the former effect dominates when the signal is sufficiently imperfect.[6] Moreover, such a signal leaves first-best welfare unchanged,

[5] We describe here our results for "intermediate" parameter values, under which the equilibrium features an interior time $t^* \in (0, \infty)$. See Section 2 for details.

[6] As elaborated in Section 2, the convexity results from the fact that an increase in the prior brings the distortion forward in time while a reduction in the prior postpones it, and discounting is convex.

so the increase in distortions causes overall welfare to go down. These results are in contrast with what we find in the good news setting: releasing a public signal about project quality at the beginning of time always increases welfare when the manager learns through good news.

Our analysis highlights the role of dynamic aspects in shaping the effects of career concerns and yields different predictions for different applications. Returning to the banking industry, we find that the nature of learning in this industry explains why banks generate distortions especially in good times, and moreover implies that releasing information about the quality of credit may not help but rather exacerbate distortions. An implication of our analysis is the need for policy that pays close attention to refinancing during boom periods as well as the effects of supervisory tests and disclosure requirements. Simple restrictions are unlikely to do the job; in fact, policy aimed at reducing bad debt must deal with the problem that banks are often "creative" when it comes to refinancing loans. For the case of China, The Wall Street Journal (2013) explains that "banks need a reason to justify rolling over a loan, particularly if a company can't repay it. (...) When they do roll over loans, Chinese banks sometimes do it in creative ways. To skirt restrictions on rolling over loans, banks cooperate with informal lenders that provide bank customers with short-term loans with high interest rates. That borrowing is used to repay a bank loan on the understanding that the bank will issue a new loan two or three weeks later. Such behavior can, in some instances, lead to bigger corporate-debt burdens."

Related literature. There is a large literature on career concerns. One strand of this literature, in the tradition of Holmström (1999), studies moral hazard models in which career concerns are beneficial because they incentivize agents to exert effort. These are models where outcomes are observable but actions are not.[7] Our paper fits into a different strand of this literature, in which career concerns are detrimental because they lead to perverse incentives. Here actions are observable but outcomes are not. A seminal paper is Prendergast and Stole (1996), where an agent has private information about his ability to understand the state of the world and distorts his decisions over time to look as a fast learner. Related issues are studied in Scharfstein and Stein (1990), Zwiebel (1995), Majumdar and Mukand (2004), Prat (2005), Ottaviani and Sørensen (2006), and Aghion and Jackson (2016).[8]

Within finance, as mentioned above, Rajan (1994) studies a static model in which a career-concerned bank manager chooses whether to implement a liberal credit policy that makes a bad loan less visible. Makarov and Plantin (2015) show that a fund manager will want to take

[7] Bonatti and Hörner (2015) consider a version of Holmström's model with exponential learning.

[8] Reputational concerns also lead to bad outcomes in Morris (2001) and Ely and Välimäki (2003), where an agent's type determines whether his preferences are aligned with those of the principal.

on hidden tail risk when concerned with investors' perception of his ability to generate excess returns above a fair compensation for risk. A number of papers point out that managers with stock-based compensation may still have a conflict of interest with shareholders (see Bond, Edmans and Goldstein, 2012, Section 3 for a survey), although only a few study how this conflict can result in actions directly aimed at concealing or revealing information. Among these, Benmelech, Kandel and Veronesi (2010) show that to prevent stock price reductions, managers may use suboptimal investment policies that conceal slowdowns in the firm's growth opportunities. Frenkel (2017) examines a dynamic bargaining model in which managers with stock-sensitive compensation distort the price and timing of over-the-counter asset sales.

Our model is more closely related to those in Grenadier, Malenko and Strebulaev (2014), Bobtcheff and Levy (2017), and Thomas (2016), all of which consider environments with exponential learning.[9] Grenadier et al. (2014) examine an experimentation setting with public good news in which an agent is privately informed about his value of project success. The agent delays stopping because, unlike in our model, stopping at a later time signals a higher type.[10] Bobtcheff and Levy (2017) analyze a real option model in which a cash-constrained agent may learn bad news prior to investing. The agent wants to convey that his privately known learning intensity is high to raise capital more cheaply, and this can lead to hurried or delayed investment. Thomas (2016) studies a career concerns problem similar to ours, in which an agent learns about a project over time and can choose to abandon it. In her model, however, successes and failures are public, the agent learns privately from partially informative signals without cash-flow consequences, and the agent receives a reputation payoff only when the project succeeds, fails, or is terminated. The focus is also different: Thomas examines the conditions under which efficiency obtains,[11] whereas we study the pattern of distortions, specifically how distortions vary with expected project quality and information about quality.

More broadly, our paper is related to a sizable literature on exponential-bandit learning, including the seminal work of Keller, Rady and Cripps (2005). Most of this literature studies learning through good news, but there are exceptions: in addition to the papers described above, Bonatti and Hörner (2017) and Keller and Rady (2015) consider bad news learning, and Che and Hörner (2017), Frick and Ishii (2015), and Khromenkova (2015) compare good news

[9]Also related is Bar-Isaac (2003), where a monopolist sells units over time whose (observable) success depends on the monopolist's fixed quality. While there is no private learning, the monopolist may have initial superior information about his quality, and thus his decision to continue trading can serve as a signal.

[10]The paper shows that if a public shock forces some agents to stop, then others will blend with the crowd and stop strategically at the same time. Related papers of strategic delay include Acharya, DeMarzo and Kremer (2011) and Grenadier and Malenko (2011). See also Gratton, Holden and Kolotilin (2017).

[11]In her model, whether the first best is implementable depends on the intensity of the agent's reputational concern and the sign and relative informativeness of public and private news.

and bad news learning in various contexts.[12] Some articles study how information disclosure affects experimentation, although this information is typically about outcomes rather than the underlying state as in our paper. In a career concerns setting, see Pei (2015) and, outside the exponential bandit framework, Hörner and Lambert (2016).[13]

1 Model

Players and actions. Consider an agent and a market. Time is continuous, the horizon is infinite, and the discount rate is $r > 0$. The agent has a project and, at each time $t \geq 0$, decides whether to continue working on the project or to stop. To simplify the exposition, we assume that stopping is irreversible.[14]

The quality of the agent's project is either "good" or "bad", a fully persistent state. Working on the project yields the agent an instantaneous payoff $x \in \mathbb{R}$, capturing the instantaneous cost of working and any deterministic flow revenue from the project (so that x may be positive or negative). In addition, if the agent works at time t and the project is bad, the agent receives a lump-sum payoff of -1 at t with instantaneous probability $\lambda_B \geq 0$; if the agent works at time t and the project is good, he receives a lump-sum payoff of 1 at t with instantaneous probability $\lambda_G \geq 0$. This structure allows us to embed both a bad news setting and a good news setting, as we describe below. We assume $x + \lambda_G > 0 > x - \lambda_B$, i.e. the expected payoff from the project is positive if the project is known to be good and negative if it is known to be bad. The payoff from not working is normalized to zero.

Information. The quality of the agent's project is initially unknown to both the agent and the market. We denote by μ_t the agent's time-t belief that the project is good. The exogenous prior belief is $\mu_0 \in (0, 1)$, commonly known also to the market.

The market only observes the agent's decision at each point of whether to continue or to stop working on the project; the realized payoffs from the project are nonverifiable and privately observed by the agent. Hence, at any time $t > 0$, the market's belief about the project may differ from the agent's belief, as it is updated based on the agent's actions only. We denote by $\widehat{\mu}_t$ the market's time-t belief that the project is good.

[12] See also Abreu, Milgrom and Pearce (1991) and Board and Meyer-ter-Vehn (2012).

[13] Rosenberg, Solan and Vieille (2007) study a model of social experimentation in which agents observe their opponents' actions but not their opponents' payoffs, and the decision to abandon the risky project is irreversible. Quah and Strulovici (2009) examine the comparative statics of optimal stopping time problems in a general framework.

[14] One can show that given our solution concept, this irreversibility assumption will be without loss.

Since the quality of the project is uncertain and may be learnt when the agent works, we say that the agent "experiments" when he runs the project.

Bad news versus good news. We focus on a setting in which the agent learns about project quality through bad news events: $\lambda_B > \lambda_G = 0$, and, to avoid trivialities, we then assume $x > 0$ (where, as noted, $x - \lambda_B < 0$). The agent therefore earns small profits so long as he does not experience a "failure," namely a lump-sum payoff of -1. Learning in this bad news environment takes the form of slow improvement in the agent's belief μ_t until the agent fails and learns that the project is bad. We contrast this setting with one in which learning occurs through good news events: $\lambda_G > \lambda_B = 0$, and (again to avoid trivialities) we then assume $x < 0$ (where, as noted, $x + \lambda_G > 0$). Here the agent incurs small losses so long as he does not experience a "success," namely a lump-sum payoff of 1. Learning in this good news environment takes the form of slow deterioration of the agent's belief μ_t until the agent succeeds and learns that the project is good.

Payoffs. The agent cares not only about the payoff from the project but also about how his project is perceived by the market. The quality of the project reflects the agent's skills and potential to select and successfully work on new projects. A career-concerned agent will therefore want the market to believe that his project is good rather than bad.

Following Rajan (1994) and Prendergast and Stole (1996), we take the agent's payoff at each time $t \geq 0$ to be a weighted sum of the project payoff and the reputation payoff the agent receives from the market's perception, $\widehat{\mu}_t R$, where $R \geq 0$. Specifically, if the agent works on the project only until time τ, his time-0 payoff is

$$\int_0^{\infty} e^{-rt} \left\{ \left[x + \mu_0 \lambda_G - (1 - \mu_0) \lambda_B \right] \mathbf{1}_{t<\tau} + \widehat{\mu}_t R \right\} dt, \tag{1}$$

where $\mathbf{1}_{t<\tau}$ is an indicator function taking the value 1 if $t < \tau$ and zero otherwise. This formulation captures, in a simple reduced form, any benefits the agent may enjoy from having a high reputation, such as outside options that increase the agent's wage or the possibility of working on additional projects.[15] For example, like money managers in Dasgupta and Piacentino (2015), banks differ in their ability to pick good investments, and they benefit from having a reputation for being good pickers as this allows them to attract more capital. Entrepreneurs and venture capital funds have similar career concerns, especially since they must repeatedly raise money from investors, as emphasized by Baker (2000).[16]

[15] The agent's reputation payoff could also arise from a dependence of his compensation on stock prices which reflect the market's valuation of the project.

[16] One can consider a variant of our model in which, as in Rajan (1994), the quality of the project is

We will refer to expression (1) when $R = 0$ as social welfare. Controlling for project quality, society does not benefit from the agent having a high reputation, and the efficient allocation of productive resources is the one that maximizes the profits from the project. For example, if the project is a publicly traded company, then welfare corresponds to the utility of the investors.

To focus the exposition, throughout the paper we assume parameters are such that some experimentation is socially efficient, even if the agent is myopic:

Assumption 1. *Some experimentation is always efficient:* $x + \mu_0\lambda_G - (1-\mu_0)\lambda_B > 0$.

Strategies and equilibrium. The agent's history at time t, h^t, consists of his private history of payoff realizations up to t and the public history of the agent's actions up to t. Let $\{\Omega, \mathcal{H}, \mathcal{P}\}$ be the probability space and $(\mathcal{H}_t)_{t\geq 0}$ the filtration generated by the history h^t. A pure strategy for the agent is an $\mathcal{H}_t$-adapted stopping time $\tau : \Omega \to \mathbb{R}_+ \cup \{\infty\}$. Following Shmaya and Solan (2018), we define a mixed strategy for the agent as a randomized stopping time, namely an adapted $[0,1]$-valued process $\rho = (\rho_t)_{t\geq 0}$ with right-continuous nondecreasing paths. This mixed strategy is a cumulative distribution function that measures, for each $t \geq 0$, the probability to stop before or at time t.

The agent's beliefs $\mu = (\mu_t)_{t\geq 0}$ are adapted to $\mathcal{H}_t$. The market's beliefs $\widehat{\mu} = (\widehat{\mu}_t)_{t\geq 0}$ are adapted to the filtration generated by the agent's stopping time.

An equilibrium is defined as a pair $\{\rho, \widehat{\mu}\}$ such that: (i) given $\widehat{\mu}$, every stopping time τ in the support of the randomized stopping time represented by ρ solves

$$\sup_{\tau} \int_0^{\infty} e^{-rt} \left\{ [x + \mu_0\lambda_G - (1-\mu_0)\lambda_B] \mathbf{1}_{t<\tau} + \widehat{\mu}_t R \right\} dt,$$

and (ii) given ρ, $\widehat{\mu}$ is computed by Bayes' rule for all public histories on the equilibrium path.[17] As is standard in signaling games, we use a refinement to rule out equilibria that can arise only due to unreasonable beliefs off the equilibrium path. Let $\widehat{\mu}_t^1$ be the market's time-t belief that the project is good conditional on the agent not having stopped by t, and let $\widehat{\mu}_t^0(t')$ be the market's time-t belief conditional on the agent having stopped at time $t' \leq t$. We require the following *belief monotonicity* property on off-the-equilibrium-path beliefs: if $\tau < t$ with probability one (that is, $\rho_t = 1$), then $\widehat{\mu}_t^1 \geq \widehat{\mu}_t^0(t')$ for all $t' \leq t$.

determined by the state of the economy and the ability of the agent, and the agent cares about the market's perception of his ability. Our main results would apply to such a setting. Details are available upon request.

[17] We note that in our game, the agent has no private information at time 0. Hence, the market's belief cannot change upon observing the agent's decision of whether to start or not the project. This is sometimes described as a "no signaling what you don't know" condition (see Watson, 2016).

Our refinement is in the same spirit as the Divinity refinement introduced by Banks and Sobel (1987),[18] and is similar to the belief monotonicity requirements used in other continuous-time signaling models such as Daley and Green (2012), Gul and Pesendorfer (2012), and Strebulaev, Zhu and Zryumov (2016). In our game, for any given beliefs of the market, continuing with the project at a time $t > 0$ is always more attractive to an agent who has not failed/has succeeded by t than to one who has failed/has not succeeded by t. Hence, we require that if the agent continues at a time by which the candidate equilibrium strategy specifies having stopped, the market should assign weakly higher probability to the agent not having failed/having succeeded, and hence to the agent's project being good, than if the agent had indeed stopped.[19] Without belief monotonicity, one could construct equilibria in which the market's off-the-equilibrium-path beliefs "punish" the agent for deviating to continuing with the project, therefore forcing him to abandon the project by an arbitrary time.[20]

We will indicate where we use belief monotonicity in our discussions in the text and in the proofs in the Appendix. From now on, equilibrium refers to an equilibrium as defined above satisfying this refinement.

2 Bad news

Consider a setting in which the agent learns about project quality from the arrival of a failure: $\lambda_B > x > \lambda_G = 0$. With a slight abuse of notation, denote by μ_t the agent's belief that the project is good at time t given that he has run the project and not failed up to t. By Bayes' rule:

$$\mu_t = \frac{\mu_0}{\mu_0 + (1-\mu_0)e^{-\lambda_B t}}. \tag{2}$$

The evolution of this belief is governed by

$$\dot{\mu}_t = \mu_t (1-\mu_t) \lambda_B. \tag{3}$$

As the agent works without failing, his belief that the project is good goes up. If at any time the agent fails, his belief jumps down to zero.

[18] We cannot apply Divinity directly as we study an infinite horizon continuous-time game.

[19] As we will see, Bayes' rule implies that the market's belief weakly increases upon observing an agent's on-path decision to continue, and belief monotonicity extends this property to off the equilibrium path.

[20] Such equilibria would fail forward induction reasoning: as noted, the agent who has not failed/has succeeded, and whose project is thus more likely to be good, is precisely the type of agent who has relatively more to gain from continuing with the project.

2.1 First best

Suppose $R = 0$, so the agent does not have a career concern and maximizes social welfare. Since a failure reveals that the project is bad, the first-best solution prescribes abandoning the project as soon as it fails. Moreover, since the agent's belief that the project is good increases over time in the absence of failure, an agent who starts working should continue working so long as the project has not failed. The value of starting the project at time 0 is

$$S_0^{FB} = \mu_0 \frac{x}{r} + (1 - \mu_0) \frac{x - \lambda_B}{r + \lambda_B} > 0, \tag{4}$$

where the inequality follows from Assumption 1. Hence, the first best entails working on the project absent failure and stopping immediately when a failure occurs.

2.2 Career concerns

Consider now the setting with $R > 0$, i.e. where the agent cares about both the payoff from the project and the market's belief about project quality. Because an agent with a good project cannot fail whereas one with a bad project can, the agent would like to make the market believe that a failure has not occurred.

Suppose the agent starts working on the project at time 0. We begin by showing that the agent never stops in the absence of failure.

Lemma 1. *In any equilibrium, if the agent starts the project at time* 0 *and does not fail by time* $t > 0$*, he continues with the project at* t*.*

(All proofs are in the Appendix.)

To see the logic, suppose the agent starts working at time 0 and continues without failing until time $t > 0$. Since the project's expected payoff at t is then strictly positive, the agent would choose to stop at t only if stopping gives him a reputation gain compared to continuing. However, this cannot happen on the equilibrium path: if an agent who has not failed was willing to stop at t, an agent who has failed would strictly prefer to stop by t, and therefore stopping would not increase the market's belief that the agent's project is good. Furthermore, by our belief monotonicity refinement, this cannot happen off the equilibrium path either.

Lemma 1 implies that if the agent stops at a time $t > 0$, the market learns that the agent has failed by t. Hence,

Corollary 1. *If the agent stops at time* $t > 0$ *in equilibrium, the market's belief that the project is good at* $t' \geq t$ *is zero.*

Note that the market's belief upon observing that the agent stops is independent of the time at which he stops: the market learns that the agent has failed, but the time at which the agent failed contains no information about project quality.

Consider next the market's belief that the project is good when the agent has not stopped by t, which we denote by $\widehat{\mu}_t^1$. Given that an agent who has not failed always continues with the project, this belief is determined by whether and when an agent who has failed stops. Suppose the agent were to follow the first-best strategy, i.e. stop immediately upon failure. Then the market's belief would be $\widehat{\mu}_t^1 = \mu_t$, where μ_t is given by (2); that is, the market's belief about the project would coincide with the agent's belief. This implies that $\widehat{\mu}_t^1$ would be increasing over time.

Given these market's beliefs, would the agent indeed have incentives to stop immediately upon failing at a time $t > 0$? If the agent stops at t, his continuation payoff is zero (cf. Corollary 1). Suppose instead that the agent continues working on the project after failing at t. Since the market's belief is increasing over time so long as the agent has not stopped and it is constant and independent of the time at which the agent stops after he stops, the agent continues working forever if he continues working after failing at t. Hence, the agent's expected payoff from continuing is

$$\int_t^{\infty} e^{-r(s-t)} \left(x - \lambda_B + \mu_s R\right) ds. \tag{5}$$

The agent is willing to stop at a time t at which he fails if and only if (5) is negative at this time. Since the first-best strategy requires that the agent stop whenever a failure occurs, (5) must be negative at all $t > 0$; given $\lim_{t\to\infty} \mu_t = 1$, this requires

$$R \leq -(x - \lambda_B). \tag{6}$$

In addition, the first-best strategy prescribes the agent to start working at time 0; as explained below, Assumption 1 ensures that the agent indeed has incentives to do so.

Condition (6) is satisfied if the agent's career concern is sufficiently small. If instead R is large enough that this condition fails, there will exist a time $t > 0$ at which the agent will not want to stop upon failure. Intuitively, from that time on, the agent's cost of losing his reputation would exceed his cost of continuing working on a bad project forever.

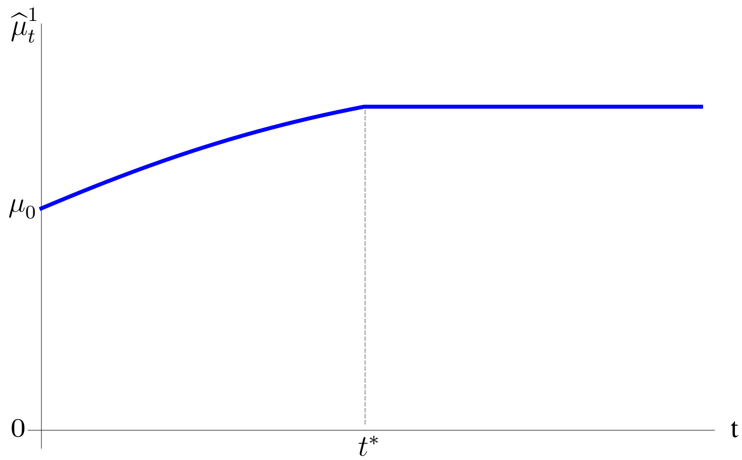

Figure 1: Market's belief and threshold time in the equilibrium of the bad news setting. Parameters are $\mu_0 = 0.6, x = 0.75, \lambda_B = 1.8, R = 1.2$, and $r = 1$.

Proposition 1 (Bad news setting)**.** *The equilibrium is unique. There exists a threshold time $t^* \geq 0$ such that the agent starts working at time 0, stops immediately if he fails at $t < t^*$, and continues working otherwise.*
The market's time-t belief conditional on the agent not having stopped by t is $\widehat{\mu}_t^1 = \mu_t$ for $t \leq t^*$ and $\widehat{\mu}_t^1 = \mu_{t^*}$ for $t > t^*$. The market's time-t belief conditional on the agent having stopped by t is 0, both on and off the equilibrium path.
The threshold time satisfies $t^* = \infty$ if $R \leq -(x - \lambda_B)$ and $t^* = 0$ if $R \geq -(x - \lambda_B)/\mu_0$. If $-(x - \lambda_B) < R < -(x - \lambda_B)/\mu_0$, then t^* is given by

$$x - \lambda_B + \mu_{t^*} R = 0. \tag{7}$$

The equilibrium is characterized by a threshold time $t^* \geq 0$ such that the agent stops at t if and only if a failure occurs at t and $t < t^*$. As discussed above, if $R \leq -(x - \lambda_B)$, the equilibrium implements the first best, so $t^* = \infty$ in this case. At the other extreme, if $R \geq -(x - \lambda_B)/\mu_0$, the agent always prefers to continue after failure, so $t^* = 0$. The threshold time is interior when $-(x - \lambda_B) < R < -(x - \lambda_B)/\mu_0$. In this case, the threshold time is the time t^* at which the agent is indifferent between stopping and continuing given that he has failed at t^*, given by equation (7). Figure 1 illustrates the market's equilibrium belief $\widehat{\mu}_t^1$ in an example with such an interior time t^*.

The agent's time-0 expected payoff from following the equilibrium strategy with threshold time t^* is

$$\pi_0(t^*) = \mu_0 \left(\frac{x}{r} + \int_0^{t^*} e^{-rt} \mu_t R dt + e^{-rt^*} \frac{\mu_{t^*} R}{r} \right)$$
$$+ (1 - \mu_0) \left[\int_0^{t^*} e^{-(r+\lambda_B)t} (x - \lambda_B + \mu_t R)\, dt + e^{-(r+\lambda_B)t^*} \frac{(x - \lambda_B + \mu_{t^*} R)}{r} \right].$$

A sufficient condition for the agent to prefer following the equilibrium strategy rather than never working is $\pi_0(0) \geq \mu_0 R/r$, which is satisfied by Assumption 1. In fact, given this assumption, we can show that a no-work equilibrium (in which the agent never works on the project) does not exist. Intuitively, if a no-work equilibrium exists, it exists when the market's beliefs are such that the agent's reputation benefit from starting the project at time 0 is minimized. These beliefs correspond to the market expecting the agent to never stop once he starts working — so that $\widehat{\mu}_t^1 = \mu_0$ for all $t > 0$ — and believing that the agent has failed if he ever stops.[21] However, since the agent's payoff from never working is $\mu_0 R/r$, it is clear that under Assumption 1 the agent would prefer working forever to never working.

We can therefore show that the unique equilibrium is the equilibrium characterized in Proposition 1. Suppose the agent's career concern R is intermediate so that the threshold time t^* is interior. Applied to the banking industry, the equilibrium says that a bank will stop rolling over a borrower's debt if it learns early enough that the borrower is in distress. However, as the borrower repays and the bank keeps lending, the bank's reputation increases. At some point, the reputational cost of admitting losses becomes high enough that the bank would choose to roll over a bad loan.[22]

Since a bank that keeps the debt until time t^* continues refinancing it regardless of its value, the market learns no information about the quality of the bank's loans after this time. Yet, note that the market expects more losses as time goes by. Specifically, let η_t be the probability that the agent has failed by time t given that the agent has not stopped by t. In the first best, $\eta_t = 0$ for all $t \geq 0$, since the agent stops immediately when a failure occurs. Instead, in an equilibrium with finite threshold time t^*, $\eta_t = 0$ for $t < t^*$ and

$$\eta_t = (1 - \mu_{t^*}) \left[1 - e^{-\lambda_B (t - t^*)}\right]$$

[21] Recall that the agent has no private information at time 0; hence, the market's belief cannot change upon observing the agent's start decision.

[22] While here the agent may keep a bad project forever, a variant of our model in which bad news become public with positive probability would yield analogous results and ensure that the agent stops in finite time. Details are available upon request.

for $t \geq t^*$, since the agent does not stop upon failing after t^*.

Corollary 2. *The market's belief that the agent has failed conditional on the agent not having stopped is increasing over time.*

For the banking industry, these equilibrium dynamics have the flavor of a "crisis buildup": banks continue rolling over (bad) loans and the market becomes increasingly concerned that banks are accumulating losses.

2.3 Expected quality and information

Social welfare in an equilibrium with threshold time t^* is equal to

$$S_0(t^*) = \mu_0 \frac{x}{r} + (1-\mu_0)(x-\lambda_B)\left[\frac{1-e^{-(r+\lambda_B)t^*}}{r+\lambda_B} + \frac{e^{-(r+\lambda_B)t^*}}{r}\right]. \tag{8}$$

Note that $S_0(t^*)$ coincides with S_0^{FB} in equation (4) if and only if $t^* = \infty$. Clearly, *ceteris paribus*, welfare $S_0(t^*)$ decreases, and the distortion $S_0^{FB} - S_0(t^*)$ increases, when the threshold time t^* declines.

The welfare effects of career concerns are then immediate from Proposition 1. The higher is the agent's concern for his reputation R, the lower is the equilibrium threshold time t^*, and hence the lower is social welfare and the larger is the distortion relative to first best.

What are the welfare effects of an increase in the expected quality of projects? Suppose the prior probability of a good project, μ_0, increases. First-best welfare then naturally goes up. Moreover, since the agent's career concern distorts actions away from first best only when the project is bad, an increase in μ_0 has a direct effect of decreasing the distortion for any fixed threshold time t^*. However, an increase in μ_0 also reduces the equilibrium threshold time t^*: when expected project quality is higher, the agent's reputation loss from ending the project at any time t increases, and the time t^* after which the agent prefers to continue upon failing decreases. We show that in net, an increase in μ_0 increases welfare, but, for interior solutions, it also increases the distortion generated by the agent's career concern.

Proposition 2 (Expected quality in bad news setting)**.** *Suppose parameters* $\{\mu_0, x, \lambda_B, R, r\}$ *satisfy* $-(x-\lambda_B)/\mu_0 > R > -(x-\lambda_B)$ *(so the equilibrium features* $t^* \in (0,\infty)$*) and consider changes in* μ_0 *that preserve this property and Assumption 1. An increase in* μ_0 *increases welfare but it also increases the distortion relative to first best.*

For intuition, note that the distortion relative to first best is equal to the losses the agent generates when he does not stop upon failing. By Proposition 1, the agent does not stop

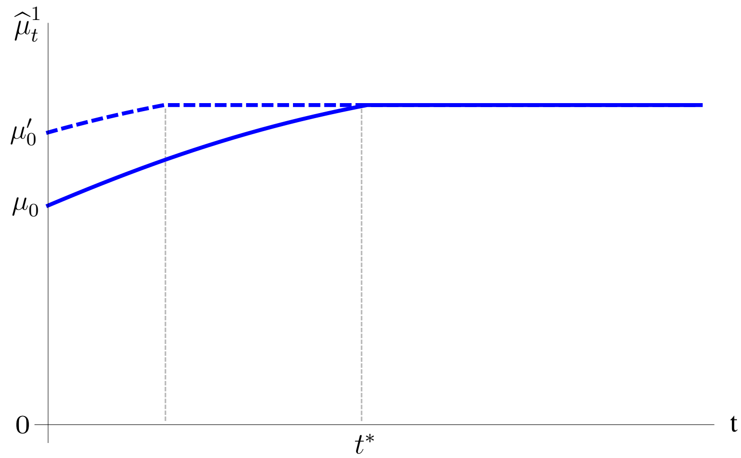

Figure 2: Effects of an increase in μ_0 in the equilibrium of the bad news setting. Parameters are the same as in Figure 1, with $\mu_0' = 0.8$.

if failure first occurs after the market's belief has reached μ_{t^*}, given by (7). Observe that μ_{t^*} is independent of μ_0. Hence, once the market's belief reaches μ_{t^*}, the probability that the project is bad $(1 - \mu_{t^*})$, and thus the expected losses the agent generates by continuing, are also independent of μ_0. An increase in μ_0 however reduces the time that it takes for the market's belief to reach μ_{t^*}; that is, as illustrated in Figure 2, t^* decreases with μ_0. This reduction in t^* has two implications: first, the agent is less likely to fail by t^*, and second, the losses occur earlier in time and are thus less heavily discounted. As a consequence, both the probability of a distortion and the present value of the distortion increase when μ_0 goes up.

Proposition 2 considers parameters under which the equilibrium features a distortion but the distortion is not extreme (i.e. $-(x - \lambda_B)/\mu_0 > R > -(x - \lambda_B)$). If instead the agent's reputational concern R is small enough that the equilibrium implements the first best (i.e. $R < -(x - \lambda_B)$), then distortions are zero regardless of μ_0 and welfare increases with μ_0. Similarly, if R is large enough that the equilibrium distortion is already extreme (i.e. $-(x - \lambda_B)/\mu_0 < R$), then an increase in μ_0 has no effect on the behavior of the agent (who never abandons the project) and the distortion decreases with μ_0. Since our interest is in studying the pattern of distortions, and the effects in the latter case are simply due to a corner solution when R is too large relative to μ_0, we focus on the intermediate case.

The results in Proposition 2 contribute to the discussion mentioned in the Introduction on how banks' behavior and distortions vary in good versus bad times. During good times, the average quality of borrowers is higher than during bad times. However, because the market's expectation of loan quality is then higher, banks' reputational loss from recognizing bad loans is also larger.[23] Proposition 2 shows that, as a result, career-concerned bank managers will be more likely to roll over their bad loans during good times compared to bad times. Furthermore, despite the proportion of bad borrowers being smaller, banks will in expectation accumulate more bad debt during good times. Consistent with the empirical findings of Schularick and Taylor (2012), one may say that banks plant the seed for the next crisis during boom periods.

Would information about project quality ameliorate the welfare distortions due to career concerns? After all, it is because the quality of the project is uncertain that career concerns lead to distorted behavior. Suppose that it is possible to release a public signal at time 0 that refines the agent and market's common prior on the project, μ_0. Absent career concerns, the signal either keeps welfare unchanged — if it does not affect the decision of whether to start the project at time 0 — or increases welfare — if it does affect this start decision. However, when the agent is career-concerned, the signal also affects distortions: as implied by Proposition 2, a high realization of the signal (i.e. a realization that increases μ_0) may increase the distortion relative to first best, whereas a low realization may lower this distortion. We find that the net welfare effect, and thus the value of information, can be negative.

Proposition 3 (Information in bad news setting)**.** *Suppose parameters $\{\mu_0, x, \lambda_B, R, r\}$ satisfy $-(x-\lambda_B)/\mu_0 > R > -(x-\lambda_B)$ (so the equilibrium features $t^* \in (0,\infty)$) and consider a public signal that refines μ_0 at time 0 while preserving this property and Assumption 1 for all of its realizations. The signal increases the distortion relative to first best and lowers welfare. If instead the signal is perfect, it eliminates distortions and increases welfare.*

The first part of the proposition considers an imperfect public signal that keeps the first-best start decision, and thus first-best welfare, unchanged. As in Proposition 2, we focus on intermediate parameters (i.e. no corner solutions). To illustrate the effects of the signal on equilibrium welfare, suppose the signal is binary, i.e. it either increases the prior to $\mu_0^h > \mu_0$ or lowers it to $\mu_0^\ell < \mu_0$. Let each realization be unconditionally equally likely, so that $\mu_0 = \frac{1}{2}\left(\mu_0^h + \mu_0^\ell\right)$. Building on our discussion of Proposition 2, the signal realization affects the time t^* that it takes for the market's belief to reach μ_{t^*}, after which the agent is no longer willing to stop upon failure. The losses that the agent generates once μ_{t^*} is reached are independent of μ_0; what matters is the probability of reaching μ_{t^*} and how heavily the losses are discounted. The probability of reaching μ_{t^*} is equal to $\frac{\mu_0}{\mu_{t^*}} = \frac{1}{2}\left(\frac{\mu_0^h}{\mu_{t^*}} + \frac{\mu_0^\ell}{\mu_{t^*}}\right)$ and is thus unchanged with the public signal. However, if $t^*(\mu_0)$ is the amount of time it takes to reach μ_{t^*} from μ_0, then the losses at μ_{t^*} are discounted by $e^{-rt^*(\mu_0)}$, which is a convex function of μ_0.[24] This means that the

[23] Loan quality in reality also depends on the bank's screening. This discussion assumes that screening does not fully eliminate differences in loan quality between good and bad times.

[24] By equation (2), $e^{-rt^*(\mu_0)} = \frac{\mu_0}{1-\mu_0}\frac{1-\mu_{t^*}}{\mu_{t^*}}$, which is convex in μ_0 for a fixed posterior μ_{t^*}.

public signal increases the expected discounted losses (since $\frac{1}{2}(e^{-rt^*(\mu_0^h)} + e^{-rt^*(\mu_0^\ell)}) > e^{-rt^*(\mu_0)}$), and as a consequence it reduces welfare.

Things of course are different if the public signal is perfect, as considered in the second part of Proposition 3. If the signal fully reveals the quality of the project at time 0, the agent's actions provide no information to the market. Hence, in this case, a career-concerned agent has no incentives to distort his behavior, and the distortion relative to first best is eliminated. Since first-best welfare increases with a fully informative signal, it follows that equilibrium welfare also increases. Combined with the first part of the proposition, this result implies that the effects of information are non-monotonic: sufficient information is beneficial, but limited information is harmful.

Our results have implications for policy, especially when applied to the banking industry. Governmental authorities conduct supervisory exams on banks to produce information on expected loan performance, and have the choice of making this information publicly available or not. Since 2009, the US and Europe incorporated into their supervisory programs the use of stress tests, which are forward-looking exams with the goal of projecting left-tail risk (e.g., Hirtle and Lehnert, 2014). Unlike with more traditional exams, the US requires that the results of individual banks' stress tests be publicly disclosed. In Europe, stress test results were not published in 2009 but public disclosure was required in subsequent years.

There is a current debate among practitioners and scholars on whether the results of banks' stress tests should be publicly disclosed. Bernanke (2013) argues that disclosure provides valuable information to market participants and the public and promotes market discipline. Goldstein and Sapra (2013) also find disclosure beneficial, although they discuss various risks and challenges associated with disclosure. Our goal is not to assess the different benefits and costs of disclosing banks' stress test results, but rather to point out a potential pitfall in the view that information is always beneficial — a view that seems to be behind much of the support for stress tests and their public disclosure. Proposition 3 shows that information on loan quality can be detrimental: when imperfect, this information can exacerbate distortions due to bank managers' career concerns and reduce overall welfare.

3 Good news

We contrast the bad news setting of Section 2 with a good news setting in which the agent learns about project quality from the arrival of a success: $\lambda_G > \lambda_B = 0 > x$ (with $\lambda_G + x > 0$). With a slight abuse of notation, we now denote by μ_t the agent's belief that the project is good at time t given that he has run the project and not succeeded up to t. By Bayes' rule:

$$\mu_t = \frac{\mu_0 e^{-\lambda_G t}}{\mu_0 e^{-\lambda_G t} + 1 - \mu_0}. \tag{9}$$

The evolution of this belief is governed by

$$\dot{\mu}_t = -\mu_t (1 - \mu_t) \lambda_G. \tag{10}$$

As the agent works without succeeding, his belief that the project is good goes down. If at any time the agent succeeds, his belief jumps up to one.

3.1 First best

Suppose $R = 0$, so the agent does not have a career concern and maximizes social welfare. Since a success reveals that the project is good, the first-best solution prescribes continuing with the project forever after it succeeds. Denote by v the present value of a success, which is equal to the payoff of 1 plus the present value of continuing working on a good project: $v = 1 + \frac{\lambda_G + x}{r}$. Absent success, the agent should continue so long as the expected marginal benefit of working is larger than the marginal cost, $\mu_t \lambda_G v \geq -x$, and should stop otherwise. The first-best stopping belief is then given by

$$\mu^{FB} := \frac{-x}{\lambda_G v}, \tag{11}$$

where, recall, $x < 0$ in this good news setting. We denote by t^{FB} the associated stopping time (derived from (9) and (11)). Note that by Assumption 1, $t^{FB} > 0$.

3.2 Career concerns

Suppose now $R > 0$, so the agent cares not only about the payoff from the project but also about the market's belief about project quality. Because an agent with a good project can succeed whereas one with a bad project cannot, the agent would like to make the market believe that his project has succeeded.

Arguments analogous to those used in Lemma 1 imply that in any equilibrium, the agent continues working forever once he has succeeded. Hence, if the agent stops at a time $t > 0$ in equilibrium, the market learns that the agent has not succeeded by t, and its belief that the project is good is equal to μ_t. Note that this belief is strictly decreasing: the later the agent stops, the longer is the period of time over which the agent has worked without obtaining success, and hence the lower is the market's belief that the agent's project is good. The market's belief remains constant at μ_t at all times $s > t$ if the agent stops at t.

Consider next the market's belief at a time t given that the agent has not stopped by this time, which we denote by $\widehat{\mu}_t^1$. Since an agent who has succeeded always continues with the project, this belief is determined by whether and when an agent who has not succeeded stops. We show that in any equilibrium, the agent must follow a mixed strategy. Suppose to the contrary that the agent were to follow a pure strategy, where absent success he stops at a finite time t with certainty. Then if the agent does not stop at t, the market believes that the agent has succeeded. That is, the market's belief that the project is good would jump up to one at t. However, this implies that if the agent was willing to work until time t, he will have a strict incentive to continue at t, a contradiction. More generally, the agent's stopping policy cannot have an atom.

Proposition 4 (Good news setting)**.** *The equilibrium is unique up to off-the-equilibrium-path beliefs. There exist threshold times $\underline{t} > 0$ and $\bar{t} \geq \underline{t}$ such that the agent works with certainty until $\underline{t}$, stops with positive probability over $[\underline{t}, \bar{t}]$ absent success, and continues otherwise. If $\bar{t} < \infty$, the agent stops with certainty by $\bar{t}$ absent success.*

The market's time-t belief conditional on the agent not having stopped by t evolves continuously and satisfies $\widehat{\mu}_t^1 = \mu_0$ for $t \leq \underline{t}$, $\widehat{\mu}_t^1 \in (\mu_0, 1)$ for $t \in (\underline{t}, \bar{t})$, and $\widehat{\mu}_t^1 = 1$ for $t \geq \bar{t}$ if $\bar{t} < \infty$. The market's time-t belief conditional on the agent having stopped at $s \leq t$ is μ_s for $s \in [\underline{t}, \bar{t}]$. For $s \notin [\underline{t}, \bar{t}]$, an off-the-equilibrium-path belief equal to μ_s supports the equilibrium.

The threshold times satisfy $\underline{t} < \infty$ if and only if $R < -x/\mu_0$ and $\bar{t} < \infty$ if and only if $R < -x$. For all parameters, $\underline{t} > t^{FB}$.

The equilibrium is characterized by two threshold times, $\underline{t} > t^{FB}$ and $\bar{t} \geq \underline{t}$. The agent always works until $\underline{t}$, over-experimenting relative to the first best due to his career concern. Absent success, the agent then implements a random stopping policy over $[\underline{t}, \bar{t}]$. Figure 3 illustrates the equilibrium beliefs of the agent and the market in an example with intermediate parameters, where $\underline{t}$ is finite and $\bar{t}$ is infinite. Note that in this case, an agent who has not succeeded continues working with strictly positive probability in the limit as $t \to \infty$.

Since the agent follows a randomized stopping time, he must be indifferent absent success at each time $t \in [\underline{t}, \bar{t}]$. That is, the agent's expected payoff from stopping at $t \in [\underline{t}, \bar{t}]$ must be equal to his expected payoff from continuing for an arbitrarily small amount of time dt and stopping at $t + dt$ absent success by then:

$$\frac{\mu_t R}{r} = \widehat{\mu}_t^1 R dt + x dt + \mu_t \lambda_G dt V_t + (1 - \mu_t \lambda_G dt - r dt) \frac{\mu_{t+dt} R}{r}, \tag{12}$$

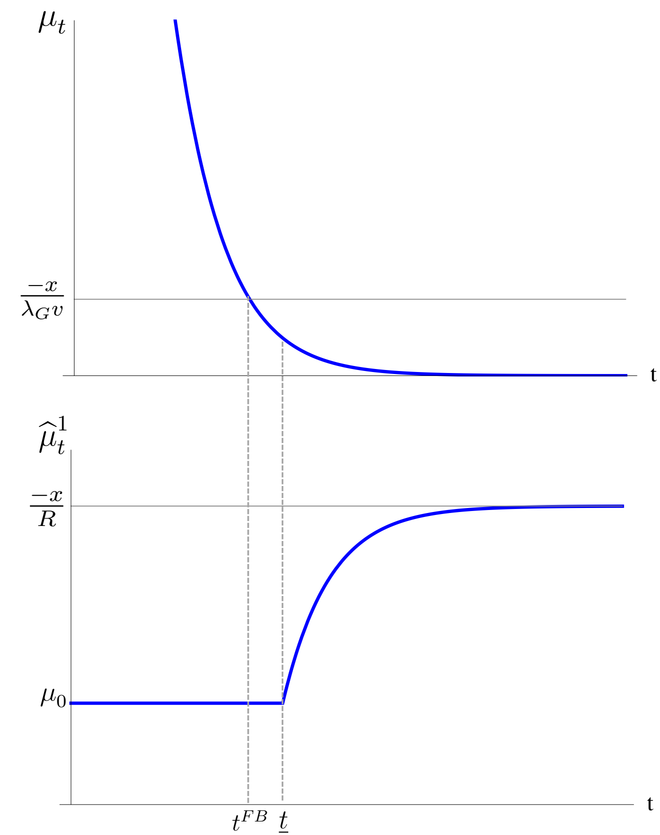

Figure 3: Beliefs and threshold times in the equilibrium of the good news setting. Parameters are $\mu_0 = 0.5, x = -0.8, \lambda_G = 6, R = 0.9$, and $r = 1$.

where $V_t = v + R\int_t^\infty e^{-r(s-t)}\widehat{\mu}_s^1 ds$ and we are ignoring $(dt)^2$ terms. Furthermore, since this condition must hold at each time $t \in [\underline{t}, \overline{t}]$, it must be that at each such time, the agent is also indifferent between stopping and continuing all the way until time $\overline{t}$ absent success. In the proof of Proposition 4, we show that the agent's indifference conditions pin down $\underline{t}$ and $\overline{t}$ and the market's belief $\widehat{\mu}_t^1$. We relegate the details to the Appendix.

A suitable application of this good news setting may be to venture capital, where an entrepreneur learns about the quality of a project from the arrival of a breakthrough. Analogous to our characterization of the bad news setting, we find that a career-concerned entrepreneur keeps the project for too long relative to the first best. Note that since the entrepreneur keeps getting pessimistic as time passes without success, he requires increasing reputation benefits not to abandon the project. This explains why the market's belief conditional on the agent continuing, $\widehat{\mu}_t^1$, must be increasing over time.

There are interesting differences between the good news and bad news settings with regards to how the agent and market's beliefs compare. Recall that in the equilibrium of the bad news setting, the agent follows a pure strategy, abandoning the project if and only if a failure arrives before a date t^*. The agent's actions are fully responsive to news up to t^* and unresponsive from t^* on; as a result, the agent and market's beliefs coincide up to t^* and diverge from t^* on. In contrast, in the equilibrium of the good news setting, the agent's actions are unresponsive to news up to $\underline{t}$ and partially responsive from $\underline{t}$ on, when the agent starts quitting with positive probability absent success. Here the agent and market's beliefs diverge from time 0 up to $\underline{t}$ and converge from $\underline{t}$ on, coinciding only when either the agent stops or a finite time $\bar{t}$ is reached.

Finally, similar to our analysis in Section 2, we can consider the market's belief that the agent has succeeded by time t given that the agent has not stopped by t. In the first best, this belief is increasing until the first-best stopping time t^{FB} and constant at one from that time on. In the presence of career concerns, this belief is also increasing: it increases strictly at all $t \leq \bar{t}$, and is constant at one from $\bar{t}$ on when $\bar{t}$ is finite. Thus, unlike under bad news learning, here the market becomes more optimistic about the outcomes of the agent's project as time passes without the agent stopping.

Corollary 3. *The market's belief that the agent has succeeded conditional on the agent not having stopped is increasing over time.*

3.3 Expected quality and information

The welfare effects of career concerns follow directly from our equilibrium characterization in Proposition 4. The larger is the weight R that the agent places on his reputation, the longer he will delay quitting the project in the absence of success. Consequently, the distortion relative to first best increases, and welfare decreases, with the agent's career concern.

We next study how expected project quality and information about quality affect distortions and welfare. Recall that in the bad news setting, we focused our attention on the case in which the agent's career concern is intermediate: the agent would prefer to continue with the project after failing if that gives him a full reputation benefit of R, but would prefer to stop upon failure if the reputation benefit is only given by the prior $\mu_0 R$. For comparison, we focus here on the analogous case, namely the case in which the agent would always prefer to continue absent success if that gives him a reputation benefit of R, but would prefer to stop if enough time has passed without success and continuing only gives him a reputation benefit of $\mu_0 R$. We obtain:

Proposition 5 (Expected quality and information in good news setting)**.** *Suppose parameters* $\{\mu_0, x,\ \lambda_G, R, r\}$ *satisfy* $-x/\mu_0 > R > -x$ *(so the equilibrium features* $\underline{t} < \infty$ *and* $\bar{t} = \infty$*). Consider changes in* μ_0 *and public signals that refine* μ_0 *at time* 0*, both preserving this property and Assumption 1. Welfare increases with* μ_0 *and with any such signal. Furthermore, there exist parameters for which the distortion relative to first best decreases with* μ_0 *and with the signal. If the signal is perfect, it eliminates distortions.*

Consider the effects of an increase in μ_0, which are illustrated in Figure 4. As in the bad news setting of Section 2, welfare increases with μ_0, but, unlike in that setting, here the distortion relative to first best can decrease with μ_0. To see why, note that the distortion in this good news setting is given by the expected losses the agent generates by continuing after his belief has reached μ^{FB}, given by (11). On the one hand, we find that if μ_0 increases, the agent stops more slowly after reaching μ^{FB}, so he generates larger losses in expectation. On the other hand, an increase in μ_0 also increases the time that it takes for the agent's belief to reach μ^{FB}; that is, as shown in Figure 4, t^{FB} is increasing in μ_0. Analogous to our discussion of Proposition 2, but now with the opposite implication, this has two effects: first, the agent is more likely to succeed by t^{FB}, and second, losses occur later in time and are thus more heavily discounted. Both of these effects imply that distortions decrease when μ_0 increases; furthermore, we show that these effects can dominate.

Regarding information, Proposition 5 shows that, unlike in the bad news setting, here releasing a public signal about project quality at time 0 always increases welfare, even if the signal is imperfect. Moreover, distortions may decline with an imperfect signal. Roughly, the main reason for the latter is that the effects of μ_0 on the agent's stopping policy are concave: a reduction in μ_0 increases the probability with which an unsuccessful agent stops by time $t > t^{FB}$ by more than what an increase in μ_0 lowers it. As for the effects on welfare, note that a signal about project quality always increases first-best welfare: while the decision to start the project may not change, the signal allows the agent to adjust the stopping time t^{FB} to avoid quitting too early or too late. Proposition 5 reveals that even if distortions were to increase with information, the effects on first-best welfare dominate. Therefore, the effects of information on equilibrium welfare are always positive when learning is from good news.

Returning to the application to venture capital, our results suggest a pattern of distortions that contrasts with that of the banking industry. We find that career-concerned entrepreneurs investing in innovation projects may generate smaller distortions during good times, when the expected quality of projects is relatively high, compared to bad times. Furthermore, no matter how imperfect, public information about the prospects of innovation opportunities is always beneficial.

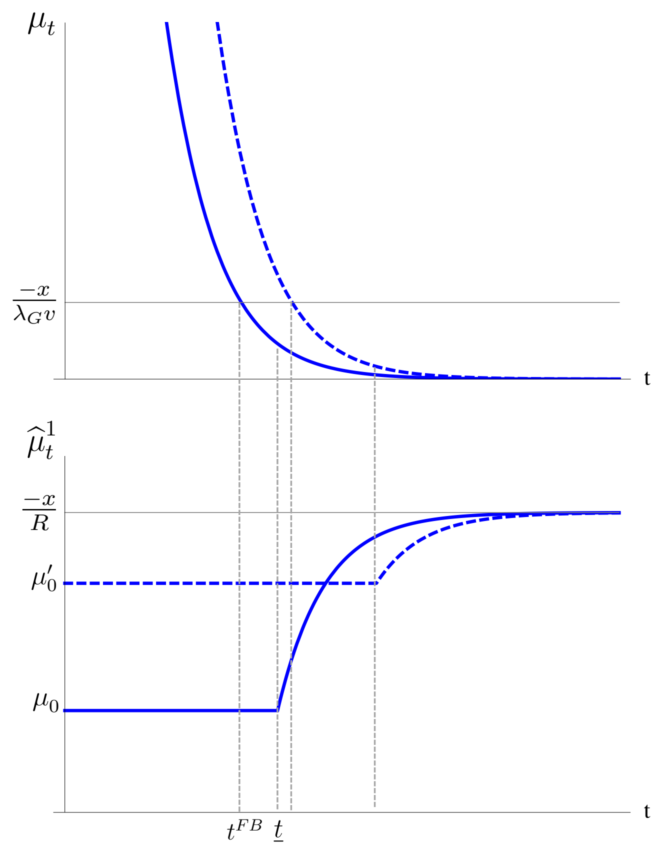

Figure 4: Effects of an increase in μ_0 in the equilibrium of the good news setting. Parameters are the same as in Figure 3, with $\mu_0' = 0.75$.

References

Abreu, Dilip, Paul Milgrom, and David Pearce, "Information and Timing in Repeated Partnerships," *Econometrica*, 1991, *59* (6), 1713–1733.

Acharya, Viral V., Peter DeMarzo, and Ilan Kremer, "Endogenous Information Flows and the Clustering of Announcements," *American Economic Review*, 2011, *101* (7), 2955–2979.

Aghion, Philippe and Matthew O. Jackson, "Inducing Leaders to Take Risky Decisions: Dismissal, Tenure, and Term Limits," *American Economic Journal: Microeconomics*, 2016, *8* (3), 1–38.

Baker, Malcolm, "Career Concerns and Staged Investment: Evidence from the Venture Capital Industry," 2000. Working Paper.

Banks, Jeffrey S. and Joel Sobel, "Equilibrium Selection in Signaling Games," *Econometrica*, 1987, *55* (3), 647–661.

Bar-Isaac, Heski, "Reputation and Survival: Learning in a Dynamic Signalling Model," *The Review of Economic Studies*, 2003, *70* (2), 231–251.

Benmelech, Efraim, Eugene Kandel, and Pietro Veronesi, "Stock-Based Compensation and CEO (Dis)Incentives," *Quarterly Journal of Economics*, 2010, *125* (4), 1769–1820.

Bernanke, Ben S., "Stress Testing Banks: What Have We Learned?," 2013. Board of Governors of the Federal Reserve System, Speech.

Board, Simon and Moritz Meyer-ter-Vehn, "Reputation for Quality," *Econometrica*, 2012, *81* (6), 2381–2462.

Bobtcheff, Catherine and Raphaël Levy, "More Haste, Less Speed: Signalling through Investment Timing," *American Economic Journal: Microeconomics*, 2017, *Forthcoming.*

Bonatti, Alessandro and Johannes Hörner, "Career Concerns with Exponential Learning," *Theoretical Economics*, January 2015, *12*, 425–475.

_ and _, "Learning to Disagree in a Game of Experimentation," *Journal of Economic Theory*, 2017, *169*, 234–269.

Bond, Philip, Alex Edmans, and Itay Goldstein, "The Real Effects of Financial Markets," *Annual Review of Financial Economics*, 2012, *4* (1), 339–360.

Caballero, Ricardo, Takeo Hoshi, and Anil K Kashyap, *American Economic Review*, 2008, *98* (5), 1943–1977.

Che, Yeon-Koo and Johannes Hörner, "Optimal Design for Social Learning," *Quarterly Journal of Economics*, 2017, *Forthcoming.*

Daley, Brendan and Brett Green, "Waiting for News in the Market for Lemons," *Econometrica*, 2012, *80* (4), 1433–1504.

Dang, Tri Vi, Gary Gorton, and Bengt Holmström, "Ignorance, Debt and Financial Crises," 2015. Working Paper.

_ , _ , and _, "The Information Sensitivity of a Security," 2015. Working Paper.

Dasgupta, Amil and Giorgia Piacentino, "The Wall Street Walk when Blockholders Compete for Flows," *Journal of Finance*, 2015, *70* (6).

Ely, Jeffrey and Juuso Välimäki, "Bad Reputation," *Quarterly Journal of Economics*, 2003, *118* (3), 785–814.

Frenkel, Sivan, "Dynamic Asset Sales with a Feedback Effect," 2017. Working Paper.

Frick, Mira and Yuhta Ishii, "Innovation Adoption by Forward-Looking Social Learners," 2015. Working Paper.

Goldstein, Itay and Haresh Sapra, "Should Banks' Stress Test Results be Disclosed? An Analysis of the Costs and Benefits," *Foundations and Trends in Finance*, 2013, *8* (1), 1–54.

Gratton, Gabriele, Richard Holden, and Anton Kolotilin, "When to Drop a Bombshell," *Review of Economic Studies*, 2017, *Forthcoming.*

Grenadier, Steven R. and Andrey Malenko, "Real Options Signaling Games with Applications to Corporate Finance," *The Review of Financial Studies*, 2011, *24* (12).

_ , _ , and Ilya A. Strebulaev, "Investment Busts, Reputation, and the Temptation to Blend in with the Crowd," *Journal of Financial Economics*, 2014, *111*, 137–157.

Gul, Faruk and Wolfgang Pesendorfer, "The War of Information," *Review of Economic Studies*, 2012, *79*, 707–734.

Hirtle, Beverly and Andreas Lehnert, "Supervisory Stress Tests," 2014. Federal Reserve Bank of New York, Staff Report No. 696.

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

Hörner, Johannes and Nicolas Lambert, "Motivational Ratings," 2016. Working Paper.

Keller, Godfrey and Sven Rady, "Breakdowns," *Theoretical Economics*, 2015, *10*, 175–202.

_ , _ , and Martin Cripps, "Strategic Experimentation with Exponential Bandits," *Econometrica*, 2005, *73* (1), 39–68.

Khromenkova, Daria, "Collective Experimentation with Breakdowns and Breakthroughs," 2015. Working Paper.

Majumdar, Sumon and Sharun W. Mukand, "Policy Gambles," *American Economic Review*, 2004, *94* (4), 1207–1222.

Makarov, Igor and Guillaume Plantin, "Rewarding Trading Skills without Inducing Gambling," *Journal of Finance*, 2015, *70* (3), 925–962.

Morris, Stephen, "Political Correctness," *Journal of Political Economy*, 2001, *109* (2), 231–265.

Ottaviani, Marco and Peter Sørensen, "Professional Advice," *Journal of Economic Theory*, 2006, *126*, 120–142.

Peek, Joe and Eric S. Rosengren, "Unnatural Selection: Perverse Incentives and the Misallocation of Credit in Japan," *American Economic Review*, 2005, *95* (4), 1144–1166.

Pei, Harry Di, "Reputation with Strategic Information Disclosure," 2015. Working Paper.

Prat, Andrea, "The Wrong Kind of Transparency," *American Economic Review*, 2005, *95* (3), 862–877.

Prendergast, Canice and Lars Stole, "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 1996, *104* (6), 1105–1134.

Quah, John and Bruno Strulovici, "Comparative Statics, Informativeness, and the Interval Dominance Order," *Econometrica*, 2009, *77* (6), 1949–1992.

Rajan, Raghuram, "Why Bank Credit Policies Fluctuate: A Theory and Some Evidence," *Quarterly Journal of Economics*, 1994, *109*, 399–441.

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille, "Social Learning in One-Arm Bandit Problems," *Econometrica*, 2007, *75* (6), 1591–1611.

Scharfstein, David and Jeremy Stein, "Herd Behavior and Investment," *American Economic Review*, 1990, *80* (3), 465–479.

Schularick, Moritz and Alan M. Taylor, "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008," *American Economic Review*, 2012, *102* (2), 1029–1061.

Sekine, Toshitaka, Keiichiro Kobayashi, and Yumi Saita, "Forbearance Lending: The Case of Japanese Firms," *Monetary and Economic Studies*, 2003, *21* (2), 69–92.

Shmaya, Eran and Eilon Solan, "Equivalence between Random Stopping Times in Continuous Time," 2018. Working Paper.

Strebulaev, Ilya, Haoxiang Zhu, and Pavel Zryumov, "Dynamic Information Asymmetry, Financing, and Investment Decisions," 2016. Working Paper.

The Economist, "The Great Hole of China," 2014. October 18.

_ , "Deleveraging Delayed," 2015. October 24.

The Wall Street Journal, "Skepticism on China's Nonperforming Loans," 2013. December 3.

Thomas, Caroline, "Career Concerns and Policy Intransigence – A Dynamic Signaling Model," 2016. Working Paper.

Townsend, Robert M., "Optimal Contracts and Competitive Markets with Costly State Verification," *Journal of Economic Theory*, 1979, *21*, 265–293.

Watson, Joel, "Perfect Bayesian Equilibrium: General Definitions and Illustrations," January 2016. Working Paper.

Zwiebel, Jeffrey, "Corporate Conservatism and Relative Compensation," *Journal of Political Economy*, 1995, *103* (1), 1–25.

A Appendix

This Appendix contains the proofs of the results stated in Section 2; the proofs of the results in Section 3 are in the Internet Appendix. We introduce some notation: we denote by $\widehat{\mu}_t^0$ the market's belief that the project is good at time t conditional on the agent stopping at this time. Note that if the agent stops at t, the market's belief is $\widehat{\mu}_t^0$ for all $s \geq t$. As in the text, we denote by $\widehat{\mu}_t^1$ the market's time-t belief that the project is good conditional on the agent not having stopped by t. To simplify the exposition, in what follows we will disregard the almost surely qualification that would account for probability zero events.

A.1 Proof of Lemma 1

We prove the lemma by proving two claims:

Claim 1: Suppose there exists an equilibrium in which an agent who has not failed by $t > 0$ stops with strictly positive probability by t. Then the agent stops with certainty at a time $t' \leq t$.

Proof of Claim 1: Suppose the claim is not true. Then there exists an equilibrium in which the agent, not having failed by t, mixes over a time interval $[t', t' + dt]$ with dt arbitrarily close to zero and $t' + dt \leq t$. Since the market's belief is left-continuous, this agent must be indifferent between stopping and continuing at t'. Now note that for any fixed continuation strategy in which the agent continues at t', the expected continuation payoff is strictly larger for an agent who has not failed than for an agent who has failed by t'. On the other hand, the expected payoff from stopping at t' is the same for both types of the agent. Thus, if the agent who has not failed is indifferent at t', the agent who has failed has strict incentives to stop and stops with probability one by t'. Bayes' rule then yields that the market's belief that the agent has not failed by t' upon observing that he continues at t' is one, which implies $\widehat{\mu}^1_{t'} \geq \widehat{\mu}^0_{t'}$.

Given these market's beliefs, the agent has a continuation strategy upon continuing at t' that gives him an expected reputation payoff no smaller than $\widehat{\mu}^1_{t'} R/r$. Furthermore, note that the expected project payoff from any such continuation strategy is strictly positive for an agent who has not failed; this follows from $x - (1 - \mu_{t'})\lambda_B > 0$ by Assumption 1 and $\mu_{t'} > \mu_0$. Therefore, the total expected continuation payoff for an agent who has not failed by t' is strictly larger than $\widehat{\mu}^1_{t'} R/r$ if the agent continues with the project at t', and it is equal to $\widehat{\mu}^0_{t'} R/r \leq \widehat{\mu}^1_{t'} R/r$ if the agent stops at t'. It follows that this agent cannot be indifferent at t', yielding a contradiction.

Claim 2: There exists no equilibrium in which the agent stops with certainty at a time $t' > 0$ by which he has not failed.

Proof of Claim 2: The proof of this claim is analogous to that of Claim 1 above. The difference is that if an agent who has not failed stops with certainty at a time $t' > 0$, then by the arguments above, there is no history following which the agent continues with the project at t'. Hence, the market's belief upon observing that the agent continues at t' is off the equilibrium path. However, note that by our belief monotonicity refinement, we must have $\widehat{\mu}^1_{t'} \geq \widehat{\mu}^0_{t'}$. The rest of the proof is identical to that of Claim 1.

A.2 Proof of Proposition 1

We begin by showing existence under the different parameter conditions considered in the proposition. We then prove uniqueness.

Existence. Suppose first $R \leq -(x - \lambda_B)$. Consider an equilibrium in which the agent starts the project at time 0 and stops at time $t > 0$ if and only if he fails at t. Let the market's belief that the project is good at time $t \geq 0$ be μ_0 if the agent has not started the project, 0 if the agent has started and stopped by t, and μ_t if the agent has started and not stopped by t. Clearly, these beliefs are consistent (and on the equilibrium path), and by the arguments in the text the agent's stopping decision at each $t > 0$ is optimal. All is left to be shown is that it is optimal for the agent to start the project. The agent's expected payoff if he does not start is $\mu_0 R/r$. The agent's expected payoff if he starts and follows the proposed strategy is

$$\begin{aligned}
&\mu_0 \int_0^\infty e^{-rt} (x + \mu_t R) \, dt + (1 - \mu_0) \int_0^\infty e^{-(r+\lambda_B)t} (x + \mu_t R - \lambda_B) \, dt \\
&\geq \int_0^\infty e^{-rt} \left[x - (1 - \mu_0)\lambda_B + \mu_t R\right] dt \\
&> \frac{\mu_0 R}{r}.
\end{aligned}$$

The first inequality follows from the agent's stopping strategy at $t > 0$ being optimal given the market's beliefs (and thus yielding a larger expected payoff than a strategy of working forever). The second inequality follows from Assumption 1 and the fact that $\mu_t > \mu_0$ for all $t > 0$. Hence, we obtain that the payoff from starting the project and following the equilibrium strategy is larger than that from not starting the project.

Suppose next $R > -(x - \lambda_B)$. Consider an equilibrium in which the agent starts the project at time 0 and stops at time $t > 0$ if and only if he fails at t and $t < t^*$, where t^* is given by (7) if $R < -(x - \lambda_B)/\mu_0$ and $t^* = 0$ otherwise. Let the market's belief at time $t \geq 0$ be μ_0 if the agent has not started the project, 0 if the agent has started and stopped by t, μ_t if the agent has started and not stopped by t and $t \leq t^*$, and μ_{t^*} if the agent has started and not stopped by t and $t > t^*$. Clearly, on-the-equilibrium-path beliefs are consistent and off-the-equilibrium-path beliefs satisfy our belief monotonicity refinement. Moreover, by the arguments in the text, the agent's stopping decision at each $t > 0$ is optimal. Finally, an analogous argument to that above implies that it is optimal for the agent to start the project at time 0.

Uniqueness. We proceed by proving three claims.

Claim 1: An equilibrium in which the agent does not start the project at time 0 does not exist.
Proof of Claim 1: Suppose by contradiction that such an equilibrium exists. Recall that the agent has no private information at time 0; hence, the market's beliefs cannot change upon observing the agent's start decision. It follows that the agent has a continuation strategy upon

starting the project which gives him an expected reputation payoff no smaller than $\mu_0 R/r$. Moreover, as argued in the proof of Lemma 1, the agent's expected project payoff from any such continuation strategy is strictly positive. Therefore, the agent's expected payoff from starting the project at time 0 is strictly larger than $\mu_0 R/r$, which is his payoff from not starting. Contradiction.

Claim 2: In any equilibrium, if the agent fails at a time $t' > 0$, he either stops at t' with certainty or continues with the project at all $t \geq t'$.

Proof of Claim 2: Suppose the claim is not true. Then there exist $t' > 0$ and $dt > 0$ such that the agent stops with strictly positive probability on the interval $(t', t' + dt]$ when having failed by t'. By Lemma 1, the market's belief conditional on the agent not having stopped is weakly increasing. Moreover, for dt arbitrarily small, Bayes' rule implies that this belief is strictly increasing over $(t', t' + dt]$. This means that the agent's expected reputation payoff from continuing at $t \in (t', t' + dt]$ is strictly larger than that from continuing at t'. Additionally, for an agent who has failed by t', the expected project payoff is the same when continuing at any $t \geq t'$. Finally, by Lemma 1, the payoff from stopping at $t \in (t', t' + dt]$ is zero and thus no larger than the payoff from stopping at t'. It follows that if the agent is willing to stop over $(t', t' + dt]$, he has strict incentives to stop at t', yielding a contradiction.

Claim 3: In any equilibrium, the agent's strategy and the market's beliefs must be as described in the proposition.

Proof of Claim 3: The equilibrium strategy of the agent follows from Lemma 1, Claims 1 and 2 above, and the arguments in the text. Given the agent's strategy, Bayes' rule then pins down $\widehat{\mu}_t^1$ for all $t \geq 0$ and $\widehat{\mu}_t^0$ for $t < t^*$. We now show that the market's belief must satisfy $\widehat{\mu}_t^0 = 0$ for $t \geq t^*$. First, note that $\widehat{\mu}_t^0$ cannot jump at t^*: if it did, an agent who fails at $t < t^*$ arbitrarily close to t^* would prefer to continue at t and stop at t^*, yielding a contradiction. Since $\widehat{\mu}_t^0 = 0$ for $t < t^*$, it follows that $\widehat{\mu}_{t^*}^0 = 0$. Second, note that the agent is indifferent between stopping and continuing upon failing at t^*, and the market's belief conditional on the agent continuing is $\widehat{\mu}_t^1 = \mu_{t^*}$ for all $t \geq t^*$. Thus, if $\widehat{\mu}_{t'}^0 > 0$ for some $t' > t^*$, an agent who fails at t' strictly prefers to stop at that time, yielding another contradiction. The claim follows.

A.3 Proof of Proposition 2

Equations (2) and (7) yield $t^* = \frac{1}{\lambda_B} \log\left(\frac{-(1-\mu_0)(x-\lambda_B)}{\mu_0(R+x-\lambda_B)}\right)$. The derivative with respect to μ_0 is

$$\frac{dt^*}{d\mu_0} = -\frac{1}{\lambda_B \mu_0 (1-\mu_0)}. \tag{13}$$

The distortion relative to first best is

$$S_0^{FB} - S_0(t^*) = (1-\mu_0)(x-\lambda_B)e^{-(r+\lambda_B)t^*}\left(\frac{1}{r+\lambda_B} - \frac{1}{r}\right). \tag{14}$$

The derivative with respect to μ_0 is

$$\frac{d(S_0^{FB} - S_0(t^*))}{d\mu_0} = -(x-\lambda_B)e^{-(r+\lambda_B)t^*}\left(\frac{1}{r+\lambda_B} - \frac{1}{r}\right)\left[1+(1-\mu_0)(r+\lambda_B)\frac{dt^*}{d\mu_0}\right]. \tag{15}$$

Hence, the distortion is strictly increasing in μ_0 if and only if $\frac{dt^*}{d\mu_0} < -\frac{1}{(1-\mu_0)(r+\lambda_B)}$, or, equivalently, substituting with (13),

$$-\frac{1}{\lambda_B\mu_0(1-\mu_0)} < -\frac{1}{(1-\mu_0)(r+\lambda_B)}.$$

This inequality is always satisfied since $\mu_0 < 1$ and $r > 0$.

Finally, we show that equilibrium welfare $S_0(t^*)$ is strictly increasing in μ_0. Differentiating $S_0(t^*)$ given in (8) with respect to μ_0 yields

$$\frac{dS_0(t^*)}{d\mu_0} = \frac{x}{r} - (x-\lambda_B)\left(\frac{1-e^{-(\lambda_B+r)t^*}}{\lambda_B+r} + \frac{e^{-(\lambda_B+r)t^*}}{r}\right) - \frac{\lambda_B(1-\mu_0)(x-\lambda_B)e^{-(\lambda_B+r)t^*}}{r}\frac{dt^*}{d\mu_0}.$$

Substituting with (13) and rearranging terms,

$$\frac{dS_0(t^*)}{d\mu_0} = \frac{x}{r} - (x-\lambda_B)\frac{\mu_0 r - e^{-(\lambda_B+r)t^*}(\lambda_B(1-\mu_0)+r)}{(\lambda_B+r)\mu_0 r}. \tag{16}$$

Suppose by contradiction that $S_0(t^*)$ is decreasing in μ_0. Given (16), this requires

$$x \le (x-\lambda_B)\frac{\mu_0 r - e^{-(\lambda_B+r)t^*}(\lambda_B(1-\mu_0)+r)}{(\lambda_B+r)\mu_0}.$$

Recall that $x - \lambda_B < 0$. It follows that this condition can hold only if the left-hand side is smaller than the right-hand side when $t^* = 0$, i.e. only if

$$x \le -(x-\lambda_B)\frac{(\lambda_B+r)(1-\mu_0)}{(\lambda_B+r)\mu_0}.$$

However, this requires $x - (1-\mu_0)\lambda_B \le 0$, which violates Assumption 1. Hence, $S_0(t^*)$ is strictly increasing in μ_0.

A.4 Proof of Proposition 3

Consider the first part of the proposition. Recall that the distortion relative to first best is given by (14). Proposition 2 shows that this distortion is increasing in μ_0. Moreover, note that first-best welfare is unchanged with the signal given that starting the project is efficient for all of its realizations. (This follows from the fact that first-best welfare, given by (4), is linear in μ_0.) To prove the first part of the proposition, it is therefore sufficient to show that the distortion is convex in μ_0. Differentiating (15) yields

$$\begin{aligned}\frac{d^2(S_0^{FB}-S_0(t^*))}{d\mu_0^2} &= 2(x-\lambda_B)e^{-(r+\lambda_B)t^*}\left(\frac{1}{r+\lambda_B}-\frac{1}{r}\right)(r+\lambda_B)\frac{dt^*}{d\mu_0} \\ &\quad +(1-\mu_0)(x-\lambda_B)e^{-(r+\lambda_B)t^*}\left(\frac{1}{r+\lambda_B}-\frac{1}{r}\right)(r+\lambda_B)^2\left(\frac{dt^*}{d\mu_0}\right)^2 \\ &\quad -(1-\mu_0)(x-\lambda_B)e^{-(r+\lambda_B)t^*}\left(\frac{1}{r+\lambda_B}-\frac{1}{r}\right)(r+\lambda_B)\frac{d^2t^*}{d\mu_0^2}.\end{aligned}$$

Hence,

$$\frac{d^2(S_0^{FB}-S_0(t^*))}{d\mu_0^2}>0 \iff 2\frac{dt^*}{d\mu_0}+(1-\mu_0)(r+\lambda_B)\left(\frac{dt^*}{d\mu_0}\right)^2-(1-\mu_0)\frac{d^2t^*}{d\mu_0^2}>0. \quad (17)$$

Differentiating (13) yields

$$\frac{d^2t^*}{d\mu_0^2}=\frac{\lambda_B(1-2\mu_0)}{[\lambda_B\mu_0(1-\mu_0)]^2}. \quad (18)$$

Substituting (13) and (18) in (17) and rearranging terms yields that $\frac{d^2(S_0^{FB}-S_0(t^*))}{d\mu_0^2}>0$ if and only if

$$-\frac{2}{\lambda_B\mu_0(1-\mu_0)}+\frac{(1-\mu_0)(r+\lambda_B)}{[\lambda_B\mu_0(1-\mu_0)]^2}-\frac{(1-\mu_0)\lambda_B(1-2\mu_0)}{[\lambda_B\mu_0(1-\mu_0)]^2}>0,$$

which reduces to $(1-\mu_0)r>0$ and is thus always satisfied.

Consider next the second part of the proposition. If the signal reveals a good project, it is efficient to continue forever and there are no distortions. If the signal reveals a bad project, it is efficient to not start the project. Since the agent's actions reveal no information, the career-concerned agent has no incentives to start, so again there are no distortions. It follows that a fully informative signal eliminates distortions relative to first best. Since first-best welfare increases with a fully informative signal (as losses are avoided when the project is bad), equilibrium welfare increases.

B Internet Appendix

This Internet Appendix contains the proofs of the results stated in Section 3 of the paper.

B.1 Proof of Proposition 4

We begin by examining the agent's indifference conditions. We then describe the equilibrium construction and show existence under the different parameter conditions considered in the proposition. Finally, we prove uniqueness.

Indifference conditions. Consider the agent's indifference condition (12) for each $t \in [\underline{t}, \overline{t}]$. Writing $\mu_{t+dt} = \mu_t + dt\dot{\mu}_t = \mu_t - \mu_t(1-\mu_t)\lambda_G dt$, dividing by dt both sides of the equation, and taking dt to zero, this condition can be rewritten as

$$R\left(\widehat{\mu}_t^1 - \mu_t\right) + x + \mu_t\lambda_G\left(V_t - \frac{R}{r}\right) = 0, \tag{19}$$

where $V_t = v + R\int_t^\infty e^{-r(s-t)}\widehat{\mu}_s^1 ds$. Since (19) holds at each $t \in [\underline{t}, \overline{t}]$, the agent must be indifferent between stopping and continuing until $\overline{t}$ absent success. For $\overline{t} < \infty$, the agent stops at $\overline{t}$ if he has not succeeded and continues forever otherwise. Thus, for each $t \in [\underline{t}, \overline{t}]$,

$$\begin{aligned}
\mu_t R &= (x+\mu_t\lambda_G)\left(1-e^{-r(\overline{t}-t)}\right) + rR\int_t^{\overline{t}} e^{-r(s-t)}\widehat{\mu}_s^1 ds \\
&\quad + e^{-r(\overline{t}-t)}\left\{\mu_t\left(1-e^{-\lambda_G(\overline{t}-t)}\right)(x+\lambda_G+R) + \left[\mu_t e^{-\lambda_G(\overline{t}-t)} + 1 - \mu_t\right]\mu_{\overline{t}}R\right\}.
\end{aligned} \tag{20}$$

Differentiating this condition yields

$$\begin{aligned}
\dot{\mu}_t R &= \dot{\mu}_t\lambda_G\left(1-e^{-r(\overline{t}-t)}\right) - r\left(x+\mu_t\lambda_G\right)e^{-r(\overline{t}-t)} - rR\left[\widehat{\mu}_t^1 - r\int_t^{\overline{t}} e^{-r(s-t)}\widehat{\mu}_s^1 ds\right] \\
&\quad + re^{-r(\overline{t}-t)}\left[\mu_t\left(1-e^{-\lambda_G(\overline{t}-t)}\right)(x+\lambda_G+R) + \left(\mu_t e^{-\lambda_G(\overline{t}-t)} + 1 - \mu_t\right)\mu_{\overline{t}}R\right] \\
&\quad + e^{-r(\overline{t}-t)}\left\{\begin{array}{c} \left[\dot{\mu}_t\left(1-e^{-\lambda_G(\overline{t}-t)}\right) - \mu_t\lambda_G e^{-\lambda_G(\overline{t}-t)}\right](x+\lambda_G+R) \\ + \left[\mu_t\lambda_G e^{-\lambda_G(\overline{t}-t)} - \dot{\mu}_t\left(1-e^{-\lambda_G(\overline{t}-t)}\right)\right]\mu_{\overline{t}}R \end{array}\right\}.
\end{aligned}$$

Substituting with equation (20), writing $\dot{\mu}_t = -\mu_t(1-\mu_t)\lambda_G$, and rearranging terms yields

$$\widehat{\mu}_t^1 = \frac{1}{rR}\left\{\begin{array}{c} r\left[-x-\mu_t(\lambda_G - R)\right] - \mu_t(1-\mu_t)\lambda_G\left[\lambda_G\left(1-e^{-r(\overline{t}-t)}\right) - R\right] \\ -e^{-r(\overline{t}-t)}\mu_t\lambda_G\left[1-\mu_t\left(1-e^{-\lambda_G(\overline{t}-t)}\right)\right]\left[x+\lambda_G+R(1-\mu_{\overline{t}})\right] \end{array}\right\}. \tag{21}$$

Existence under $-\mathbf{x} < \mathbf{R} < -\mathbf{x}/\boldsymbol{\mu}_0$. We construct an equilibrium as described in the proposition in which $\underline{t} \in (t^{FB}, \infty)$ and $\bar{t} = \infty$. Equations (20) and (21) reduce to

$$\mu_t R = x + \mu_t \lambda_G + rR \int_t^\infty e^{-r(s-t)} \widehat{\mu}_s^1 ds, \tag{22}$$

$$\widehat{\mu}_t^1 = \frac{1}{rR} \Big\{ r\left[-x - \mu_t(\lambda_G - R)\right] - \mu_t(1-\mu_t)\lambda_G\left(\lambda_G - R\right) \Big\}, \tag{23}$$

for each $t \geq \underline{t}$. Consider an equilibrium in which the agent starts the project at time 0, he continues with probability one at time t if he has succeeded by t or $t < \underline{t}$, and otherwise he follows a random stopping policy from time $\underline{t}$ on such that the market's belief $\widehat{\mu}_t^1$ satisfies (23). Let the market's belief at t be μ_0 if the agent has not started the project or has started and not stopped by t and $t < \underline{t}$, $\widehat{\mu}_t^1$ satisfying (23) if the agent has started and not stopped by t and $t \geq \underline{t}$, and μ_s if the agent has started and stopped at $s \in (0, t]$. (Note that the belief upon observing that the agent stops at a time $t < \underline{t}$ is off the equilibrium path; we show existence when this belief satisfies $\widehat{\mu}_t^0 = \mu_t$ for all such t.)

The market's on-the-equilibrium-path beliefs are consistent and the off-the-equilibrium-path beliefs satisfy our belief monotonicity refinement. We now show that given the market's beliefs, the agent's stopping plan is optimal. By construction, the agent is indifferent and thus willing to follow a mixed strategy for $t \geq \underline{t}$ in the absence of success. Note also that an agent who has succeeded strictly prefers to continue with the project at time t if an agent who has not succeeded weakly prefers to continue. Thus, all is left to be shown is that the agent has incentives to start the project at time 0 and to continue with the project at $t < \underline{t}$ in the absence of success.

For the start decision, note that the agent's payoff if he does not start is $\mu_0 R/r$. By the martingale property of beliefs, this is also the agent's expected reputation payoff if he starts the project and follows the equilibrium strategy. Moreover, by Assumption 1, the agent receives a strictly positive expected project payoff from any strategy that starts the project and continues forever upon success. Hence, the agent strictly prefers to start and follow the equilibrium strategy compared to not starting the project.

To show that it is optimal for the agent to continue at $t < \underline{t}$ absent success, it is sufficient to show that the left-hand side of (22) is smaller than its right-hand side for $t < \underline{t}$, or, equivalently,

$$\Psi_t := x + \mu_t(\lambda_G - R) + rR \int_t^\infty e^{-r(s-t)} \widehat{\mu}_s^1 ds \geq 0. \tag{24}$$

Suppose by contradiction that $\Psi_t < 0$ for some $t < \underline{t}$. Note that by the previous claims, $\Psi_0 > 0$, and by definition of $\underline{t}$, $\Psi_{\underline{t}} = 0$. Therefore, if $\Psi_t < 0$ for $t < \underline{t}$, there exist $t' < t'' < \underline{t}$

such that $\Psi_{t'} = 0$, $\dot{\Psi}_{t'} < 0$, $\Psi_{t''} < 0$, and $\dot{\Psi}_{t''} = 0$. Differentiating (24) yields that for $t < \underline{t}$,

$$\dot{\Psi}_t = -\mu_t(1-\mu_t)\lambda_G\left(\lambda_G - R\right) - rR\mu_0 + r^2 R \int_t^{\infty} e^{-r(s-t)}\widehat{\mu}_s^1 ds. \tag{25}$$

Using (24) and (25), note that $\Psi_{t'} = 0$, $\dot{\Psi}_{t'} < 0$, $\Psi_{t''} < 0$, and $\dot{\Psi}_{t''} = 0$ imply

$$-\mu_{t'}(1-\mu_{t'})\lambda_G\left(\lambda_G - R\right) - rR\mu_0 - r[x + \mu_{t'}(\lambda_G - R)] \quad < \quad 0, \tag{26}$$
$$-\mu_{t''}(1-\mu_{t''})\lambda_G\left(\lambda_G - R\right) - rR\mu_0 - r[x + \mu_{t''}(\lambda_G - R)] \quad > \quad 0. \tag{27}$$

However, given Assumption 1 and $-x < R < -x/\mu_0$, there is a unique value $\mu_{\underline{t}} \in (0, \mu_0)$ that solves (23) at $\underline{t}$ given $\widehat{\mu}_{\underline{t}}^1 = \mu_0$, and this value is given by

$$\mu_{\underline{t}} = \frac{\lambda_G + r - \sqrt{\frac{\lambda_G^2(2r+\lambda_G-R)+\lambda_G r[4(\mu_0 R+x)+r-2R]-r^2 R}{\lambda_G - R}}}{2\lambda_G}. \tag{28}$$

Hence, we cannot have $\mu_{t'}, \mu_{t''} \in (\mu_{\underline{t}}, \mu_0)$. Contradiction.

Finally, we show that the equilibrium threshold time $\underline{t}$ satisfies $\underline{t} > t^{FB}$. Note that $\underline{t}$ is uniquely pinned down by (9) and (28). To prove $\underline{t} > t^{FB}$, we verify that $\mu_{\underline{t}} < \mu^{FB}$. The latter inequality is immediate from comparing $\mu_{\underline{t}}$ given in (28) and μ^{FB} given in (11), taking into account that Assumption 1 and $R < -x/\mu_0$ imply $\lambda_G > R$.

Existence under $\mathbf{R < -x}$. We consider an equilibrium analogous to that above, except that the agent now stops with certainty by a finite time $\bar{t}$ if he has not succeeded by then. Consistently, we specify beliefs for the market as those above but with $\widehat{\mu}_t^1 = 1$ for all $t \geq \bar{t}$. Given this, the agent's indifference condition (19) at $\bar{t}$ implies

$$R(1-\mu_{\bar{t}}) + x + \mu_{\bar{t}}\lambda_G v = 0, \tag{29}$$

and hence[25]

$$\mu_{\bar{t}} = \frac{-(x+R)r}{r(\lambda_G - R) + \lambda_G(\lambda_G + x)}. \tag{30}$$

The market's on-the-equilibrium-path beliefs are consistent and the off-the-equilibrium-path beliefs satisfy our belief monotonicity refinement. (Note that the belief upon observing that the agent stops at a time $t < \underline{t}$ or $t > \bar{t}$ is off the equilibrium path; we show existence when this belief satisfies $\widehat{\mu}_t^0 = \mu_t$ for all such t.) We now verify that given the market's beliefs, the

[25]Equation (30) can equivalently be derived from (21) at $\bar{t}$.

agent's stopping plan is optimal. Since the construction is analogous to that in the previous case, all we need to verify is that the agent has no incentives to continue beyond $\bar{t}$ absent success. Note that μ_t is decreasing over time and $\lambda_G > R$ (by Assumption 1 and $R < -x$). Thus, the left-hand side of (29) evaluated at $t > \bar{t}$ instead of $\bar{t}$ is strictly negative, implying that the agent has indeed strict incentives not to continue beyond time $\bar{t}$.

Finally, we show that the above conditions pin down $\mu_{\underline{t}}$ and imply $\underline{t} > t^{FB}$. Since (30) uniquely pins down $\mu_{\bar{t}}$ and (together with (9)) $\bar{t}$, these can be substituted into equation (21) at $t = \underline{t}$ (with $\widehat{\mu}^1_{\underline{t}} = \mu_0$) to solve for $\mu_{\underline{t}}$ and $\underline{t}$. Combining (9), (21), and (30) yields

$$\begin{aligned} 0 &= r\left[x + \mu_{\underline{t}}\lambda_G + R(\mu_0 - \mu_{\underline{t}})\right] + \mu_{\underline{t}}(1 - \mu_{\underline{t}})\lambda_G(\lambda_G - R) \\ &\quad + \mu_{\underline{t}}(1 - \mu_{\underline{t}})\lambda_G\left(\frac{-(x+R)r}{(\lambda_G + x)(r + \lambda_G)}\frac{1 - \mu_{\underline{t}}}{\mu_{\underline{t}}}\right)^{\frac{r}{\lambda_G}}\frac{\lambda_G(R + x)}{r + \lambda_G}. \end{aligned} \tag{31}$$

To show that $\underline{t} > t^{FB}$, or equivalently $\mu_{\underline{t}} < \mu^{FB}$, note that $\mu_{\underline{t}}$ is continuous in R for all $R \geq 0$, and $\mu_{\underline{t}} \to \mu^{FB}$ as $R \to 0$. Moreover, in the limit as $R \to -x$, $\mu_{\underline{t}}$ coincides with (28), which implies $\mu_{\underline{t}} < \mu^{FB}$ in this limit. Thus, suppose by contradiction that $\mu_{\underline{t}} \geq \mu^{FB}$ for some $R \in (0, -x)$. Then there must exist $R' > 0$ such that $\mu_{\underline{t}}(R') = \mu^{FB}$. More precisely, let $g(R)$ be the right-hand-side of (31) when $\mu_{\underline{t}} = \mu^{FB}$, as a function of R. Algebraic manipulations yield

$$g(R) := \frac{r\begin{pmatrix} -x^2\left(\lambda_G^2 + \lambda_G(\mu_0 + 1)R + 2rR\right) - Rx(\lambda_G + r)(2\lambda_G\mu_0 + \lambda_G + r) \\ -\lambda_G\mu_0 R(\lambda_G + r)^2 + \lambda_G x(\lambda_G + x)(R + x)\left(\frac{R+x}{x}\right)^{r/\lambda_G} - \lambda_G x^3 \end{pmatrix}}{\lambda_G(\lambda_G + r + x)^2}.$$

As noted, $g(0) = 0$, and by the contradiction assumption, there exists $R' > 0$ such that $g(R') = 0$. Furthermore, there must then exist $0 < R'' < R'$ such that $g'(R'') = 0$, where differentiating $g(R)$ gives

$$g'(R) = \frac{r\begin{pmatrix} -x^2(\lambda_G\mu_0 + \lambda_G + 2r) - x(\lambda_G + r)(2\lambda_G\mu_0 + \lambda_G + r) \\ -\lambda_G\mu_0(\lambda_G + r)^2 + x(\lambda_G + r)(\lambda_G + x)\left(\frac{R+x}{x}\right)^{r/\lambda_G} \end{pmatrix}}{\lambda_G(\lambda_G + r + x)^2}.$$

Note that given a set of parameters $\{\mu_0, \lambda_G, x, r\}$, there is a unique value R'' for which $g'(R'') = 0$, and it must satisfy $g(R'') < 0$. Therefore, it must be that $g'(R') > 0$. However, one can verify that $g(R') = 0$ implies $g'(R') < 0$, yielding a contradiction.

Existence under $\mathbf{R > -x/\mu_0}$. Consider an equilibrium in which the agent starts the project and never stops. Let the market's beliefs be $\widehat{\mu}_t^1 = \mu_0$ and $\widehat{\mu}_t^0 = \mu_t$ for all $t \geq 0$. (Note that the belief upon observing that the agent stops at a time $t > 0$ is off the equilibrium path; we show existence when this belief satisfies $\widehat{\mu}_t^0 = \mu_t$ for all such t.) The market's on-the-equilibrium-path beliefs are consistent and the off-the-equilibrium-path beliefs satisfy our belief monotonicity refinement. We now show that given these beliefs, the agent always has strict incentives to continue. That is, for all $t \geq 0$,

$$\mu_t R < x + \mu_t \lambda_G + \mu_0 R.$$

Given $R > -x/\mu_0$, it is immediate that this condition always holds if $R \leq \lambda_G$. Suppose instead that $R > \lambda_G$. Since μ_t is decreasing over time, it suffices to show in this case that this condition holds at time 0. By Assumption 1, this is indeed true.

Uniqueness. We show that the equilibrium is unique up to off-the-equilibrium-path beliefs by proving the following claims.

Claim 1: An equilibrium in which the agent does not start the project at time 0 does not exist.
Proof of Claim 1: The proof of this claim is analogous to that of Claim 1 in the proof of Proposition 1 and thus omitted.

Claim 2: Suppose there exists an equilibrium in which an agent who succeeds at $t > 0$ stops with strictly positive probability after t. Then the agent stops with certainty at a time $t' \geq t$.
Proof of Claim 2: The proof of this claim is analogous to that of Claim 1 in the proof of Lemma 1 and thus omitted.

Claim 3: There exists no equilibrium in which the agent stops with certainty at a time $t' > 0$ by which he has succeeded.
Proof of Claim 3: The proof of this claim is analogous to that of Claim 2 in the proof of Lemma 1 and thus omitted.

Claim 4: If $R > -x/\mu_0$, an equilibrium in which the agent stops with strictly positive probability does not exist.
Proof of Claim 4: Our proof of existence under $R > -x/\mu_0$ shows that the agent has strict incentives to continue at a time t if the market's beliefs satisfy $\widehat{\mu}_t^0 = \mu_t$ and $\widehat{\mu}_s^1 = \mu_0$ for all $s \geq t$. It follows that the agent also has strict incentives to continue at t if $\widehat{\mu}_t^0 \leq \mu_t$ and $\widehat{\mu}_s^1 \geq \mu_0$ for all $s \geq t$. By Claims 2 and 3 above, $\widehat{\mu}_s^1 \geq \mu_0$ for all $s \geq 0$. Hence, the agent can only be willing to stop at a time t if $\widehat{\mu}_t^0 > \mu_t$. However, by Claims 2 and 3 such a belief would not be consistent. The claim follows.

Claim 5: If $R < -x/\mu_0$, an equilibrium in which the agent never stops does not exist.

Proof of Claim 5: Suppose by contradiction that such an equilibrium exists. Then the market's belief conditional on the agent continuing is $\widehat{\mu}_t^1 = \mu_0$ for all $t \geq 0$, and the agent must be willing to continue rather than stop at all times. However, since the agent's payoff from stopping is weakly positive and $\mu_t \to 0$ as $t \to \infty$, this requires $\mu_0 R + x \geq 0$. Contradiction.

Claim 6: If $R < -x$, an equilibrium in which the agent continues with the project with strictly positive probability absent success in the limit as $t \to \infty$ does not exist.

Proof of Claim 6: Suppose by contradiction that such an equilibrium exists. The agent must be willing to continue rather than stop absent success in the limit as $t \to \infty$. Since the agent's payoff from stopping is weakly positive and $\mu_t \to 0$ as $t \to \infty$, this requires that for some $\widehat{\mu}_\infty^1 \leq 1$, $\widehat{\mu}_\infty^1 R + x \geq 0$. This inequality however cannot be satisfied when $R < -x$.

Claim 7: In any equilibrium, the market's belief conditional on the agent not having stopped, $\widehat{\mu}_t^1$, must be continuous.

Proof of Claim 7: Suppose by contradiction that an equilibrium in which $\widehat{\mu}_t^1$ is discontinuous exists. Let $\hat{t}$ be the earliest time at which this belief jumps. By Claims 2 and 3, $\widehat{\mu}_t^1$ is weakly increasing and can only jump up. Suppose the belief jumps at $\hat{t}$ from $\widehat{\mu}_{\hat{t}-}^1 = \widehat{\mu}^{1-}$ to $\widehat{\mu}_{\hat{t}+}^1 = \widehat{\mu}^{1+} > \widehat{\mu}^{1-}$. This requires the agent stopping with strictly positive probability, and by consistency of beliefs and Claims 2 and 3, the market's belief must satisfy $\widehat{\mu}_{\hat{t}}^0 = \mu_{\hat{t}}$. Observe also that the market's belief satisfies $\widehat{\mu}_t^0 \geq \mu_t$ for all $t > 0$, on and off the equilibrium path. (That is, the most pessimistic belief at t corresponds to no success having arrived by t.)

Consider now the agent's incentives. In the absence of success, the agent must be willing to stop at $\hat{t}$ rather than continue for an arbitrarily small amount of time dt and stop at $\hat{t} + dt$ if no success is obtained over $[\hat{t}, \hat{t} + dt]$. Following similar steps to those used to derive (19), taking dt to 0, this condition is

$$R\left(\widehat{\mu}^{1+} - \mu_{\hat{t}}\right) + x + \mu_{\hat{t}} \lambda_G \left(V_{\hat{t}} - \frac{R}{r}\right) \leq 0. \tag{32}$$

In the absence of success, the agent must also be willing to continue working over $[\hat{t} - dt, \hat{t}]$ and stop at $\hat{t}$ if no success is obtained over $[\hat{t} - dt, \hat{t}]$ rather than stop at $\hat{t} - dt$. This condition can be written as

$$R\left(\widehat{\mu}^{1-} - \widehat{\mu}_{\hat{t}-}^0\right) + x + \mu_{\hat{t}} \lambda_G \left(V_{\hat{t}} - \frac{R}{r}\right) \geq 0. \tag{33}$$

However, $\widehat{\mu}^{1+} > \widehat{\mu}^{1-}$ and $\widehat{\mu}_{\hat{t}-}^0 \geq \mu_{\hat{t}}$ imply that (32) and (33) cannot be simultaneously satisfied. Contradiction.

Claim 8: Suppose there exists an equilibrium in which, absent success, the agent stops with *strictly positive probability by a time $\bar{t} < \infty$ and with zero probability at all times $t > \bar{t}$. Then $\widehat{\mu}_t^1 = 1$ for all $t \geq \bar{t}$.*

Proof of Claim 8: Suppose the claim is not true. Then there exists an equilibrium in which the agent quits with strictly positive probability absent success, he ceases quitting at a time $\bar{t} < \infty$, and the market's belief satisfies $\widehat{\mu}_t^1 < 1$ for some $t \geq \bar{t}$. Since the agent continues with certainty after $\bar{t}$ if he has not stopped by then, the market's belief $\widehat{\mu}_t^1$ must be constant at some value, call it $\overline{\mu}$, for all $t \geq \bar{t}$. The agent's indifference condition (19) at $\bar{t}$ yields

$$R\overline{\mu} + x + \mu_{\bar{t}} \left(\lambda_G - R + \lambda_G \frac{\lambda_G + x}{r} - \lambda_G (1 - \overline{\mu}) \frac{R}{r} \right) = 0. \tag{34}$$

Note that $\overline{\mu} < 1$ requires that an agent who has not succeeded by $\bar{t}$ be willing to continue beyond this time. Since μ_t is decreasing over time, equation (34) implies that the agent is willing to continue after $\bar{t}$ absent success if and only if the expression in parenthesis is negative. That is, rearranging terms, the equilibrium requires

$$\lambda_G - R + \frac{\lambda_G}{r} (\lambda_G - R + x + \overline{\mu} R) \leq 0.$$

By Claim 4, the agent stopping with strictly positive probability in equilibrium requires $R < -x/\mu_0$. Together with Assumption 1, this implies $\lambda_G > R$. Hence, the above inequality can hold only if $x + \overline{\mu} R < 0$. However, if the parenthesis in (34) is negative and $x + \overline{\mu} R < 0$, (34) cannot hold. Contradiction.

Claim 9: Suppose $R > -x$. There is no equilibrium in which, absent success, the agent stops with strictly positive probability by a time $\bar{t} < \infty$ and with zero probability at all times $t > \bar{t}$.

Proof of Claim 9: Suppose by contradiction that such an equilibrium exists. By Claim 4, the agent stopping with strictly positive probability in equilibrium requires $R < -x/\mu_0$. Moreover, as shown in Claim 8, if the agent's quitting ceases by a time $\bar{t} < \infty$, the market's belief must be $\widehat{\mu}_t^1 = 1$ for all $t \geq \bar{t}$, and hence equation (30) must hold at $\bar{t}$. However, if $R > -x$, this equation yields $\mu_{\bar{t}} < 0$ (recall $\lambda_G > R$ by Assumption 1 and $R < -x/\mu_0$), a contradiction.

Claim 10: Suppose there exists an equilibrium in which, absent success, the agent stops with strictly positive probability over $[t_1, t_2]$ and with zero probability over $[t_2, t_3]$, for some $0 < t_1 < t_2 < t_3$. Then the agent stops with zero probability at all $t \geq t_2$.

Proof of Claim 10: Suppose the claim is not true. Then there exists an equilibrium in which, absent success, the agent stops with strictly positive probability over $[t_1, t_2]$, with zero probability over $[t_2, t_3]$, and with strictly positive probability over $[t_3, t_4]$, for some $0 < t_1 < t_2 < t_3 < t_4$. Let $\bar{t} > 0$ be such that either $\bar{t} < \infty$ and the probability of stopping absent success is zero at all $t > \bar{t}$, or $\bar{t} = \infty$. By construction, around both times t_2 and t_3, the agent must be indifferent between stopping and continuing until $\bar{t}$ absent success. It follows that equation (21) must hold at t_2 and t_3, where note that if $\bar{t} < \infty$, then (30) uniquely pins down $\mu_{\bar{t}}$ and (together with (9)) $\bar{t}$. However, since the agent stops with zero probability between t_2 and t_3, we must have $\widehat{\mu}^1_{t_2} = \widehat{\mu}^1_{t_3} \geq \mu_0$. Given $\mu_{t_2} > \mu_{t_3}$, (21) cannot simultaneously hold at t_2 and t_3, yielding a contradiction.

Claim 11: Up to off-the-equilibrium-path beliefs, the equilibrium is unique.
Proof of Claim 11: This follows from the claims above and the fact that the solutions for the threshold times and $\widehat{\mu}^1_t$ shown in the proofs of existence above are unique.

B.2 Proof of Proposition 5

Preliminaries. Consider parameters with $-x/\mu_0 > R > -x$. As shown in Proposition 4, the equilibrium features $\underline{t} \in (t^{FB}, \infty)$ and $\bar{t} = \infty$, where $\underline{t}$ is given by (9) and (28), and equations (22) and (23) hold at each $t \geq \underline{t}$. The market's belief conditional on the agent not having stopped is $\widehat{\mu}^1_t = \mu_0$ for $t < \underline{t}$, and, by equation (23), this belief can be written as a function of μ_t independent of μ_0 for $t \geq \underline{t}$. Note also that the posterior belief at which the agent starts quitting, $\mu_{\underline{t}}$, is decreasing in μ_0; this can be verified using (28).

The equilibrium therefore implies that, for any $\delta \geq 0$, $\widehat{\mu}^1_{t^{FB}(\mu_0)+\delta}$ is increasing and convex in μ_0. To see this, fix a prior μ_0' and a posterior belief $\mu' \leq \mu^{FB}$. For any prior $\mu_0 \geq \mu_0'$, consider the market's belief that corresponds to such a posterior, $\widehat{\mu}^1(\mu', \mu_0)$. The construction implies that if $\mu' > \mu_{\underline{t}(\mu_0')}$, $\widehat{\mu}^1(\mu', \mu_0)$ increases one-for-one as μ_0 increases from μ_0'. If $\mu' < \mu_{\underline{t}(\mu_0')}$, then as μ_0 increases from μ_0', the belief $\widehat{\mu}^1(\mu', \mu_0)$ is invariant to μ_0 up to $\mu_0 = \widehat{\mu}^1(\mu', \mu_0')$, and increases one-for-one with μ_0 for $\mu_0 > \widehat{\mu}^1(\mu', \mu_0')$. Figure 5 provides an illustration.

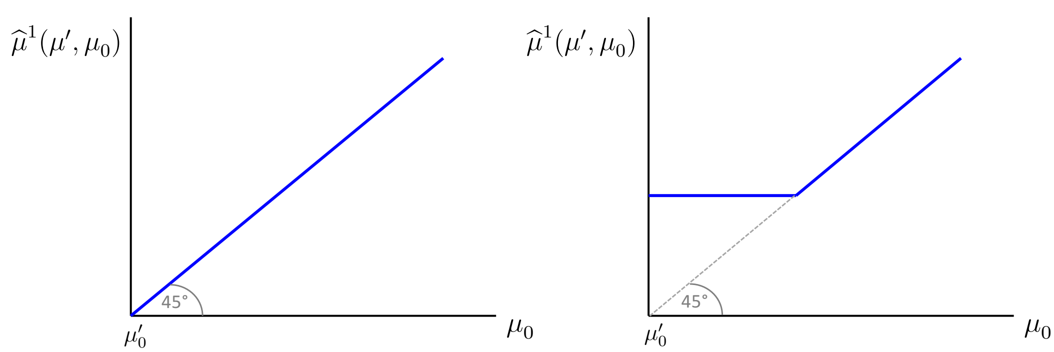

Figure 5: Market's belief conditional on the agent not having stopped, as a function of μ_0 for a fixed posterior belief μ'. The left graph corresponds to $\mu' > \mu_{\underline{t}(\mu_0')}$, where μ_0' is the lowest prior considered in the figure. The right graph corresponds to $\mu' < \mu_{\underline{t}(\mu_0')}$.

Hence, for any $\delta \geq 0$, we have

$$\frac{\partial \widehat{\mu}^1_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \geq 0, \quad \frac{\partial^2 \widehat{\mu}^1_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \geq 0. \tag{35}$$

Let $\eta_{t^{FB}(\mu_0)+\delta}$ denote the probability that the agent has not succeeded by time $t^{FB}(\mu_0)+\delta$ conditional on the agent continuing until this time:

$$\eta_{t^{FB}(\mu_0)+\delta} = \frac{\Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}.$$

The market's belief satisfies

$$\widehat{\mu}^1_{t^{FB}(\mu_0)+\delta} = 1 - \eta_{t^{FB}(\mu_0)+\delta} + \eta_{t^{FB}(\mu_0)+\delta}\ \mu_{t^{FB}(\mu_0)+\delta}. \tag{36}$$

Since $\mu_{t^{FB}(\mu_0)+\delta}$ is independent of μ_0, differentiating (36) yields

$$\begin{aligned} -\frac{\partial \widehat{\mu}^1_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} &= \frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}(1-\mu_{t^{FB}(\mu_0)+\delta}), \\ -\frac{\partial^2 \widehat{\mu}^1_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} &= \frac{\partial^2 \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}(1-\mu_{t^{FB}(\mu_0)+\delta}). \end{aligned}$$

Combining this with (35), we obtain that for any $\delta \geq 0$,

$$\frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \leq 0, \quad \frac{\partial^2 \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq 0. \tag{37}$$

Effects of μ_0 on welfare. We show that welfare is increasing in μ_0. Since first-best welfare is increasing in μ_0, it suffices to show that flow welfare at any time $t > t^{FB}(\mu_0)$ is increasing in μ_0. For any $\delta > 0$, welfare at time $t^{FB}(\mu_0)+\delta$ is

$$\begin{aligned} &\Pr(\text{succeeded \& cont})_{t^{FB}(\mu_0)+\delta}\ (\lambda_G + x) \\ &+ \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}\ (\mu_{t^{FB}(\mu_0)+\delta}\lambda_G + x). \end{aligned} \tag{38}$$

Suppose for the purpose of contradiction that for some $\delta > 0$, (38) is decreasing in μ_0, that is (using the fact that $\mu_{t^{FB}(\mu_0)+\delta}$ is independent of μ_0),

$$\frac{\partial \Pr(\text{succeeded \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} (\lambda_G + x) + \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} (\mu_{t^{FB}(\mu_0)+\delta}\lambda_G + x) < 0. \tag{39}$$

We can rewrite (39) as

$$\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} (\lambda_G + x) < \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \lambda_G(1 - \mu_{t^{FB}(\mu_0)+\delta}). \tag{40}$$

Note that the derivative on the left-hand side is positive.[26] Moreover, Assumption 1 implies $(\lambda_G + x)/\lambda_G > 1 - \mu_0$. Hence, (40) implies

$$\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} (1 - \mu_0) < \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} (1 - \mu_{t^{FB}(\mu_0)+\delta}).$$

Substituting with $1 - \mu_{t^{FB}(\mu_0)+\delta} = \frac{1-\mu_0}{\mu_0 e^{-\lambda_G(t^{FB}(\mu_0)+\delta)} + 1 - \mu_0}$, this can be rewritten as

$$\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \left(\mu_0 e^{-\lambda_G(t^{FB}(\mu_0)+\delta)} + 1 - \mu_0\right) < \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}. \tag{41}$$

Finally, note that $\eta_{t^{FB}(\mu_0)+\delta} \le \mu_0 e^{-\lambda_G(t^{FB}(\mu_0)+\delta)} + 1 - \mu_0$, as an agent who has succeeded does not stop. Thus, (41) implies

$$\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \eta_{t^{FB}(\mu_0)+\delta} < \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}. \tag{42}$$

We now show that (42) contradicts (37), namely the fact that $\eta_{t^{FB}(\mu_0)+\delta}$ is decreasing in

[26] To see why, take $\mu_0' > \mu_0''$. It is clear that $\Pr(\text{cont})_{t^{FB}(\mu_0')+\delta} \ge \Pr(\text{cont})_{t^{FB}(\mu_0'')+\delta}$ for $t^{FB}(\mu_0') + \delta \le \underline{t}(\mu_0')$, as $\mu_{\underline{t}}$ is decreasing in μ_0. Moreover, since $\eta_{t^{FB}(\mu_0')+\delta} = \eta_{t^{FB}(\mu_0'')+\delta}$ for $t^{FB}(\mu_0') + \delta \ge \underline{t}(\mu_0')$, the percentage change over time in $\Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}$ must be the same under μ_0' and μ_0'' at all $t^{FB}(\mu_0') + \delta \ge \underline{t}(\mu_0')$, and hence we also obtain $\Pr(\text{cont})_{t^{FB}(\mu_0')+\delta} \ge \Pr(\text{cont})_{t^{FB}(\mu_0'')+\delta}$ for all those times.

μ_0. The derivative of $\eta_{t^{FB}(\mu_0)+\delta}$ with respect to μ_0 being negative implies

$$\frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta} - \frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta} \leq 0,$$

or, equivalently,

$$\frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \leq \frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \eta_{t^{FB}(\mu_0)+\delta}.$$

This inequality is in contradiction with (42).

Effects of information on welfare. We show that the welfare effects of information are positive. Consider a public signal at time 0 that refines μ_0 while satisfying $-x/\mu_0 > R$ and Assumption 1 for all of its realizations. Since first-best welfare increases with information, it suffices to show that flow welfare at any time $t > t^{FB}(\mu_0)$ is convex in μ_0. Note that

$$\frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} = \left\{ \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} - \frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \eta_{t^{FB}(\mu_0)+\delta} \right\} \frac{1}{\Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}.$$

By (37), $\frac{\partial^2 \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq 0$. Hence,

$$0 \geq \left\{ \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} - \frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \eta_{t^{FB}(\mu_0)+\delta} - \frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \right\} \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta} - \frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}.$$

Equivalently,

$$2\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}\frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \geq \left\{ \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} - \frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}\eta_{t^{FB}(\mu_0)+\delta} \right\}.$$

Since $\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \geq 0$ and $\frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \leq 0$, the left-hand side is negative, which implies

$$\frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq \frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}\eta_{t^{FB}(\mu_0)+\delta}. \tag{43}$$

If the derivative on the left-hand side of (43) is negative, the distortion relative to first best is concave in μ_0 and thus welfare is convex in μ_0.

Suppose instead that the derivative on the left-hand side of (43) is strictly positive. Then this equation implies $\frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} > 0$. Suppose for the purpose of contradiction that welfare is concave in μ_0, that is:

$$\frac{\partial^2 \Pr(\text{succeeded \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}\,(\lambda_G + x) + \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}\,(\mu_{t^{FB}(\mu_0)+\delta}\lambda_G + x) < 0. \tag{44}$$

We can rewrite (44) as

$$\frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}(\lambda_G + x) < \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial^2 \mu_0}\lambda_G(1 - \mu_{t^{FB}(\mu_0)+\delta}). \tag{45}$$

Recall that we are considering the case in which the derivative on the left-hand side is strictly positive. Hence, we can follow analogous steps to those in (40)-(42) to show that (45) implies

$$\frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}\eta_{t^{FB}(\mu_0)+\delta} < \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}. \tag{46}$$

This inequality is in contradiction with (43).

Finally, consider a fully informative signal that reveals at time 0 whether the project is good or bad. Arguments analogous to those in the proof of Proposition 3 imply that this signal eliminates distortions and increases welfare.

Effects of μ_0 and information on distortion relative to first best. We show by example that an increase in the prior μ_0 and imperfect information that refines μ_0 can reduce the distortion relative to first best. To do this, we compute numerically the equilibrium for different prior beliefs.

We approximate the continuous time outcome by taking a discrete time model with periods of small length. Specifically, discretize time in periods of dt length, so $t \in \{0, dt, 2dt, ...\}$, and assume t^{FB} and $\underline{t}$ are on the grid (i.e., t^{FB}/dt and $\underline{t}/dt$ are integers). The probability that a good project succeeds over a period of length dt is $\lambda_G dt$. The probability that a good project succeeds before time $\underline{t}$ is $1-(1-\lambda_G dt)^{\frac{t}{dt}}$.

Recall that the agent continues with certainty until time $\underline{t}$, and we can compute the market's belief $\widehat{\mu}_t^1$ for each time $t \geq \underline{t}$ using equation (23). Using $\widehat{\mu}_t^1$, we can then solve for the probability with which the agent continues at each time. Call $\gamma_{\underline{t}} dt$ the probability that the agent stops over $[\underline{t}, \underline{t}+dt]$ absent success. Then

$$\widehat{\mu}_{\underline{t}+dt}^1 = \frac{\mu_0\left[1-(1-\lambda_G dt)^{\frac{t}{dt}}\right]+\mu_0(1-\lambda_G dt)^{\frac{t}{dt}}(1-\gamma_{\underline{t}} dt)}{\mu_0\left[1-(1-\lambda_G dt)^{\frac{t}{dt}}\right]+(\mu_0(1-\lambda_G dt)^{\frac{t}{dt}}+1-\mu_0)(1-\gamma_{\underline{t}} dt)}.$$

Similarly, call $\gamma_{\underline{t}+dt} dt$ the probability that the agent stops over $[\underline{t}+dt, \underline{t}+2dt]$ absent success. The probability that an agent who has a bad project will stay until time $\underline{t}+2dt$ is $(1-\gamma_{\underline{t}} dt)(1-\gamma_{\underline{t}+dt} dt)$. The probability that an agent who has a good project and had not succeeded by time $\underline{t}$ will stay until $\underline{t}+2dt$ is $(1-\gamma_{\underline{t}} dt)[\lambda_G dt+(1-\lambda_G dt)(1-\gamma_{\underline{t}+dt} dt)]$. Thus,

$$\widehat{\mu}_{\underline{t}+2dt}^1 = \frac{\mu_0\left[1-(1-\lambda_G dt)^{\frac{t}{dt}}\right]+\mu_0(1-\lambda_G dt)^{\frac{t}{dt}}(1-\gamma_{\underline{t}} dt)[\lambda_G dt+(1-\lambda_G dt)(1-\gamma_{\underline{t}+dt} dt)]}{\left[\begin{array}{c}\mu_0\left[1-(1-\lambda_G dt)^{\frac{t}{dt}}\right]+\mu_0(1-\lambda_G dt)^{\frac{t}{dt}}(1-\gamma_1 dt)[\lambda_G dt+(1-\lambda_G dt)(1-\gamma_{\underline{t}+dt} dt)] \\ +(1-\mu_0)(1-\gamma_{\underline{t}} dt)(1-\gamma_{\underline{t}+dt} dt)\end{array}\right]}.$$

We perform analogous computations for $\underline{t}+3dt$, $\underline{t}+4dt$, and so on. Equilibrium welfare can then be written as

$$\begin{aligned}
&\mu_0 \sum_{t=0}^{\underline{t}-dt}(1-rdt)^{\frac{t}{dt}}(1-\lambda_G dt)^{\frac{t}{dt}}(xdt+\lambda_G dt\ v)+(1-\mu_0)\sum_{t=0}^{\underline{t}-dt}(1-rdt)^{\frac{t}{dt}} xdt \\
&+\mu_0 \sum_{t=\underline{t}}^{\infty}(1-rdt)^{\frac{t}{dt}}(1-\lambda_G dt)^{\frac{t}{dt}}\Pi_{s=\underline{t}}^t(1-\gamma_s dt)(xdt+\lambda_G dt\ v) \\
&+(1-\mu_0)\sum_{t=\underline{t}}^{\infty}(1-rdt)^{\frac{t}{dt}}\Pi_{s=\underline{t}}^t(1-\gamma_s dt)xdt,
\end{aligned}$$

where $\Pi_{s=\underline{t}}^{t}(1-\gamma_s dt) = (1-\gamma_{\underline{t}}dt)(1-\gamma_{\underline{t}+dt}dt)(1-\gamma_{\underline{t}+2dt}dt)\ldots(1-\gamma_{\underline{t}+ndt}dt)$ for $\underline{t}+ndt = t$. We take a large time T (on the grid) such that γ_t is virtually zero for $t > T$ and approximate welfare by computing

$$\begin{aligned} S &= \mu_0 \sum_{t=0}^{\underline{t}-dt}(1-rdt)^{\frac{t}{dt}}(1-\lambda_G dt)^{\frac{t}{dt}}(xdt+\lambda_G dt\ v) + (1-\mu_0)\sum_{t=0}^{\underline{t}-dt}(1-rdt)^{\frac{t}{dt}}xdt \\ &+\mu_0 \sum_{t=\underline{t}}^{T-dt}(1-rdt)^{\frac{t}{dt}}(1-\lambda_G dt)^{\frac{t}{dt}}\Pi_{s=\underline{t}}^{t}(1-\gamma_s dt)(xdt+\lambda_G dt\ v) \\ &+(1-\mu_0)\sum_{t=\underline{t}}^{T-dt}(1-rdt)^{\frac{t}{dt}}\Pi_{s=\underline{t}}^{t}(1-\gamma_s dt)xdt \\ &+(1-rdt)^{\frac{T}{dt}}\Pi_{s=\underline{t}}^{T}(1-\gamma_s dt)\left[\mu_0(1-\lambda_G dt)^{\frac{T}{dt}}\frac{(x+\lambda_G v)}{r+\lambda_G - r\lambda_G dt} + (1-\mu_0)\frac{x}{r}\right]. \end{aligned}$$

Finally, we compute first-best welfare,

$$S^{FB} = \mu_0 \sum_{t=0}^{t^{FB}}(1-rdt)^{\frac{t}{dt}}(1-\lambda_G dt)^{\frac{t}{dt}}(xdt+\lambda_G dt\ v) + (1-\mu_0)\sum_{t=0}^{t^{FB}}(1-rdt)^{\frac{t}{dt}}xdt,$$

where $v = 1+\frac{x+\lambda_G}{r}$, and we compute the distortion, $D = S^{FB} - S$.

Consider the parameters reported in the example of Figure 3, with a prior $\mu_0 = 0.5$, and take periods of length $dt = 0.001$. We verify that γ_t becomes virtually zero after a large enough number of periods; accordingly, we compute equilibrium welfare S above for $T = 3,000$. Let $\mu_0' = 0.75$ and $\mu_0'' = 0.25$, and denote by $D(\mu_0)$ the distortion given a prior belief μ_0. We obtain $D(\mu_0) = 0.0426$, $D(\mu_0') = 0.0409$, and $D(\mu_0'') = 0.0286$. Hence, an increase in the prior from μ_0 to μ_0' reduces the distortion relative to first best. Furthermore, take a binary public signal that increases the prior to μ_0' when the realization is high and decreases the prior to μ_0'' when the realization is low, with each realization occurring with equal probability. Since $D(\mu_0) > 0.5D(\mu_0') + 0.5D(\mu_0'')$, releasing this public signal at time 0 reduces the distortion relative to first best.