Experimenting with Career Concerns*

Marina Halac† Ilan Kremer‡

August 29, 2016

Abstract

A manager who learns privately about a project over time may be reluctant to quit it if recognizing failure/lack of success hurts his reputation. In the banking industry, managers may want to roll over bad loans that would reflect poorly on their skills. 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 the arrival of bad news, i.e. a default. We show that in such an environment, the distortions due to career concerns tend to increase with expected quality. Furthermore, while perfect information about quality increases welfare, imperfect information can be detrimental. These results resonate with evidence on banks' behavior and have policy implications. We find that dynamic aspects matter: results differ when managers learn about project quality from the arrival of good news, as in venture capital.

*We would like to thank Simon Board, Patrick Bolton, Tri Vi Dang, Peter DeMarzo, Wouter Dessein, Marco Di Maggio, Zhiguo He, Johannes Hörner, Matt Jackson, Navin Kartik, Qingmin Liu, Matteo Maggiori, George Mailath, Konrad Mierendorff, Pauli Murto, Andrea Prat, Andy Skrzypacz, Lars Stole, Phillip Strack, and various seminar and conference audiences for helpful comments. Yifeng Guo provided excellent research assistance.

†Columbia University and University of Warwick. Email: `mhalac@columbia.edu`.

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

1 Introduction

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, 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."[1] 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 higher 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*

[1] 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."

priori unclear: perfect public 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.[2] This contrasts, for example, with the case of a venture capital fund that invests in a startup firm, or an entrepreneur who invests in an innovation, where learning occurs through the "good news" that arrive when the startup or innovation succeeds.

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.[3] 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.

[2] 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.

[3] In an extension, we consider a variant of this model in which, as in Rajan (1994), the quality of the project is determined by the ability of the manager and the state of the economy, and the manager cares about the market's perception of his ability. We find similar results as in our baseline model. Our qualitative results also apply to a model extension in which, rather than having a reputational concern, the manager is compensated based on how the market values the project. See Section 5.

[4] We study perfect Bayesian equilibria satisfying Divinity. See Section 2.

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 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 (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 find however that in the bad news setting, the effects of information are non-monotonic: a sufficiently imperfect signal both increases the distortion relative to first best and reduces overall welfare. Intuitively, a high signal realization increases

[5] We describe here our results for 'intermediate" parameter values, under which the equilibrium features an interior time $t^* \in (0, \infty)$. See Section 3 for details.

the prior that the project is good and (by the result described above) increases the distortion relative to first best, while 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 information is sufficiently imperfect. These results are also 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. Bonatti and Hörner (2015a) consider a version of Holmström's model with exponential learning. 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), and Aghion and Jackson (2016).[6]

[6] 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.

More closely related to our paper are Grenadier, Malenko and Strebulaev (2014), Bobtcheff and Levy (2015), and Thomas (2015), all of which use exponential learning.[7] 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.[8] Bobtcheff and Levy (2015) 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 (2015) 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, a bad project never yields a payoff, and a good project yields a payoff when a publicly observable breakthrough occurs, at a random time. Before a breakthrough arrives, the agent learns privately from partially informative good or bad news which have no cash-flow consequences. Thomas's motivation stems from political leaders who are reluctant to effect changes in their policies; her setup is less natural to study distortions in the banking sector, where bad news take the form of default and breakthroughs play no role. The focus is also different: Thomas examines the conditions under which the first best is implementable, with the goal of assessing the extent to which distortions in policymaking can be attributed to leaders' career concerns.[9] Instead, our interest is in analyzing the pattern of distortions, specifically how distortions due to career concerns vary with expected project quality and information about quality.

More broadly, there is 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 (2015b) and Keller and Rady (2015) consider bad news learning, and Che and Hörner (2014), Frick and Ishii (2015), and Khromenkova (2015) compare good news and bad news learning in various contexts.[10] 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

[7] 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.

[8] 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 (2016).

[9] Her findings are different from ours, owing to the different modeling assumptions.

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

the exponential bandit framework, Hörner and Lambert (2016).[11]

Finally, there is a related literature in finance. As mentioned above, Rajan (1994) studies a static model in which a career-concerned bank manager observes the quality of a loan and chooses whether to implement a liberal credit policy, namely a policy that makes a bad loan less likely to be reflected in the bank's earnings. The paper assumes that the loan is good if and only if both the manager's ability and the state of the economy are high, an assumption that we adopt as an extension of our model in Section 5. Benmelech, Kandel and Veronesi (2010) consider a manager who chooses hidden effort to improve a firm's growth opportunities and an investment strategy given the realized, privately observed opportunities. They show that stock-based compensation induces the manager to exert effort but also to conceal growth slowdowns and choose inefficient investment strategies. Makarov and Plantin (2015) examine a model in which a fund manager's compensation depends on investors' perception of his ability to generate excess returns above a fair compensation for risk (alpha). The paper shows that the manager will take on hidden tail risk in order to distort his perceived ability temporarily. Winton and Yerramilli (2015) analyze a repeated game (without learning) of originate-to-distribute lending, in which banks' incentives to monitor the loans they sell depend on the extent future investors can punish bad loan performance. They find that if booms are uncommon, monitoring is harder to sustain in boom times.

2 Model

Players and actions. There is an agent and a market. Time is continuous and 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, $a_t = 1$, or stop, $a_t = 0$. To simplify the exposition, we assume that stopping is irreversible: if $a_t = 0$, then $a_\tau = 0$ for all $\tau > t$.

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

[11]Other papers include Bimpikis and Drakopoulos (2014), Che and Hörner (2014), Heidhues, Rady and Strack (2014), Kremer, Mansour and Perry (2014), Bimpikis, Ehsani and Mostagir (2015), and Halac, Kartik and Liu (2016). Quah and Strulovici (2009) examine the comparative statics of optimal stopping time problems in a general framework.

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 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.

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 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 $R \geq 0$. This specification captures, in a reduced form, any benefits the agent may enjoy from having a high reputation, such as outside options that increase his wage or the possibility of working on additional projects. Given a sequence of actions $\{a_t\}_{t\geq 0}$ and market beliefs $\{\widehat{\mu}_t\}_{t\geq 0}$, the agent's expected payoff at time zero is

$$\int_0^\infty e^{-rt} \left\{ \left[x + \mu_0 \lambda_G - (1-\mu_0)\lambda_B \right] a_t + \widehat{\mu}_t R \right\} dt. \tag{1}$$

We will refer to the expected payoff from the project — given by expression (1) when $R = 0$ — as social welfare. 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.[12]

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:* $x + \mu_0 \lambda_G - (1-\mu_0)\lambda_B > 0$.

Model variants. In Section 5, we describe a variant of this model in which the quality of the agent's project is determined by the agent's ability and the state of the economy, both initially unknown. The agent cares about the market's perception of his ability: he receives a reputation benefit at each time that is proportional to the market's belief that his ability is high. Our baseline model does not distinguish between ability and state of the economy, and instead assumes that the agent's reputation benefits are proportional to the market's belief of his project's quality. In addition to being simpler, an appealing feature of our model is that it recognizes commonly observed frictions in managerial compensation. In particular, there is evidence that managers are often rewarded for absolute rather than relative performance (e.g., Bertrand and Mullainathan, 2001), and a number of theories explain why markets and firms may not filter out common factors (e.g., Oyer, 2004, and references therein). Accordingly, we posit that an agent's reputation will increase with perceived project quality, even if quality may be partly driven by the state of the economy. As discussed in Section 5, we find similar results in the two model formulations.

Our model can also be modified to reflect the effects of stock-based compensation. Specifically, as argued in Section 5, our qualitative results continue to apply if, rather than having a reputational concern, the agent is compensated based on how the market values the project.

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 consisting of the agent's actions

[12] One may posit that society also values acquiring information, beyond the effects on the project's profits; see Section 5 for a discussion.

up to t. A strategy for the agent is a measurable function that specifies, for each h^t, a stopping probability at t. The agent maximizes his expected payoff given his observed history and the prior, and the market forms its belief about the agent's project given the public history and prior.

We solve for the perfect Bayesian equilibria of this game.[13] 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. Specifically, we adopt Divinity, a refinement introduced by Banks and Sobel (1987) that requires off-the-equilibrium-path beliefs to place relatively more weight on types that gain more from deviating from a fixed equilibrium. Recall that in our game the agent has no private information at time 0, and his type at any time $t > 0$ is given by his private history of payoff realizations. Divinity imposes that if either stop ($a_t = 0$) or continue ($a_t = 1$) is played with zero probability at time t, then upon observing this action the market assigns (weakly) higher probability to the agent types that would find deviating to such action more attractive. We will indicate where we use this refinement, both in our discussions in the text and in the proofs in the Appendix. Divinity is weaker than D1 (Cho and Kreps, 1987) and other refinements used in the literature. From now on, we refer to perfect Bayesian equilibria satisfying Divinity as simply equilibria.

3 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

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

[13] We should note that perfect Bayesian equilibrium (Fudenberg and Tirole, 1991) requires not only that beliefs be updated according to Bayes' rule whenever possible, but also what is known as the "no signaling what you don't know" condition. In our game, since the agent has no private information at time 0, this condition implies that the market's belief cannot change upon observing the agent's decision of whether to start or not the project. See Watson (2016) for a general definition of perfect Bayesian equilibrium in extensive form games.

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.

3.1 First best

Suppose $R = 0$, so the agent does not have a career concern and maximizes social welfare. Since the agent's belief that the project is good increases over time in the absence of failure, the first-best solution is straightforward: if the agent works on the project for any positive amount of time, he works so long as he has not failed, and stops immediately if he fails. 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.

3.2 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 the project's 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 by contradiction that there exists a time $t > 0$ at which an agent who has not failed stops with strictly positive probability. Note that the expected payoff from continuing with the project is always strictly smaller for an agent who has failed than for one who has not, whereas the payoff from stopping is the same for both types of the agent. Hence, an agent who has failed by t stops with certainty by t, and for any given market beliefs he finds continuing with the project strictly less attractive than the agent who has not failed. As a result, the market's belief that the agent has not failed — and hence its belief that the project is good — (weakly) increases upon observing that the agent continues at t: this follows from Bayes' rule if the agent who has not failed stops with strictly interior probability at t, and from Divinity if he stops with probability one. Now this implies that there exists a continuation strategy for the agent in which he continues with the project at t and his expected reputation payoff is higher than if he stops.[14] Moreover, note that for the agent who has not failed, the expected payoff from the project is strictly positive for any such continuation strategy (by Assumption 1 and this type's belief at t being $\mu_t > \mu_0$), whereas the project payoff is zero if the agent stops. Therefore, there exists a continuation strategy in which the agent continues with the project at time t and which gives the agent who has not failed a strictly larger expected payoff than that from stopping at t. This contradicts our initial assumption that this agent stops with strictly positive probability at t.

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 belief that the project is good at $\tau \geq t$ is zero.*

Note that the market 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 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, so long as the agent continues with the project, $\widehat{\mu}_t^1$ would increase over time.

Given these market 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 that the project is good 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.

[14] Plainly, the agent can follow the continuation strategy after t that the market conjectures he would follow in the absence of failure. Such a strategy ensures a continuation reputation payoff no smaller than that corresponding to a market belief that remains constant after the agent continues at t, which by the arguments above is higher than the market belief after the agent stops at t.

Hence, the agent's expected payoff from continuing is

$$\int_t^\infty e^{-r(\tau-t)} \left(x - \lambda_B + \mu_\tau R\right) d\tau. \tag{5}$$

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, on the other hand, 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: the agent's cost of losing his reputation exceeds his cost of continuing working on a bad project forever.

Proposition 1 (Bad news setting)**.** *The equilibrium is unique. If $R \leq -(x - \lambda_B)$, the equilibrium implements the first best. If instead $R > -(x - \lambda_B)$, the equilibrium is a threshold equilibrium with threshold time $t^* \geq 0$: the agent starts working at time 0, stops immediately upon failure if he fails at $t < t^*$, and continues working forever otherwise. If $R \geq -(x - \lambda_B)/\mu_0$, then $t^* = 0$; otherwise, t^* is given by*

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

In a threshold equilibrium with threshold time t^*, the agent stops at a time t if and only if a failure occurs at t and $t < t^*$. As illustrated in Figure 1, the market's time-t equilibrium belief that the agent's project is good is $\widehat{\mu}_t^1 = \mu_t$ if the agent has not stopped by t and $t \leq t^*$, $\widehat{\mu}_t^1 = \mu_{t^*}$ if the agent has not stopped by t and $t > t^*$, and zero if the agent has stopped by t. If $R \geq -(x - \lambda_B)/\mu_0$, the agent always prefers to continue after failure, so the threshold is $t^* = 0$. If $R < -(x - \lambda_B)/\mu_0$, as in the example of Figure 1, then the threshold 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).

The agent's expected payoff at time 0 from following the equilibrium threshold strategy

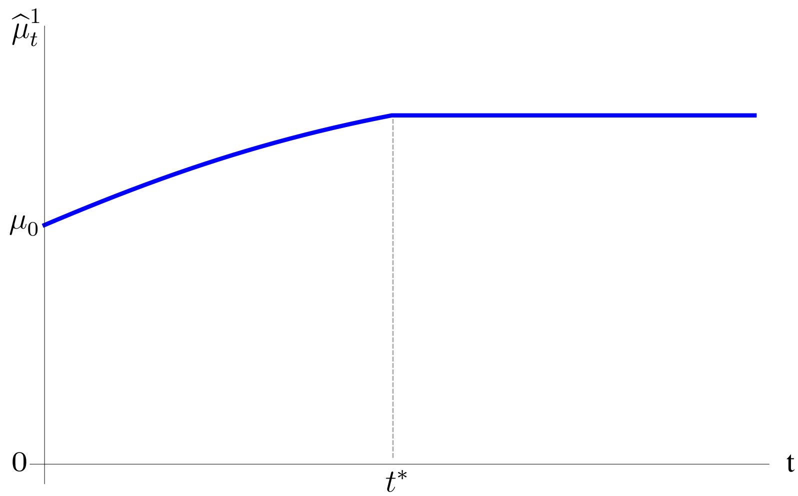

Figure 1: Beliefs and threshold time in the equilibrium of the bad news setting. Parameters are $\mu_0 = 0.6, x = 0.75, \lambda_B = 1.8, R = 1.2$, and $r = 1$.

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 threshold strategy rather than never working is $\pi_0(0) \geq \mu_0 R/r$, which is satisfied by Assumption 1. In fact, given Assumption 1, we can show that a no-work equilibrium 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 are those that correspond to the market expecting the agent to never stop after failure once he starts working (so that the market's belief conditional on the agent continuing is constant at μ_0), and believing that the agent has failed if he ever stops.[15] However, since the market's belief is constant at μ_0 when the agent does not start the project, it is clear that under Assumption 1 the agent would then prefer following a threshold strategy with threshold $t^* = 0$ to never working.

We can therefore show that if $R > -(x - \lambda_B)$, the unique equilibrium is the threshold equilibrium characterized in Proposition 1. Applied to the banking industry, this 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

[15]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.

bank's reputation goes up, and at some point the reputational cost of admitting losses becomes high enough that the bank prefers to roll over a bad loan.[16]

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. Note though 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 a threshold equilibrium with threshold time t^*, $\eta_t = 0$ for $t < t^*$ and

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

for $t \geq t^*$, as 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.

3.3 Quality and information

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

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

As $t^* \to \infty$, $S_0(t^*) \to S_0^{FB}$, where first-best welfare S_0^{FB} is given by equation (4). Clearly, *ceteris paribus*, welfare $S_0(t^*)$ decreases, and the distortion $S_0^{FB} - S_0(t^*)$ increases, when the threshold time t^* decreases.

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

[16] While here the agent may keep a bad project forever, an extension of our model in which bad news become public with positive probability yields analogous results and ensures that the agent stops in finite time. See Section 5 for a discussion.

up. Moreover, since the agent's career concern affects outcomes 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 threshold time t^* at which the agent is indifferent between stopping and continuing upon failing decreases. We show, in fact, that this effect dominates: if the equilibrium threshold time t^* is strictly positive, then the decline in this time when μ_0 (marginally) increases causes the distortion relative to first best to increase.

Proposition 2 (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.*

To see the 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 if failure occurs after the market's belief has reached μ_{t^*}, given by (7). Observe that this posterior belief μ_{t^*} is independent of the prior μ_0. Hence, once the market's belief has reached μ_{t^*}, the probability that the project is bad $(1 - \mu_{t^*})$, and therefore the expected losses the agent will generate, 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^* is decreasing in μ_0. A higher prior μ_0 therefore has two implications: first, the agent is less likely to fail before the belief reaches μ_{t^*}, and second, the losses occur earlier in time and are thus less heavily discounted. As a consequence, both the probability that a distortion arises and the present value of the distortion increase when μ_0 increases.

Proposition 2 considers parameters for which the equilibrium features a distortion but this distortion is not extreme; that is, we take parameters satisfying $-(x - \lambda_B)/\mu_0 > R > -(x - \lambda_B)$, so t^* is finite and strictly positive. 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 an increase in this prior increases welfare. 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 cannot affect the agent's behavior further, so $t^* = 0$ remains unchanged 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 a consequence of a corner solution when reputational concerns are 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,

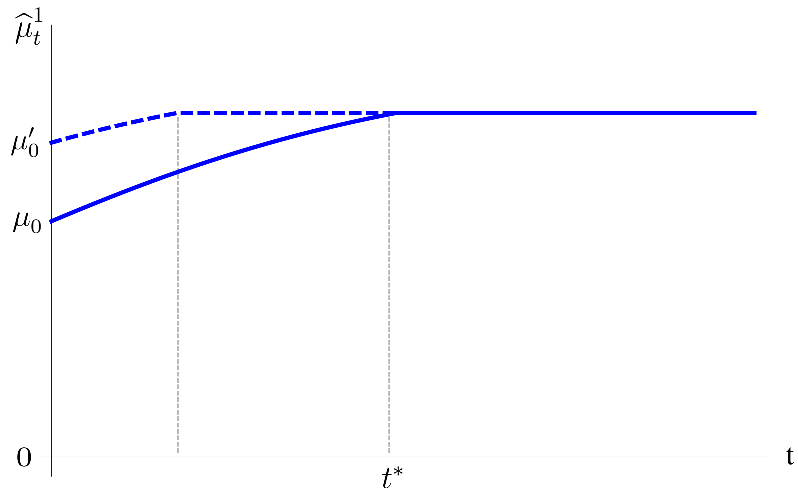

Figure 2: Effects of an increase in μ_0 in the equilibrium of the bad news setting. Parameters are the same as in Figure 1, with $\mu_0' = 0.8$.

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. Proposition 2 shows that career-concerned bank managers will as a result be more likely to roll over their bad loans during good times. Furthermore, despite the proportion of bad borrowers being smaller, banks will in expectation accumulate more bad debt during good times. Consistent with the ideas in Rajan (1994) and the empirical findings in 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 equilibrium 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 a public signal that keeps the efficient 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 on equilibrium welfare, suppose the signal is binary, i.e. its realization either is high and increases the prior μ_0, or is low and decreases μ_0. Building on our discussion of Proposition 2, a high realization of the signal reduces the time t^* that it takes for the market's belief to reach μ_{t^*}, whereas a low realization increases t^*. Once t^* is reached, the distortion due to the agent not stopping upon failure is independent of μ_0; what matters is with which probability t^* is reached and how heavily the expected future losses at t^* are discounted. The key point is that both of these effects are convex: the probability that the agent fails before t^* and discounting of the distortion at t^* decrease when t^* goes down by more than they increase when t^* goes up. While a complete argument must also take into account by how much t^* goes down and up following a high and low signal realization respectively, the proof of Proposition 3 shows that the net effect of the signal is to increase the distortion and reduce welfare.

Things of course are different if the public signal is perfect. 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 Proposition 3, 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.

4 Good news

We contrast the bad news setting of Section 3 with a good news setting in which the agent learns about project quality through 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.

4.1 First best

Suppose $R = 0$, so the agent does not have a career concern and maximizes social welfare. Denote by v the present value of a success, which is equal to the payoff of 1 plus the present value of working in the future on a project that is known to be good: $v := 1 + \frac{\lambda_G + x}{r}$. The agent always continues working on the project if he has succeeded; if no success has occurred, the agent works so long as the expected marginal benefit of effort is larger than the marginal cost, $\mu_t \lambda_G v + x \geq 0$, and he stops otherwise. The first-best solution is therefore a stopping strategy with stopping belief

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

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, $t^{FB} > 0$.

4.2 Career concerns

Suppose now $R > 0$, so the agent cares not only about the payoff from the project but also about the market's perception of the project's 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 $\tau > t$ 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 $\widehat{\mu}_t^1$. This belief depends on the agent's strategy. Suppose the agent were to follow a pure stopping strategy where he runs the project until a finite time t and, in the absence of success, stops at t with certainty. Then if the agent does not stop at t, the market's belief that the agent has succeeded, and thus its belief that the project is good, jumps 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. In fact, we show:

Proposition 4 (Good news setting). *The equilibrium is unique up to off-the-equilibrium-path beliefs. There exist $\underline{t} > 0$ and $\overline{t} \geq \underline{t}$ such that the agent runs the project until $\underline{t}$ with certainty and mixes between stopping and continuing at each $t \in [\underline{t}, \overline{t}]$ absent success. The agent continues forever following a success at any time, and he stops with certainty by $\overline{t}$ if $\overline{t}$ is finite and he has not succeeded by this time. The agent's stopping policy is such that the market's belief $\widehat{\mu}_t^1$ evolves continuously at all $t > 0$. The threshold times satisfy $\underline{t} < \infty$ if and only if $R < -x/\mu_0$ and $\overline{t} < \infty$ if and only if $R < -x$. Furthermore, for all parameters, $\underline{t} > t^{FB}$, so the agent over-experiments relative to first best.*

The equilibrium is characterized by two threshold times, $\underline{t} > t^{FB}$ and $\overline{t} \geq \underline{t}$. The agent continues working on the project until $\underline{t}$ with certainty. From $\underline{t}$ on, the agent implements a

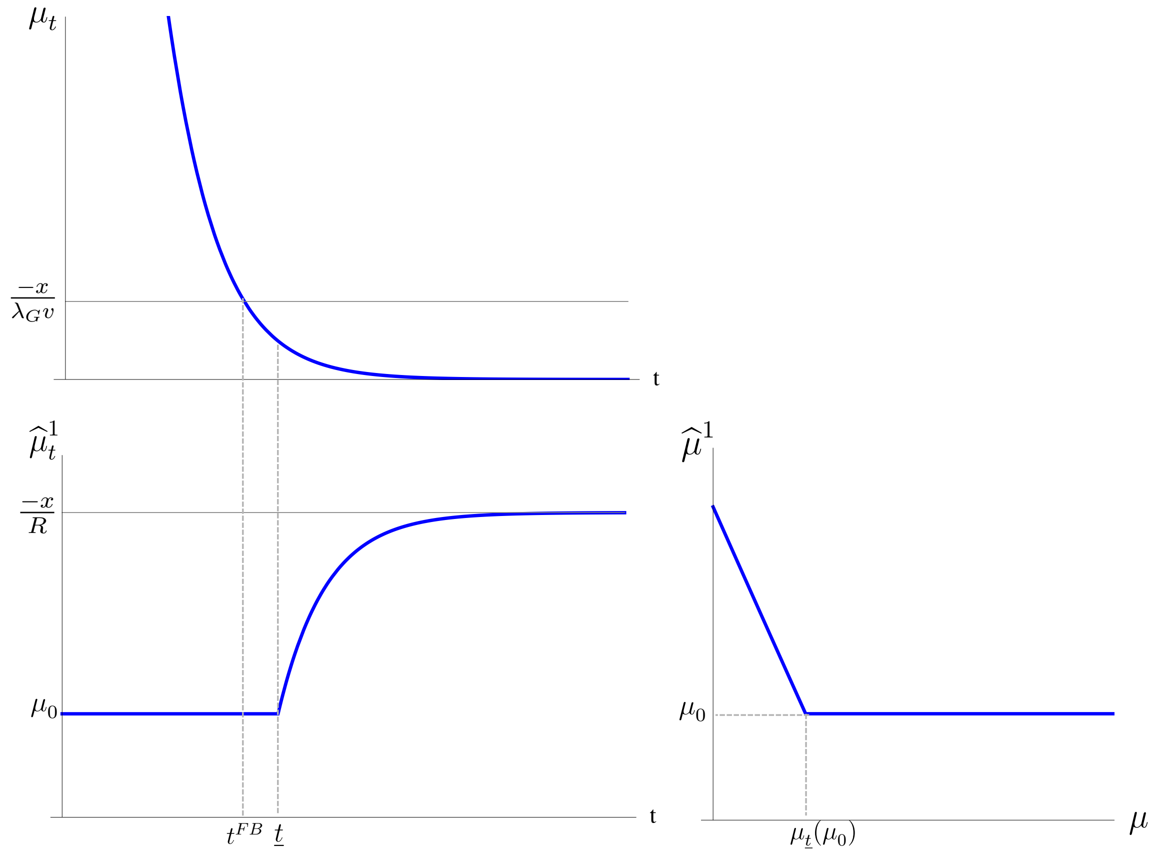

Figure 3: Beliefs and threshold times in the equilibrium of the good news setting. Parameters are $\mu_0 = 0.5, x = -0.8, \lambda_G = 6, R = 0.9$, and $r = 1$.

random stopping policy so that the market's perception evolves continuously. If $R < -x$, the probability with which the agent stops at each time absent success is such that the agent stops with certainty by a finite time $\bar{t}$ (if he has not succeeded by then). If $R > -x$, the agent continues mixing forever, and the probability that an agent who has not succeeded works at time t is strictly positive in the limit as $t \to \infty$. For all parameters, the equilibrium with career concerns induces inefficient experimentation: the agent over-experiments to increase his reputation.

The market's equilibrium belief that the project is good conditional on the agent continuing through time t is $\widehat{\mu}^1_t = \mu_0$ for $t \leq \underline{t}$, $\widehat{\mu}^1_t \in (\mu_0, 1)$ for $t \in (\underline{t}, \bar{t})$, and $\widehat{\mu}^1_t = 1$ for $t \geq \bar{t}$ if $\bar{t}$ is finite. The market's belief conditional on the agent having stopped at time $t \in [\underline{t}, \bar{t}]$ is μ_t for all $\tau \geq t$. Figure 3 illustrates with an example. The graphs on the left show the evolution of the agent and market's beliefs over time: the first best dictates that the agent stop at time t^{FB} in the absence of success, but the career-concerned agent continues with certainty until $\underline{t}$ and, in this example, mixes between stopping and continuing (absent success) at all $t \geq \underline{t}$. The proof of Proposition 4 shows that the market's belief conditional on the agent not having stopped, $\widehat{\mu}^1$, can be written as a function of the belief conditional on no success, μ, independent of time. The right graph in Figure 3 depicts this function. For $\mu > \mu_{\underline{t}}$, the market's belief is constant at $\widehat{\mu}^1 = \mu_0$; as μ falls below $\mu_{\underline{t}}$, the market's belief increases towards its long-run value.

To solve for the threshold times $\underline{t}$ and $\bar{t}$ and the evolution of $\widehat{\mu}^1_t$ between these times, consider the agent's incentives. At each $t \in [\underline{t}, \bar{t}]$ at which the agent stops with an interior probability (given no success by t), the agent must be indifferent between stopping and continuing. Specifically, the agent's payoff from stopping at t must be equal to his expected payoff from continuing working for an arbitrarily small amount of time dt and stopping at $t + dt$ if no success is obtained over $[t, t + dt]$:

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

where $V_t = v + R \int_t^\infty e^{-r(s-t)} \widehat{\mu}^1_s ds$ and we are ignoring $(dt)^2$ terms. Furthermore, note that since this condition must hold at each time $t \in [\underline{t}, \bar{t}]$, it must be that at each such time, the agent is also indifferent between stopping and continuing all the way until time $\bar{t}$. In the proof of Proposition 4, we show that the agent's indifference conditions yield a closed-form solution for $\underline{t}$ and $\bar{t}$ and the market's belief $\widehat{\mu}^1_t$. Moreover, this solution is unique. 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. Unlike under bad news learning, however, this entrepreneur keeps learning about project quality as he continues beyond the first-best stopping time. In particular, the entrepreneur keeps getting pessimistic about the project as time passes without success, and hence he requires increasing reputation benefits not to stop. This explains why the market's belief about the project conditional on the agent continuing, $\widehat{\mu}^1_t$, must be increasing over time.

Similar to our analysis in Section 3, we can also 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.*

4.3 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 (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 $\underline{t} \in (t^{FB}, \infty)$ and $\bar{t} = \infty$) and consider changes in μ_0 and public signals that preserve this property and Assumption 1. Welfare increases with μ_0 and with a signal that refines μ_0 at time 0. 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 3, welfare increases with μ_0, but, unlike in that setting, here distortions can decrease with μ_0. To see why, note that the distortion relative to first best in this good news setting is given by the expected losses the agent generates when he continues with the project after his belief has reached μ^{FB}, given by (11). On the one hand, we find that an increase in μ_0 makes the agent stop more slowly after reaching μ^{FB}; in fact, Figure 4 shows that $\mu_{\underline{t}}$ is decreasing in μ_0, so the agent's stopping begins at a lower posterior belief when the prior is higher. This implies that the expected losses the agent generates once μ^{FB} is reached are increasing in μ_0. On the other hand, however, we also find that an increase in μ_0 increases the time that it takes for the agent's belief to reach μ^{FB}; that is, as also shown in Figure 4, t^{FB} is increasing in μ_0. Analogous to our discussion of Proposition 2, but now

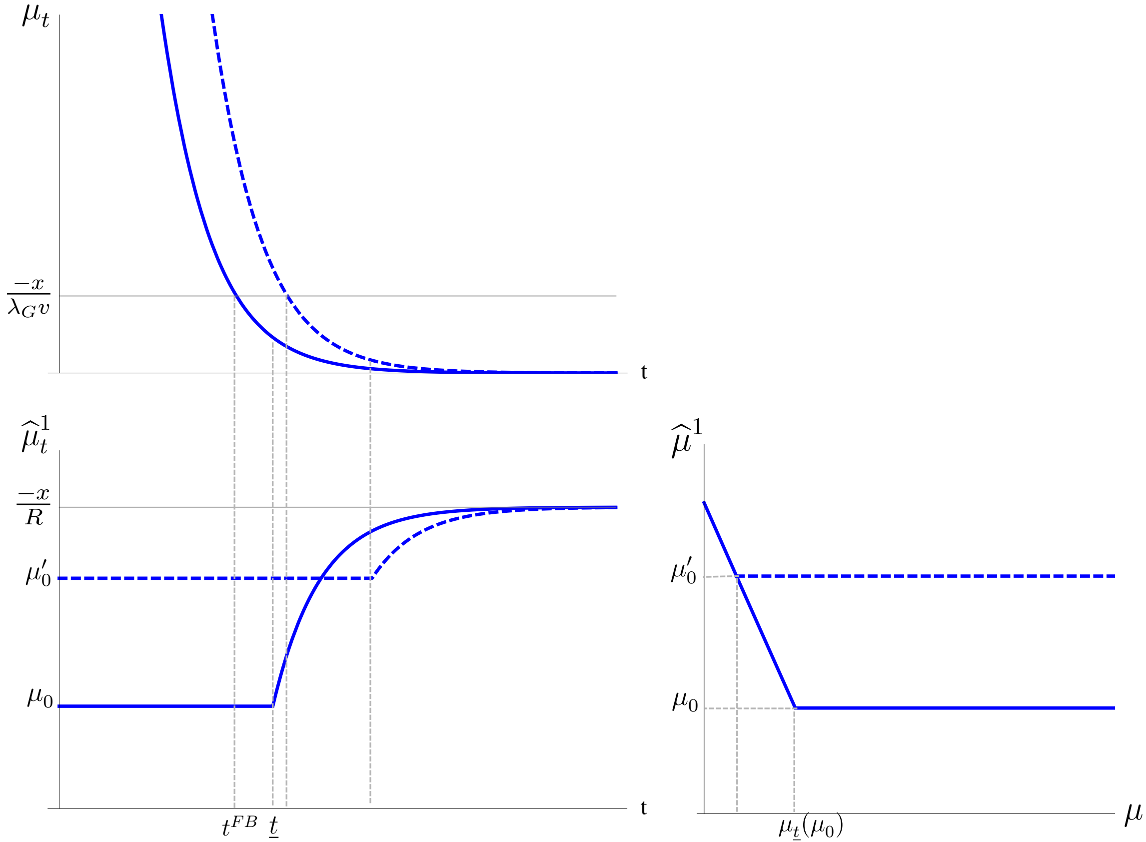

Figure 4: Effects of an increase in μ_0 in the equilibrium of the good news setting. Parameters are the same as in Figure 3, with $\mu_0' = 0.75$.

with the opposite implication, this has two effects: first, the agent is more likely to succeed before time t^{FB} and thus less likely to reach the posterior belief μ^{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; moreover, 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 — and it can also lower distortions. Roughly, building on the intuitions above, the main reason why distortions can decrease with an imperfect signal is that the effects of μ_0 on the agent's stopping policy are concave. That is, an increase in μ_0 reduces the probability with which an unsuccessful agent stops at each time $t > t^{FB}$ by less than what a reduction in μ_0 increases it. As a result, in this good news environment, an imperfect signal that refines μ_0 can reduce distortions in the agent's behavior.

As for the effects of information on welfare, note that releasing a signal about project quality always increases first-best welfare in this setting: while the decision of whether to start the project may remain unchanged, the first-best stopping time is a function of μ_0 and thus changes with a signal that refines this prior. Intuitively, the signal allows the agent to avoid stopping 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, unlike when learning is through bad news, the effects of information on equilibrium welfare are always positive when learning is through 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. Moreover, no matter how imperfect, public information about the prospects of innovation opportunities is always beneficial.

5 Discussion

Agent ability and state of the economy. Consider a variant of our model in which the quality of the agent's project is determined by the agent's ability, $\theta \in \{\underline{\theta}, \overline{\theta}\}$, and the state of the economy, $\omega \in \{\underline{\omega}, \overline{\omega}\}$. As in Rajan (1994), assume that the project is good if $\theta = \overline{\theta}$ and $\omega = \overline{\omega}$, and it is bad otherwise. Both θ and ω are initially unknown to the agent and the market, with prior beliefs $p_0 \in (0, 1)$ on a high-ability agent ($\theta = \overline{\theta}$) and $q_0 \in (0, 1)$ on a high state of the economy ($\omega = \overline{\omega}$). The prior that the project is good is $\mu_0 = p_0 q_0$. The agent cares about the payoff from the project and the market's belief about his ability. Specifically, the agent's expected payoff at time 0 is given by expression (1) but with the reputation benefit at time t being $\widehat{p}_t R$ (instead of $\widehat{\mu}_t R$), where $\widehat{p}_t$ is the market's time-t belief that the agent's ability is high.

This formulation yields similar equilibrium characterizations as our baseline model. To illustrate, consider the bad news setting. For the same reasons discussed in Section 3, one can show that the agent never quits the project in the absence of failure, and thus the market learns that the agent has failed when he stops. This means that upon observing that the agent stops, the market's belief that the project is good is zero. The market's belief that the agent's ability is high, on the other hand, may be strictly positive, as a bad project may be due to a low state rather than a low ability when $q_0 < 1$. Nevertheless, note that the market's belief about the agent's ability conditional on the agent stopping will still be independent of the time at which the agent stops, and the market's belief conditional on the agent continuing will be increasing over time. As a result, we find that the equilibrium is analogous to that in Proposition 1: the agent stops working on the project if and only if a failure arrives before a given threshold time.

This model formulation also yields analogous comparative statics. In particular, motivated by the banking application, consider the effects of changes in the expected state of the economy and public information about the state (and take intermediate parameters as in Section 3 and Section 4). One can show that in the bad news setting, distortions increase with the prior on the high state q_0 and information that refines this prior, whereas in the good news setting distortions can decrease with q_0 and information. The logic behind these results is slightly different from that in our baseline model. Take for instance the bad news setting. Unlike in Proposition 2, here the agent's initial reputation (p_0) is unchanged when q_0 increases; instead, the main effect of an increase in q_0 is to reduce the agent's reputation when the agent stops. Intuitively, upon observing that the agent quits the project, the market learns that the agent has failed, and the higher is q_0 the more the market attributes this bad outcome to a low ability rather than a low state. As a consequence, if q_0 increases, the agent is more reluctant to stop upon failing and the distortion relative to first best increases.

This extension of our model can be viewed as a dynamic version of Rajan (1994), which stresses banks' tendency to hide bad loans during good times. Our analysis reveals that dynamic aspects are important: the pattern of distortions in the banking industry owes to the nature of learning in this industry, namely the fact that banks learn about the quality of loans from the arrival of bad news.

Market valuation. In our model, the agent is concerned about the market's perception of the quality of his project. As we have discussed, the project's quality may reflect the agent's ability, and the agent may receive reputation benefits arising from outside offers that increase his pay or opportunities to work on additional projects.

It is worth noting that our model can be extended to the case in which rather than having a reputational concern, the agent is compensated based on how the market values the project. The agent in this case cares not about the market's perception of the project's quality, but about its perception of the project's value, namely the market's expected future discounted cash flows. This extension applies for example to public companies where at each point in time a manager is compensated based on the stock price.

There are two main differences relative to our baseline model. First, the value of the project depends on both the quality of the project and the agent's actions, so the market's perception at any time would be a function of both its belief about quality and its (correct) conjecture of the agent's equilibrium behavior. Second, as an implication, an agent who stops working on the project would now always receive zero future reputation payoff, in both the bad news and the good news setting. Despite these differences and the analysis becoming more complicated, one can show that equilibrium constructions analogous to ours, and therefore our (qualitative) comparative statics, remain valid under this alternative specification.

Public news. Consider an extension of our model in which the agent's (bad or good) news become public with some positive probability. Specifically, if the agent works at time t and the project is bad, he receives a *privately observed* lump-sum payoff of -1 with instantaneous probability $\alpha\lambda_B$ and a *publicly observed* lump-sum payoff of -1 with instantaneous probability $(1-\alpha)\lambda_B$, for $\alpha \in (0,1)$. If the agent works at time t and the project is good, the instantaneous probability of a *privately observed* lump-sum payoff of 1 is $\alpha\lambda_G$ and that of a *publicly observed* lump-sum payoff of 1 is $(1-\alpha)\lambda_G$. For large enough α, our main insights are unchanged. A realistic feature of this model extension, however, is that a career-concerned agent may now keep a bad project for too long but not forever. This is immediate in the bad news setting: conditional on a bad project, public bad news arrive in finite time, and at that point the agent has strict incentives to stop working on the project. As for the good news setting, note that the market's belief conditional on the agent not having stopped will decrease over time in the absence of public good news unless the agent stops with a large enough probability absent success. As a consequence, an agent who has not succeeded cannot continue working on the project forever in equilibrium.

In this formulation, the probability α parameterizes the agent's degree of private information. A potentially interesting question that is beyond the scope of our paper concerns the effects of changes in α, and how these changes may interact with other parameters such as the prior on a good project, μ_0.

We note that related issues are studied in Thomas (2015). In her model, a project breakthrough is always public, but the agent can learn from privately observed signals (without cash flow consequences) before a breakthrough arrives.

Social value of screening. In our analysis, social welfare is equal to the expected payoff from the project. Society does not benefit from the agent having a high reputation, and efficiency calls for maximizing profits. An alternative would be to posit that, in addition to valuing profits, society values learning about quality. That is, learning may be socially valuable beyond its implications on the project's profits, for example because screening good from bad agents can allow for a better allocation of agents to new projects. What are the welfare effects of changes in μ_0 and information that refines μ_0 when screening is socially valuable? Note that releasing a public signal about project quality at time 0 would now have a new direct positive effect on welfare. Yet, because distortions may increase, it is not *a priori* clear whether the market will indeed be better informed in equilibrium when the public signal is available. We leave these questions for future research.

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.

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*, 1769–1820.

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

Bertrand, Marianne and Sendhil Mullainathan, "Are CEOs Rewarded for Luck? The Ones Without Principals Are," *Quarterly Journal of Economicsa*, 2001, *116* (3), 901–932.

Bimpikis, Kostas and Kimon Drakopoulos, "Disclosing Information in Strategic Experimentation," 2014. Unpublished.

_ , Shayan Ehsani, and Mohamed Mostagir, "Information Provision in Dynamic Innovation Tournaments," 2015. Unpublished.

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," 2015. Unpublished.

Bonatti, Alessandro and Johannes Hörner, "Career Concerns and Market Structure," *Theoretical Economics*, January 2015, *Forthcoming.*

_ and _ , "Learning to Disagree in a Game of Experimentation," January 2015. Unpublished.

Che, Yeon-Koo and Johannes Hörner, "Optimal Design for Social Learning," 2014. Unpublished.

Cho, In-Koo and David M. Kreps, "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 1987, *102* (2), 179–221.

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

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

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

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

Fudenberg, Drew and Jean Tirole, "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 1991, *53* (2), 236–60.

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," 2016. Unpublished.

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.

Halac, Marina, Navin Kartik, and Qingmin Liu, "Contests for Experimentation," *Journal of Political Economy*, 2016, *Forthcoming.*

Heidhues, Paul, Sven Rady, and Philipp Strack, "Strategic Experimentation with Private Payoffs," *Journal of Economic Theory*, 2014, *159* (5), 531–551.

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. Unpublished.

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. Unpublished.

Kremer, Ilan, Yishay Mansour, and Motty Perry, "Implementing the 'Wisdom of the Crowd'," *Journal of Political Economy*, 2014, *122* (5), 988–1012.

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.

Oyer, Paul, "Why Do Firms Use Incentives That Have No Incentive Effects?," *Journal of Finance*, 2004, *59* (4), 1619–1649.

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

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.

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.

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," 2015. Unpublished.

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. Unpublished.

Winton, Andrew and Vijay Yerramilli, "Lender Moral Hazard and Reputation in Originate-to-Distribute Markets," 2015. Unpublished.

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 3; the proofs of the results in Section 4 are in the Online 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 $\tau \geq t$. 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.

A.1 Proof of Lemma 1

We prove the lemma by proving two claims:

Claim 1: There exists no equilibrium in which the agent stops with strictly interior probability at a time $t' > 0$ by which he has not failed.

Proof of Claim 1: Suppose by contradiction that such an equilibrium exists. Since the agent who has not failed stops with strictly interior probability at t', he must be indifferent between stopping and continuing at this time. Note that for any fixed continuation strategy in which the agent continues with the project 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'. Hence, the expected payoff from continuing with the project at t' (and following an optimal continuation plan) is strictly larger for an agent who has not failed. On the other hand, the expected payoff from stopping at t' is the same for both types of the agent. Thus, since the agent who has not failed is indifferent at t', the agent who has failed by t' has strict incentives to stop and stops with probability one by t'. Bayes' rule then implies that the market's belief that the agent has not failed by t' upon observing that the agent continues at t' is one. Therefore, the market's belief about the project at t' is higher if the agent continues than if he stops, i.e. $\widehat{\mu}_{t'}^1 \geq \widehat{\mu}_{t'}^0$.

It follows that given any fixed market conjecture about the agent's stopping strategy upon continuing at t', there must exist a continuation strategy for the agent that yields an expected reputation payoff no smaller than $\widehat{\mu}_{t'}^1 R/r$: plainly, the agent can just follow the stopping plan that the market conjectures he would follow in the absence of failure. Furthermore, note that for an agent who has not failed, the expected payoff from the project is strictly positive if the agent continues working forever ($x-(1-\mu_t)\lambda_B > 0$ by Assumption 1 and $\mu_t > \mu_0$) and it is zero whenever the agent stops. Thus, the expected project payoff from any continuation strategy that continues at time t' is strictly positive for an agent who has not failed. But then the total expected continuation payoff for an agent who has not failed by t' is strictly larger than $\widehat{\mu}_{t'}^{1}R/r$ if the agent continues with the project at t', and it is equal to $\widehat{\mu}_{t'}^{0}R/r \leq \widehat{\mu}_{t'}^{1}R/r$ if the agent stops at t'. This contradicts the assumption that the agent who has not failed by t' stops with strictly positive probability at t'.

Claim 2: There exists no equilibrium in which the agent stops with probability one 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 probability one at a time $t' > 0$, then by the arguments above there is no history following which the agent continues with the project at t', and hence the market belief upon observing that the agent continues at t' is off the equilibrium path. Now note that for any such market belief, the expected continuation payoff from continuing at time t' is strictly larger for an agent who has not failed than for an agent who has failed by t'. Thus, by Divinity, the market's belief that the agent has not failed by t' must weakly increase when the market observes that the agent chooses to continue rather than stop at t'. This in turn implies that the market beliefs that the project is good satisfy $\widehat{\mu}_{t'}^{1} \geq \widehat{\mu}_{t'}^{0}$. The rest of the proof is identical to that of Claim 1.

A.2 Proof of Proposition 1

Existence. We begin by showing existence. Suppose $R \leq -(x - \lambda_B)$. Consider this equilibrium: the agent starts the project at time 0 and stops at time $t > 0$ if and only if he fails at t; the market belief that the project is good at time $t \geq 0$ is 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, beliefs are consistent, 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 at time 0. 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 equilibrium strategy is

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

The first inequality follows from the agent's stopping strategy at $t > 0$ being optimal given the market 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 this equilibrium: 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; the market belief at time $t \geq 0$ is 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, beliefs are consistent, and 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. By Lemma 1, an agent who has not failed does not stop. The following claims are therefore sufficient to prove uniqueness.

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. Consider any arbitrary market conjecture about the agent's continuation strategy when he starts the project. Suppose the agent starts the project and follows the stopping plan that the market conjectures he would follow in the absence of failure. Then the agent ensures himself an expected payoff from reputation no smaller than $\mu_0 R/r$. Moreover, as argued in the proof of Lemma 1, the agent's expected payoff from the project from following any such stopping plan is strictly positive. Hence, the agent's expected payoff from starting the project at time 0 is strictly larger than $\mu_0 R/r$. Since the payoff from not starting the project is $\mu_0 R/r$, we reach a contradiction.

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

Proof of Claim 2: Suppose by contradiction that an agent who fails at a time $t' > 0$ does not stop at t' and stops with strictly positive probability at some time $t'' > t'$. By Lemma 1, the market's belief about the agent's project conditional on the agent not having stopped is weakly increasing. Moreover, if the agent stops with strictly positive probability at t'', the market's equilibrium belief must strictly increase upon observing that the agent does not stop at t''. This means that the agent's expected payoff from reputation from continuing at t'' is strictly larger than that from continuing at t'. On the other hand, the agent's expected continuation payoff from the project is the same when continuing at t' and t'', and by Lemma 1 the expected payoff from stopping at t'' is zero and thus no larger than the expected payoff from stopping at t'. Therefore, if the agent is willing to stop at t'', he has strict incentives to stop immediately after failing at t', yielding a contradiction.

Claim 3: In any equilibrium, the market beliefs satisfy $\widehat{\mu}_t^0 = 0$ for all $t \geq 0$.

Proof of Claim 3: By Lemma 1 and the claims above, the agent's strategy in any equilibrium must be as described in the proposition, with the agent stopping at a time t if and only if he fails at t and $t < t^*$. Bayes' rule then pins down $\widehat{\mu}_t^1$ for all $t \geq 0$ and $\widehat{\mu}_t^0 = 0$ for $t < t^*$. We now show that the market belief must satisfy $\widehat{\mu}_t^0 = 0$ for all $t \geq t^*$. First, note that $\widehat{\mu}_t^0$ cannot jump at t^*: if it did, an agent who fails at time $t < t^*$ 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 belief conditional on the agent continuing is constant at $\widehat{\mu}_{t^*}^1$ for $t \geq t^*$. 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

$$t^* = \frac{1}{\lambda_B} \log \left(\frac{-(1-\mu_0)(x-\lambda_B)}{\mu_0 \left(R + x - \lambda_B\right)} \right). \tag{12}$$

The derivative with respect to μ_0 is

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

Hence, the threshold time t^* is strictly decreasing in μ_0.

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 increasing in μ_0. Differentiating $S_0(t^*)$ given in (8) with respect to μ_0 yields

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

Substituting with (13),

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

Hence, $S_0(t^*)$ is decreasing in μ_0 only if

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

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

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

But this requires $x - (1-\mu_0)\lambda_B \leq 0$, contradicting Assumption 1. Thus, $S_0(t^*)$ is 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) and the proof of Proposition 2 shows that this distortion is increasing in μ_0. Note also that first-best welfare is unchanged with the signal if starting the project is efficient for all realizations of the signal. This follows from the fact that first-best welfare, given in (4), is linear in μ_0. To prove this part, it is therefore sufficient to show that the distortion given in (14) 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 is always satisfied.

Consider next the second part of the proposition. If the signal reveals the project to be good, it is efficient to continue forever and thus there are no distortions. If the signal reveals the project to be bad, it is efficient to not start the project at time 0. Since the agent's actions reveal no information about the project, 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 signal reveals a bad project —

equilibrium welfare increases.

B Online Appendix

This Online Appendix contains the proofs of the results stated in Section 4 of the paper.

B.1 Proof of Proposition 4

Existence. We begin by describing the equilibrium construction and showing existence. Recall that for each $t \in [\underline{t}, \overline{t}]$, the agent's payoff from stopping at t must be equal to his expected payoff from continuing working for an arbitrarily small amount of time dt and stopping at $t + dt$ if no success is obtained over $[t, t + dt]$:

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

where $V_t = v + R \int_t^{\infty} e^{-r(s-t)} \widehat{\mu}_s^1 ds$ and we are ignoring $(dt)^2$ terms. Writing $\mu_{t+dt} = \mu_t + dt \dot{\mu}_t = \mu_t - \mu_t(1 - \mu_t)\lambda_G dt$ and dividing by dt and taking dt to zero yields

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

Additionally, as argued in the text, if the agent is indifferent (absent success) between stopping and continuing for a dt amount of time at each time $t \in [\underline{t}, \overline{t}]$, then he must be indifferent at each such time between stopping and continuing all the way until $\overline{t}$, where $\overline{t}$ is possibly infinite. That is, for each $t \in [\underline{t}, \overline{t})$,

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

Since this condition must hold at each $t \in [\underline{t}, \overline{t}]$, differentiating it yields

$$\dot{\mu}_t R = \dot{\mu}_t \lambda_G \left[1 - e^{-r(\overline{t} - t)} \right] - r (x + \mu_t \lambda_G) e^{-r(\overline{t} - t)} - r R \left[\widehat{\mu}_t^1 - r \int_t^{\overline{t}} e^{-r(s-t)} \widehat{\mu}_s^1 ds \right]. \tag{21}$$

Equation (20) implies $-(x + \mu_t \lambda_G) e^{-r(\overline{t} - t)} + r R \int_t^{\overline{t}} e^{-r(s-t)} \widehat{\mu}_s^1 ds = \mu_t R - x - \mu_t \lambda_G$. Substituting with this equality and $\dot{\mu}_t = -\mu_t(1-\mu_t)\lambda_G$ in (21) and rearranging terms yields

$$\widehat{\mu}_t^1 = \frac{1}{rR}\Big\{r\left[-x-\mu_t(\lambda_G - R)\right] - \mu_t(1-\mu_t)\lambda_G\Big[\lambda_G\Big(1-e^{-r(\bar{t}-t)}\Big) - R\Big]\Big\}. \tag{22}$$

Take first the case in which $-x < R < -x/\mu_0$. We have argued that in this case $\underline{t} \in (t^{FB}, \infty)$ and $\bar{t}$ is infinite. Thus, equations (20) and (22) reduce to

$$\begin{aligned} \mu_t R &= x + \mu_t\lambda_G + rR\int_t^\infty e^{-r(s-t)}\widehat{\mu}_s^1 ds, &(23)\\ \widehat{\mu}_t^1 &= \frac{1}{rR}\Big\{r\left[-x-\mu_t(\lambda_G - R)\right] - \mu_t(1-\mu_t)\lambda_G\left(\lambda_G - R\right)\Big\}, &(24)\end{aligned}$$

for each $t \geq \underline{t}$. Consider this equilibrium: 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 he follows a random stopping policy after time $\underline{t}$ (absent success) such that the market belief $\widehat{\mu}_t^1$ satisfies (24); the market belief at t is given by μ_0 if the agent has not started the project or has started and not stopped by t and $t < \underline{t}$, by $\widehat{\mu}_t^1$ satisfying (24) if the agent has started and not stopped by t and $t \geq \underline{t}$, and by μ_τ if the agent has started and stopped at $\tau \in (0, t]$.

Note that by Assumption 1, $\lambda_G > -x/\mu_0$, and hence $R < -x/\mu_0$ implies $\lambda_G > R$. Since μ_t is decreasing, it follows from (23) that the market belief $\widehat{\mu}_t^1$ is increasing at each time $t \geq \underline{t}$ and thus $\widehat{\mu}_t^1 \geq \mu_0$ for all $t \geq 0$. Equation (23) also shows that $\widehat{\mu}_t^1 \to -x/R$ as $t \to \infty$, where note $-x/R \in (0,1)$.

It can be verified that given Assumption 1 and $-x < R < -x/\mu_0$, there is a unique value $\mu_{\underline{t}} \in (0, \mu_0)$ that solves (24) at $\underline{t}$ (where $\widehat{\mu}_{\underline{t}}^1 = \mu_0$), and this value is given by

$$\mu_{\underline{t}} = \frac{\lambda_G + r - \sqrt{\frac{\lambda_G^2(2r+\lambda_G - R)+\lambda_G r[4(\mu_0 R + x)+r-2R]-r^2R}{\lambda_G - R}}}{2\lambda_G}. \tag{25}$$

The equilibrium defines consistent market beliefs. We next show that given the market beliefs, the agent's stopping plan is optimal. By construction, the agent is indifferent and thus willing to mix between stopping and continuing with any probability at each time $t \geq \underline{t}$ 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 continue with the project at $t < \underline{t}$ in the absence of success. For the start decision, note that the agent's expected payoff if he does not start is $\mu_0 R/r$. The agent's expected payoff if he starts the project and follows the equilibrium strategy is weakly larger than his expected payoff if he starts the project and continues forever, which is given by

$$\int_0^\infty e^{-rt}(x+\mu_0\lambda_G+\widehat{\mu}_t^1 R)dt. \tag{26}$$

As argued above, $\widehat{\mu}_t^1 \geq \mu_0$ for all $t \geq 0$. Moreover, recall that $x+\mu_0\lambda_G > 0$ by Assumption 1. Hence, (26) is greater than $\mu_0 R/r$, implying that it is optimal for the agent to start the project at time 0. Finally, 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 (23) is smaller than its right-hand side for $t < \underline{t}$, or, equivalently,

$$\Psi_t := x + \mu_t(\lambda_G - R) + rR\int_t^\infty e^{-r(s-t)}\widehat{\mu}_s^1 ds \geq 0. \tag{27}$$

By the previous claims, $\Psi_0 > 0$, and by definition of $\underline{t}$, $\Psi_{\underline{t}} = 0$. Therefore, if $\Psi_t < 0$ for some $t < \underline{t}$, then there exist $t' < t'' < \underline{t}$ such that $\Psi_{t'} = 0$, $\dot{\Psi}_{t'} < 0$, $\Psi_{t''} < 0$, and $\dot{\Psi}_{t''} = 0$. Differentiating (27) yields that for $t < \underline{t}$,

$$\dot{\Psi}_t = -\mu_t(1-\mu_t)\lambda_G\left(\lambda_G - R\right) - rR\mu_0 + r^2 R\int_t^\infty e^{-r(s-t)}\widehat{\mu}_s^1 ds. \tag{28}$$

Using (27) and (28), note that $\Psi_{t'} = 0$, $\dot{\Psi}_{t'} < 0$, $\Psi_{t''} < 0$, and $\dot{\Psi}_{t''} = 0$ imply

$$\begin{aligned}
-\mu_{t'}(1-\mu_{t'})\lambda_G\left(\lambda_G - R\right) - rR\mu_0 - r[x+\mu_{t'}(\lambda_G - R)] &< 0, &(29)\\
-\mu_{t''}(1-\mu_{t''})\lambda_G\left(\lambda_G - R\right) - rR\mu_0 - r[x+\mu_{t''}(\lambda_G - R)] &> 0. &(30)
\end{aligned}$$

However, as claimed above, there exists a unique solution $\mu_{\underline{t}} \in (0,\mu_0)$ to equation (24) when $\widehat{\mu}_t^1 = \mu_0$, and therefore we cannot have $\mu_{t'}, \mu_{t''} \in (\mu_{\underline{t}}, \mu_0)$. Contradiction.

Finally, we show that $\underline{t} > t^{FB}$. Note that (19) at $\underline{t}$ yields

$$R(\mu_0 - \mu_{\underline{t}}) + x + \mu_{\underline{t}}\lambda_G\left(V_{\underline{t}} - \frac{R}{r}\right) = 0, \tag{31}$$

where, by (23), $V_{\underline{t}} = v + \frac{\mu_{\underline{t}}R-(x+\mu_{\underline{t}}\lambda_G)}{r}$. Given $V_{\underline{t}}$ and using (9), $\underline{t}$ is therefore given by

$$\underline{t} = \frac{1}{\lambda_G}\log\left(\frac{\mu_0}{1-\mu_0}\left[\frac{R-\lambda_G\left(V_{\underline{t}} - R/r\right)}{\mu_0 R + x} - 1\right]\right). \tag{32}$$

By equations (9) and (11), the first-best stopping time is

$$t^{FB} = \frac{1}{\lambda_G} \log\left(\frac{\mu_0}{1-\mu_0}\left(-\frac{\lambda_G v}{x} - 1\right)\right). \tag{33}$$

To show $\underline{t} > t^{FB}$, it then suffices to show that $\frac{R-\lambda_G\left(V_{\underline{t}}-R/r\right)}{\mu_0 R+x} > -\frac{\lambda_G v}{x}$, which substituting and simplifying terms is equivalent to

$$\mu_0\lambda_G v + x > -x\lambda_G\left[\frac{1}{r} - \int_{\underline{t}}^{\infty} e^{-r(s-\underline{t})}\widehat{\mu}_s^1 ds\right]. \tag{34}$$

Since $\widehat{\mu}_t^1 \geq \mu_0$ for all $t \geq 0$, the right-hand side is smaller than $-x\lambda_G\left(\frac{1-\mu_0}{r}\right)$. Hence, it suffices to show that $\mu_0\lambda_G v + x > -x\lambda_G\left(\frac{1-\mu_0}{r}\right)$. Substituting with the definition of v and rearranging terms, this is equivalent to

$$\mu_0\lambda_G + x + \frac{\lambda_G}{r}(\mu_0\lambda_G + x) > 0,$$

which by Assumption 1 holds.

Take next the case in which $R < -x$. Here the equilibrium construction is analogous to that above, except that the agent stops mixing at a finite time $\bar{t}$, by which an agent who has not succeeded has stopped with certainty. The agent's indifference condition at $\bar{t}$ pins down the threshold time $\bar{t}$: since $\widehat{\mu}_t^1 = 1$ for all $t \geq \bar{t}$, (19) implies

$$R(1-\mu_{\bar{t}}) + x + \mu_{\bar{t}}\lambda_G v = 0. \tag{35}$$

Equation (20) in this case implies $V_{\underline{t}} = v + \frac{\mu_{\underline{t}}R-(x+\mu_{\underline{t}}\lambda_G)\left(1-e^{-r(\bar{t}-\underline{t})}\right)}{r} + e^{-r(\bar{t}-\underline{t})}\frac{R}{r}$. Since $v = V_{\bar{t}} - \frac{R}{r} > V_{\underline{t}} - \frac{R}{r} > 1$, $R < -x$, and $\lambda_G > R$ (by Assumption 1 and $R < -x/\mu_0$), one can then verify that the left-hand side of (35) is strictly negative if $\mu_{\bar{t}} = 0$ and it is strictly positive if $\mu_{\bar{t}} = \mu_{\underline{t}}$. Hence, there exists a posterior belief $\mu_{\bar{t}} \in (0, \mu_{\underline{t}})$, and thus a threshold time $\bar{t} \in (\underline{t}, \infty)$, at which (35) is satisfied. Using (9) and (35), this time is given by

$$\bar{t} = \frac{1}{\lambda_G}\log\left(\frac{\mu_0}{1-\mu_0}\left[\frac{R-\lambda_G v}{R+x} - 1\right]\right). \tag{36}$$

Given μ_t decreasing over time (and the aforementioned inequalities), it also follows that the agent has strict incentives not to continue beyond time $\bar{t}$ if he has not succeeded by then: the left-hand side of (35) evaluated at $t > \bar{t}$ instead of $\bar{t}$ is strictly negative. The market's belief of $\widehat{\mu}_t^1 = 1$ for $t > \bar{t}$ is therefore consistent. Finally, note that the market's belief upon observing that the agent stops at a time $t > \bar{t}$ is off the equilibrium path; we have shown existence when this belief satisfies $\widehat{\mu}_t^0 = \mu_t$ for all $t > 0$.

Lastly, consider the case in which $R > -x/\mu_0$. Suppose the market beliefs satisfy $\widehat{\mu}_t^1 = \mu_0$ and $\widehat{\mu}_t^0 = \mu_t$ for all $t \geq 0$. We show that the agent always has strict incentives to continue with the project given these beliefs, 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. Therefore, when $R > -x/\mu_0$, there is an equilibrium in which the agent never stops (i.e. $\underline{t}$ is infinite) and the market's belief is $\widehat{\mu}_t^1 = \mu_0$ for all $t \geq 0$. Note that the market's belief upon observing that the agent stops at a time $t > 0$ is off the equilibrium path; we have shown existence when this belief satisfies $\widehat{\mu}_t^0 = \mu_t$ for all $t > 0$.

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: There exists no equilibrium in which the agent stops with strictly interior probability at a time $t' > 0$ by which he has succeeded.

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 probability one 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 at some time $t > 0$ does not exist.

Proof of Claim 4: The last paragraph in the Existence section above shows that when $R > -x/\mu_0$, the agent has strict incentives to continue with the project at 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 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 stopping at t in equilibrium requires $\widehat{\mu}_t^0 > \mu_t$, but 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 stops with zero probability at all times 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 with the project is $\widehat{\mu}_t^1 = \mu_0$ for all $t \geq 0$, and the agent must be willing to continue with the project rather than stop at all times. However, since $\mu_t \to 0$ as $t \to \infty$ (and the agent's expected payoff from stopping is weakly positive), 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 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. Then as $t \to \infty$, absent success, the agent must weakly prefer to continue with the project rather than stop. Since $\mu_t \to 0$ as $t \to \infty$ (and the agent's expected payoff from stopping is weakly positive), 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. Hence, if the agent stops with positive probability at a time t, he must be indifferent between stopping and continuing at t.

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 always weakly increasing and can only jump up. Suppose the belief jumps at $\hat{t}$ from $\widehat{\mu}_{\hat{t}^-}^1 = \widehat{\mu}^{1-}$ to $\widehat{\mu}_{\hat{t}^+}^1 = \widehat{\mu}^{1+} > \widehat{\mu}^{1-}$. Then the probability with which an agent who has not succeeded by $\hat{t}$ stops at $\hat{t}$ must be strictly positive and larger than the probability with which the agent stops right before $\hat{t}$. Moreover, the market's belief conditional on the agent stopping at $\hat{t}$ must be consistent, and thus by Claims 2 and 3 it must be $\widehat{\mu}_{\hat{t}}^0 = \mu_{\hat{t}}$. Note also that the lowest belief the market can have upon observing that the agent stops at a time $t \geq 0$ (even when stopping at such a time is off the equilibrium path) is μ_t, i.e. the belief that the agent has not succeeded 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 in the Existence section, taking dt to 0, this condition is

$$R\left(\widehat{\mu}^{1+} - \mu_{\hat{t}}\right) + x + \mu_{\hat{t}}\lambda_G\left(V_{\hat{t}} - \frac{R}{r}\right) \leq 0. \tag{37}$$

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

$$R\left(\widehat{\mu}^{1-} - \widehat{\mu}^{0}_{\hat{t}-}\right) + x + \mu_{\hat{t}}\lambda_G\left(V_{\hat{t}} - \frac{R}{r}\right) \geq 0. \tag{38}$$

However, $\widehat{\mu}^{1+} > \widehat{\mu}^{1-}$ and $\widehat{\mu}^{0}_{\hat{t}-} \geq \mu_{\hat{t}}$ imply that (37) and (38) cannot be simultaneously satisfied. Contradiction.

Claim 8: An equilibrium in which the agent mixes between stopping and continuing only until a finite time $\bar{t}$ *and* $\widehat{\mu}^1_t < 1$ *for* $t \geq \bar{t}$ *does not exist.*

Proof of Claim 8: Suppose by contradiction that such an equilibrium exists. By Claim 4, if the agent mixes between stopping and continuing over a period of time, we must have $R < -x/\mu_0$. In the proposed equilibrium, the agent's quitting ceases at time $\bar{t}$, so the market's belief $\widehat{\mu}^1_t$ is constant at some value, call it $\overline{\mu}$, for $t \geq \bar{t}$. The agent's indifference condition (19) at $\bar{t}$ then reduces to

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

Substituting with $v = 1 + \frac{\lambda_G + x}{r}$ and rearranging terms yields

$$R\overline{\mu} + x + \mu_{\bar{t}}\left(\lambda_G - R + \lambda_G\frac{\lambda_G + x}{r} - \lambda_G(1-\overline{\mu})\frac{R}{r}\right) = 0. \tag{39}$$

Note that $\overline{\mu} < 1$ requires that an agent who has not succeeded by $\bar{t}$ be willing to continue with the project beyond this time. Since μ_t is decreasing over time, equation (39) implies that the agent is willing to continue after time $\bar{t}$ absent success if and only if the expression in parenthesis is negative. That is, rearranging terms, the equilibrium requires

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

Note that Assumption 1 and $R < -x/\mu_0$ imply $\lambda_G > R$. Thus, the above inequality can hold only if $x + \overline{\mu}R < 0$. However, if the parenthesis in (39) is negative and $x + \overline{\mu}R < 0$,

(39) cannot hold. Contradiction.

Claim 9: If $R > -x$, an equilibrium in which the agent mixes between stopping and continuing absent success only until a finite time $\bar{t}$ does not exist.

Proof of Claim 9: Suppose by contradiction that such an equilibrium exists. By Claim 4, if the agent mixes between stopping and continuing over a period of time, we must have $R < -x/\mu_0$. As shown in Claim 8, if the agent's quitting ceases by a finite time $\bar{t}$, the market belief must satisfy $\widehat{\mu}_t^1 = 1$ for all $t \geq \bar{t}$, and hence equation (35) must hold at time $\bar{t}$. However, if $R > -x$, the left-hand side of this equation is strictly greater than zero for all finite $\bar{t}$: this follows from the fact that $\lambda_G > R$ (by Assumption 1 and $R < -x/\mu_0$) and $v > 1$. Contradiction.

Claim 10: Up to off-the-equilibrium-path beliefs, the equilibrium is unique.

Proof of Claim 10: This follows from the claims above and the fact that the solution for $\widehat{\mu}_t^1$ shown in the Existence part of the proof is 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 $\bar{t}$ infinite, where $\underline{t}$ satisfies (31) and equations (23) and (24) hold at each $t \geq \underline{t}$. The market's belief conditional on the agent not having stopped is $\widehat{\mu}_t^1 = \mu_0$ for $t < \underline{t}$, and, by equation (24), this belief can be written as a function of μ_t independent of μ_0 for $t \geq \underline{t}$. Note also that, as illustrated in Figure 4, the posterior belief at which the agent starts mixing, $\mu_{\underline{t}}$, is decreasing in μ_0; this can be verified using (25).

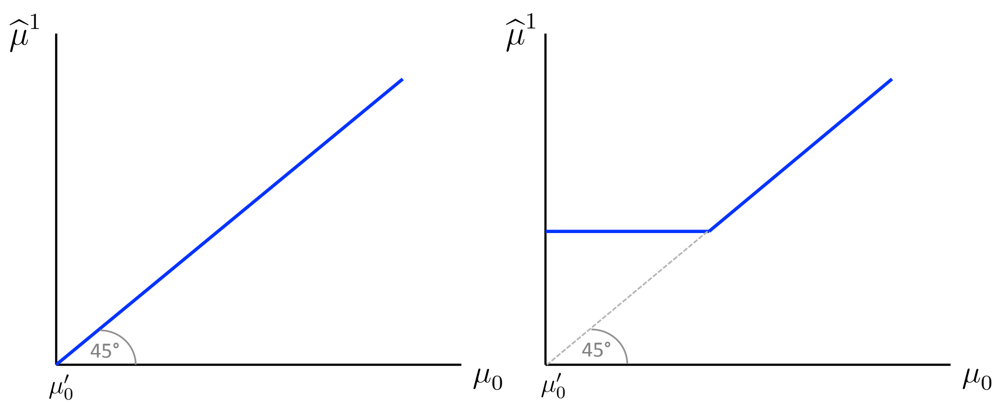

Figure 5: Market belief $\widehat{\mu}^1$ as a function of μ_0 for a fixed posterior belief μ. The left graph corresponds to a fixed posterior belief $\mu > \mu_{\underline{t}(\mu_0')}$, where μ_0' is the lowest prior considered in the figure. The right graph corresponds to a posterior belief $\mu < \mu_{\underline{t}(\mu_0')}$.

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 belief that corresponds to such 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.

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

$$\frac{\partial \widehat{\mu}^1_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \geq 0, \quad \frac{\partial^2 \widehat{\mu}^1_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \geq 0. \tag{40}$$

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

$$\widehat{\mu}^1_{t^{FB}(\mu_0)+\delta} = 1 - \eta_{t^{FB}(\mu_0)+\delta} + \eta_{t^{FB}(\mu_0)+\delta}\ \mu_{t^{FB}(\mu_0)+\delta}. \tag{41}$$

Since $\mu_{t^{FB}(\mu_0)+\delta}$ is independent of μ_0, differentiating (41) 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 (40), we obtain that for any $\delta \geq 0$,

$$\frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \leq 0, \quad \frac{\partial^2 \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq 0. \tag{42}$$

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

$$\Pr(\text{succeeded \& cont})_{t^{FB}(\mu_0)+\delta}\ (\lambda_G + x) \\ + \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}\ (\mu_{t^{FB}(\mu_0)+\delta}\lambda_G + x). \tag{43}$$

Suppose for the purpose of contradiction that for some $\delta > 0$, (43) is decreasing in μ_0, that is (using the fact that $\mu_{t^{FB}(\mu_0)+\delta}$ is independent of μ_0),

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

We can rewrite (44) as

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

Note that the derivative on the left-hand side is positive.[17] Moreover, Assumption 1 implies $(\lambda_G + x)/\lambda_G > 1 - \mu_0$. Hence, (45) 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

$$\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}\ \left(\mu_0 e^{-\lambda_G(t^{FB}(\mu_0)+\delta)} + 1 - \mu_0\right) < \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}. \tag{46}$$

Finally, note that $\eta_{t^{FB}(\mu_0)+\delta} \leq \mu_0 e^{-\lambda_G(t^{FB}(\mu_0)+\delta)} + 1 - \mu_0$, as an agent who has succeeded does not stop. Thus, (46) implies

$$\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}\ \eta_{t^{FB}(\mu_0)+\delta} < \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}. \tag{47}$$

We now show that (47) contradicts (42), namely the fact that $\eta_{t^{FB}(\mu_0)+\delta}$ is decreasing in

[17] To see why, take $\mu_0' > \mu_0''$. It is clear that $\Pr(\text{cont})_{t^{FB}(\mu_0')+\delta} \geq \Pr(\text{cont})_{t^{FB}(\mu_0'')+\delta}$ for $t^{FB}(\mu_0') + \delta \leq \underline{t}(\mu_0')$, as $\mu_{\underline{t}}$ is decreasing in μ_0. Moreover, since $\eta_{t^{FB}(\mu_0')+\delta} = \eta_{t^{FB}(\mu_0'')+\delta}$ for $t^{FB}(\mu_0') + \delta \geq \underline{t}(\mu_0')$, 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 $t^{FB}(\mu_0') + \delta \geq \underline{t}(\mu_0')$, and hence we also obtain $\Pr(\text{cont})_{t^{FB}(\mu_0')+\delta} \geq \Pr(\text{cont})_{t^{FB}(\mu_0'')+\delta}$ for all those times.

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

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

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 (47).

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

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

By (42), $\frac{\partial^2 \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq 0$. Hence,

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

Equivalently,

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

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

$$\frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq \frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}\eta_{t^{FB}(\mu_0)+\delta}. \tag{48}$$

If the derivative on the left-hand side of (48) 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 (48) 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:

$$\frac{\partial^2 \Pr(\text{succeeded \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}(\lambda_G + x)$$
$$+ \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}(\mu_{t^{FB}(\mu_0)+\delta}\lambda_G + x) < 0. \tag{49}$$

We can rewrite (49) as

$$\frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}(\lambda_G + x) < \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial^2 \mu_0}\lambda_G(1 - \mu_{t^{FB}(\mu_0)+\delta}). \tag{50}$$

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 (45)-(47) to show that (50) implies

$$\frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}\eta_{t^{FB}(\mu_0)+\delta} < \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}. \tag{51}$$

This inequality is in contradiction with (48).

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 the project with certainty until time $\underline{t}$, and we can compute the market's posterior belief $\widehat{\mu}_t^1$ for each time $t \geq \underline{t}$ using equation (24). Using $\widehat{\mu}_t^1$, we can then solve for the probability with which the agent continues with the project 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

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

where $\Pi_{\tau=\underline{t}}^{t}(1-\gamma_\tau dt)=(1-\gamma_{\underline{t}}dt)(1-\gamma_{\underline{t}+dt}dt)(1-\gamma_{\underline{t}+2dt}dt)\dots(1-\gamma_{\underline{t}+ndt}dt)$ 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

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

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 (from μ_0) 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.