

DISCUSSION PAPER SERIES

DP12569

EXPERIMENTING WITH CAREER CONCERNS

Marina Halac and Ilan Kremer

**FINANCIAL ECONOMICS and
INDUSTRIAL ORGANIZATION**



EXPERIMENTING WITH CAREER CONCERNS

Marina Halac and Ilan Kremer

Discussion Paper DP12569

Published 05 January 2018

Submitted 05 January 2018

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **FINANCIAL ECONOMICS and INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Marina Halac and Ilan Kremer

EXPERIMENTING WITH CAREER CONCERNS

Abstract

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

JEL Classification: C73, D83, G21

Keywords: strategic experimentation, career concerns, Dynamic games, private learning, banks, bad loans

Marina Halac - mhalac@columbia.edu
Columbia University and CEPR

Ilan Kremer - ikremer@huji.ac.il
The Hebrew University and University of Warwick

Acknowledgements

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

Experimenting with Career Concerns*

Marina Halac[†]

Ilan Kremer[‡]

December 31, 2017

Abstract

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

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

[†]Columbia University and CEPR. Email: mhalac@columbia.edu.

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

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

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

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

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

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

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

In our model, a manager decides at each time whether to continue to invest in a given project or abandon it. The quality of the project can be either good or bad and is initially unknown to all parties. The manager cares not only about the payoffs from the project but also about the market’s perception of the project’s quality. The market’s perception is based on the publicly observable actions of the manager, namely whether he continues or not with the project. The manager learns about the project’s quality from privately observed lump sum payoffs that arrive at random times. We contrast two scenarios. Our main focus is on a bad news setting, where the manager learns from the arrival of a negative payoff that indicates that the project is bad (and thus expected to generate losses). Here “no news is good news”: over time, in the absence of a negative payoff, the manager becomes more optimistic about the quality of the project. As a benchmark we examine a good news setting, where the manager learns from the arrival of a positive payoff that indicates that the project is good (and thus expected to generate positive profits). In this case “no news is bad news”: the manager becomes more pessimistic as time passes and a positive payoff does not arrive.⁴

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.⁵ We show that the manager’s career concern generates an inefficiency: the manager runs the project for too long relative to the first-best solution that maximizes the expected payoff from the project. In the setting in which learning about project quality is through the arrival of bad news, the manager follows a pure strategy: he abandons the project if and only if bad news arrive before a given date t^* . As the manager continues with the project before reaching t^* , the market’s belief that

³As [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.

⁴We assume that the arrival of news is not verifiable. As discussed in [Section 4](#), however, our main qualitative results continue to hold if the manager is able to verifiably disclose news.

⁵We study perfect Bayesian equilibria satisfying a version of Divinity. See [Section 1](#).

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.⁶ Distortions in this setting take the form of the manager keeping the project after learning that it is bad, namely when bad news first arrive after time t^* . When the prior probability of a good project increases, t^* decreases, meaning that the manager is even less likely to quit following bad news. We show that this higher tendency not to terminate bad projects more than compensates for the fact that a bad project is less likely, so the overall distortion increases when expected quality rises. This result contrasts with what we find in the good news setting. Distortions in that setting take the form of the manager keeping the project after enough time has passed and good news have not arrived, namely after he has reached the efficient stopping time without good news. When the prior probability of a good project increases, the manager is more likely to keep the project for a longer period of time; however, the efficient stopping time also increases, so good news are more likely to arrive prior to this time. We show that as a consequence, distortions can decrease with expected quality in the good news setting.

Our second main result concerns the effects of information. Suppose that it is possible to release a public signal at the beginning of time that makes the market and manager's common prior on the project more precise. This signal allows for a better assessment of the quality of projects and thus always weakly increases first-best welfare. The effects on equilibrium welfare depend on how the signal affects distortions. Naturally, if the signal is perfect, it eliminates distortions, as the manager's actions cannot influence the market's belief when the project's quality is known. We show however that in the bad news setting, the effects of information are non-monotonic: a sufficiently imperfect signal increases the distortion relative to first best and reduces welfare.⁷ Intuitively, a high signal realization increases the prior that the project is good and (by the result described above) increases the distortion relative to first best, whereas a low signal realization reduces the prior and thus reduces the distortion. We find that distortions are convex in the prior, and therefore the former effect dominates when the

⁶We describe here our results for “intermediate” parameter values, under which the equilibrium features an interior time $t^* \in (0, \infty)$. See [Section 2](#) for details.

⁷These results extend to public signals that arrive after time 0. See [Section 4](#).

signal is sufficiently imperfect.⁸ Moreover, such a signal leaves first-best welfare unchanged, so the increase in distortions causes overall welfare to go down. These results are in contrast with what we find in the good news setting: releasing a public signal about project quality at the beginning of time always increases welfare when the manager learns through good news.

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

Related literature. There is a large literature on career concerns. One strand of this literature, in the tradition of [Holmström \(1999\)](#), studies moral hazard models in which career concerns are beneficial because they incentivize agents to exert effort. These are models where outcomes are observable but actions are not.⁹ Our paper fits into a different strand of this literature, in which career concerns are detrimental because they lead to perverse incentives. Here actions are observable but outcomes are not. A seminal paper is [Prendergast and Stole \(1996\)](#), where an agent has private information about his ability to understand the state of the world and distorts his decisions over time to look as a fast learner. Related issues are studied in [Scharfstein and Stein \(1990\)](#), [Zwiebel \(1995\)](#), [Majumdar and Mukand \(2004\)](#), [Prat \(2005\)](#), [Ottaviani and Sørensen \(2006\)](#), and [Aghion and Jackson \(2016\)](#).¹⁰

Within finance, as mentioned above, [Rajan \(1994\)](#) studies a static model in which a career-

⁸As elaborated in [Section 2](#), the convexity results from the fact that an increase in the prior brings the distortion forward in time while a reduction in the prior postpones it, and discounting is convex.

⁹[Bonatti and Hörner \(2015\)](#) consider a version of [Holmström’s](#) model with exponential learning.

¹⁰Reputational 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.

concerned bank manager chooses whether to implement a liberal credit policy that makes a bad loan less visible. The paper assumes that the bank 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 4](#). [Makarov and Plantin \(2015\)](#) show that a fund manager will want to take on hidden tail risk when concerned with investors’ perception of his ability to generate excess returns above a fair compensation for risk. A number of papers point out that managers with stock-based compensation may still have a conflict of interest with shareholders (see [Bond, Edmans and Goldstein, 2012](#), Section 3 for a survey), although only a few study how this conflict can result in actions directly aimed at concealing or revealing information. Among these, [Benmelech, Kandel and Veronesi \(2010\)](#) show that to prevent stock price reductions, managers may use suboptimal investment policies that conceal slowdowns in the firm’s growth opportunities. [Frenkel \(2017\)](#) examines a dynamic bargaining model in which managers with stock-sensitive compensation distort the price and timing of over-the-counter asset sales.

Our model is more closely related to those in [Grenadier, Malenko and Strebulaev \(2014\)](#), [Bobtcheff and Levy \(2017\)](#), and [Thomas \(2016\)](#), all of which consider environments with exponential learning.¹¹ [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.¹² [Bobtcheff and Levy \(2017\)](#) analyze a real option model in which a cash-constrained agent may learn bad news prior to investing. The agent wants to convey that his privately known learning intensity is high to raise capital more cheaply, and this can lead to hurried or delayed investment. [Thomas \(2016\)](#) studies a career concerns problem similar to ours, in which an agent learns about a project over time and can choose to abandon it. In her model, however, successes and failures are public, the agent learns privately from partially informative signals without cash-flow consequences, and the agent receives a reputation payoff only when the project succeeds, fails, or is terminated. The focus is also different: [Thomas](#) examines the conditions under which efficiency obtains,¹³ whereas we study the pattern of distortions, specifically how distortions vary with expected project quality and information about quality.

More broadly, our paper is related to a sizable literature on exponential-bandit learning,

¹¹Also 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.

¹²The paper shows that if a public shock forces some agents to stop, then others will blend with the crowd and stop strategically at the same time. Related papers of strategic delay include [Acharya, DeMarzo and Kremer \(2011\)](#) and [Grenadier and Malenko \(2011\)](#). See also [Gratton, Holden and Kolotilin \(2017\)](#).

¹³In her model, whether the first best is implementable depends on the intensity of the agent’s reputational concern and the sign and relative informativeness of public and private news.

including the seminal work of Keller, Rady and Cripps (2005). Most of this literature studies learning through good news, but there are exceptions: in addition to the papers described above, Bonatti and Hörner (2017) and Keller and Rady (2015) consider bad news learning, and Che and Hörner (2017), Frick and Ishii (2015), and Khromenkova (2015) compare good news and bad news learning in various contexts.¹⁴ Some articles study how information disclosure affects experimentation, although this information is typically about outcomes rather than the underlying state as in our paper. In a career concerns setting, see Pei (2015) and, outside the exponential bandit framework, Hörner and Lambert (2016).¹⁵

1 Model

Players and actions. There are an agent and a market. Time is continuous, the horizon is infinite, and the discount rate is $r > 0$. The agent has a project and, at each time $t \geq 0$, decides whether to continue working on the project, $a_t = 1$, 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 $\tau > t$. (Given our solution concept, this irreversibility assumption will be without loss; see fn. 20.)

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

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

The market only observes the agent’s decision at each point of whether to continue or stop working on the project; the realized payoffs from the project are nonverifiable and privately

¹⁴See also Abreu, Milgrom and Pearce (1991) and Board and Meyer-ter-Vehn (2012).

¹⁵Rosenberg, Solan and Vieille (2007) study a model of social experimentation in which agents observe their opponents’ actions but not their opponents’ payoffs, and the decision to abandon the risky project is irreversible. Quah and Strulovici (2009) examine the comparative statics of optimal stopping time problems in a general framework.

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 $\hat{\mu}_t$ 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 select and successfully work on new projects. A career-concerned agent will therefore want the market to believe that his project is good rather than bad.

Following [Rajan \(1994\)](#) and [Prendergast and Stole \(1996\)](#), we take the agent’s payoff at each time $t \geq 0$ to be a weighted sum of the project payoff and the reputation payoff the agent receives from the market’s perception, $\hat{\mu}_t R$, 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. For example, like money managers in [Dasgupta and Piacentino \(2015\)](#), banks differ in their ability to pick good investments, and they benefit from having a reputation for being good pickers as this allows them to attract more capital. Entrepreneurs and venture capital funds have similar career concerns, especially since they must repeatedly raise money from investors, as emphasized by [Baker \(2000\)](#). Given a sequence of actions $\{a_t\}_{t \geq 0}$ and market beliefs $\{\hat{\mu}_t\}_{t \geq 0}$, the agent’s expected payoff at time zero is

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

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.¹⁶

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 4, 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 4, 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 4, 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 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 the prior.

¹⁶One may posit that society also values acquiring information, beyond the effects on the project’s profits; see Section 4 for a discussion.

We solve for the perfect Bayesian equilibria of this game.¹⁷ As is standard in signaling games, we use a refinement to rule out equilibria that can arise only due to unreasonable beliefs off the equilibrium path. We adopt a version of the Divinity refinement introduced by [Banks and Sobel \(1987\)](#), requiring that off-the-equilibrium-path beliefs place relatively more weight on types that gain more from deviating from a fixed equilibrium.¹⁸ 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. Moreover, note that for any given market beliefs, continuing with the project at a time $t > 0$ is always weakly more attractive to an agent type who has not failed/has succeeded by t than to one who has failed/has not succeeded by t . We therefore require that off-the-equilibrium-path beliefs satisfy the following property, often referred to as *belief monotonicity*: if continue ($a_t = 1$) is played with zero probability at time $t > 0$, then the market’s belief that the agent has not failed/has succeeded weakly increases upon observing that he continues at t . This refinement is similar to those used in other continuous-time signaling games such as [Daley and Green \(2012\)](#), [Gul and Pesendorfer \(2012\)](#), and [Strebulaev, Zhu and Zryumov \(2016\)](#), and it ensures that the market’s belief that the project is good weakly increases upon observing a deviation to continuing with the project. Naturally, as we will see, Bayes’ rule yields that the market’s belief that the project is good weakly increases upon observing an on-path decision to continue; belief monotonicity extends this property to off the equilibrium path.

Without belief monotonicity, one could construct equilibria in which the market’s off-the-equilibrium-path beliefs “punish” the agent for deviating to continuing with the project — by assigning a sufficiently high belief to the agent having failed/not succeeded and thus to his project being bad — and in which therefore the game can be forced to end at arbitrary times. Such equilibria would fail forward induction reasoning, as the agent types who have not failed/have succeeded, and whose project is thus more likely to be good, are precisely the types who have relatively more to gain from continuing with the project.

We will indicate where we use belief monotonicity in our discussions in the text and in the proofs in the Appendix. From now on, we refer to perfect Bayesian equilibria satisfying this refinement as simply equilibria.

¹⁷We 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.

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

2 Bad news

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

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

The evolution of this belief is governed by

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

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

2.1 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

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

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

2.2 Career concerns

Consider now the setting with $R > 0$, i.e. where the agent cares about both the payoff from the project and the market's belief about 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. Since the expected payoff from continuing with the project is 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, it must be that an agent who has failed by t stops with certainty by t . Hence, the market's belief that the agent has not failed — and thus 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 belief monotonicity 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.¹⁹ 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,²⁰

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

¹⁹Plainly, 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 (weakly) higher than the market belief after the agent stops at t .

²⁰Note that [Lemma 1](#) is valid even if stopping is reversible. To see the idea, suppose that an agent who has not failed by t were to stop at t and restart the project at $t' > t$. By the arguments above, an agent who has failed by t must stop by t with certainty, and the market belief upon observing that the agent continues through time t' must be strictly lower than that upon observing that the agent restarts at t' after having stopped at t (otherwise an agent who has not failed would not be willing to stop and restart). However, given belief monotonicity, this construction requires that a positive mass of types who have failed by t stop at t and do not restart at t' , which is not possible: these types would prefer either to stop before time t or to restart at t' . Since [Corollary 1](#) follows from [Lemma 1](#), we obtain that an agent who chooses to stop does not restart, implying that the irreversibility of stopping is without loss.

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 $\hat{\mu}_t^1$. 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 $\hat{\mu}_t^1 = \mu_t$, 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, $\hat{\mu}_t^1$ 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 . Hence, the agent's expected payoff from continuing is

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

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 $\lim_{t \rightarrow \infty} \mu_t = 1$, this requires

$$R \leq -(x - \lambda_B). \quad (6)$$

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.

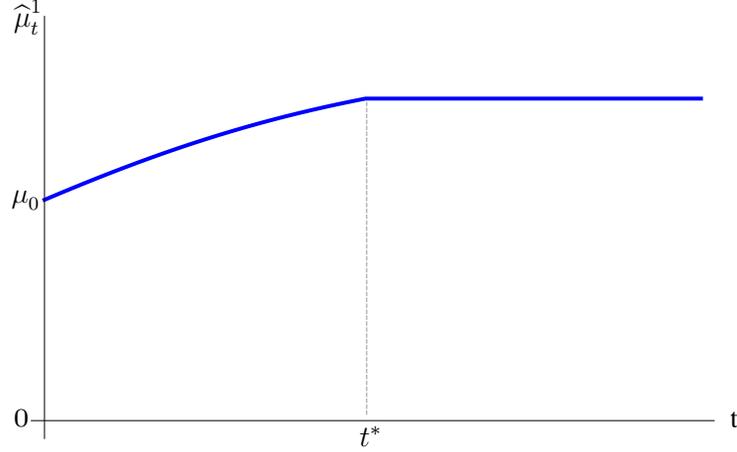


Figure 1: Beliefs and threshold time in the equilibrium of the bad news setting. Parameters are $\mu_0 = 0.6, x = 0.75, \lambda_B = 1.8, R = 1.2,$ and $r = 1$.

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*

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

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 $\hat{\mu}_t^1 = \mu_t$ if the agent has not stopped by t and $t \leq t^*$, $\hat{\mu}_t^1 = \mu_{t^*}$ 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 with threshold time t^* is

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

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 (in which the agent never works on the project) does not exist. Intuitively, if a no-work equilibrium exists, it exists when the market’s beliefs are such that the agent’s reputation benefit from starting the project at time 0 is minimized. These beliefs 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.²¹ 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 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.²²

Since a bank that keeps the debt until time t^* continues refinancing it regardless of its value, the market learns no information about the quality of the bank’s loans after this time. 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

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

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

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

²²While 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 4](#) for a discussion.

banks are accumulating losses.

2.3 Expected quality and information

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

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

As $t^* \rightarrow \infty$, $S_0(t^*) \rightarrow S_0^{FB}$, 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^* declines.

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

What are the welfare effects of an increase in the expected quality of projects? Suppose the prior probability of a good project, μ_0 , increases. First-best welfare then naturally goes up. Moreover, since the agent's career concern distorts actions away from first best only when the project is bad, an increase in μ_0 has a direct effect of decreasing the distortion for any fixed threshold time t^* . However, an increase in μ_0 also reduces the equilibrium threshold time t^* : when expected project quality is higher, the agent's reputation loss from ending the project at any time t increases, and the threshold time t^* after which the agent prefers to continue upon failing decreases. We show, in fact, that the agent's higher tendency not to terminate bad projects more than compensates for the fact that a bad project is less likely: 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 (Expected quality in bad news setting). *Suppose parameters $\{\mu_0, x, \lambda_B, R, r\}$ satisfy $-(x - \lambda_B)/\mu_0 > R > -(x - \lambda_B)$ (so the equilibrium features $t^* \in (0, \infty)$) and consider changes in μ_0 that preserve this property and [Assumption 1](#). An increase in μ_0 increases welfare but it also increases the distortion relative to first best.*

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

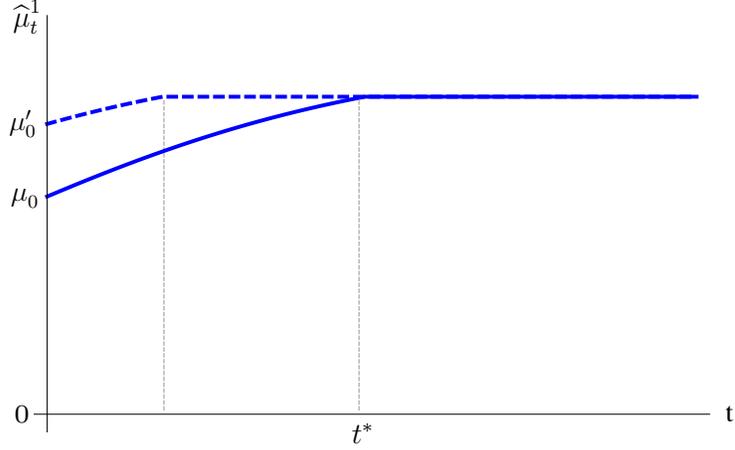


Figure 2: Effects of an increase in μ_0 in the equilibrium of the bad news setting. Parameters are the same as in Figure 1, with $\mu_0' = 0.8$.

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, 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 empirical findings of [Schularick and Taylor \(2012\)](#), one may say that banks plant the seed for the next crisis during boom periods.

Would information about project quality ameliorate the welfare distortions due to career concerns? After all, it is because the quality of the project is uncertain that career concerns lead to distorted behavior. Suppose that it is possible to release a public signal at time 0 that refines the agent and market’s common prior on the project, μ_0 . Absent career concerns, the signal either keeps welfare unchanged — if it does not affect the decision of whether to start the project at time 0 — or increases welfare — if it does affect this start decision. However, when the agent is career-concerned, the signal also affects distortions: as implied by [Proposition 2](#), a high realization of the signal (i.e. a realization that increases μ_0) may increase the 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 . Let each realization of the signal be unconditionally equally likely, with the high realization increasing the prior to $\mu_0^h > \mu_0$ and the low realization reducing the prior to $\mu_0^\ell < \mu_0$, so that $\mu_0 = \frac{1}{2}(\mu_0^h + \mu_0^\ell)$. Building on our discussion of [Proposition 2](#), the high realization reduces the time t^* that it takes for the market’s belief to reach the posterior μ_{t^*} (given by [\(7\)](#)), whereas the low realization increases t^* . Once μ_{t^*} is reached, the distortion due to the agent not stopping upon failure is independent of μ_0 ; what matters is with which probability μ_{t^*} is reached and how heavily the expected future losses at t^* are discounted. Given a prior μ_0 , the probability of reaching μ_{t^*} is $\frac{\mu_0}{\mu_{t^*}}$,²³ which is linear in μ_0 . Hence, the expected probability of the distortion is unchanged with the public signal: $\frac{1}{2} \left(\frac{\mu_0^h}{\mu_{t^*}} + \frac{\mu_0^\ell}{\mu_{t^*}} \right) = \frac{\mu_0}{\mu_{t^*}}$. However, if the amount of time that it takes to reach μ_{t^*} starting from a prior μ_0 is $t^*(\mu_0)$, then discounting of the distortion at μ_{t^*} is given by $e^{-rt^*(\mu_0)}$, which as implied by [\(2\)](#) is a

²³This follows from the fact that the market’s posterior at time t^* is either μ_{t^*} (if no failure occurs by t^*) or 0 (if a failure occurs by t^*), and the prior μ_0 is correct.

convex function of μ_0 .²⁴ That is, discounting of the expected losses at μ_{t^*} decreases when μ_0 goes up by more than it increases when μ_0 goes down by the same magnitude. As a result, $\frac{1}{2}(e^{-rt^*(\mu_0^h)} + e^{-rt^*(\mu_0^l)}) > e^{-rt^*(\mu_0)}$, implying that the expected discounted losses increase, and thus welfare decreases, when the public signal is available.

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.

[Section 4](#) discusses an extension in which a public signal of project quality may be released at a time $t > 0$. While new considerations come into play when information arrives after time 0, the effects described above are still present, and imply that this information can be detrimental to welfare if (and only if) sufficiently imperfect.

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.

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

3 Good news

We contrast the bad news setting of [Section 2](#) with a good news setting in which the agent learns about project quality 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:

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

The evolution of this belief is governed by

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

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

3.1 First best

Suppose $R = 0$, so the agent does not have a career concern and maximizes social welfare. 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

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

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

3.2 Career concerns

Suppose now $R > 0$, so the agent cares not only about the payoff from the project but also about the market's 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 $\hat{\mu}_t^1$. 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 $\bar{t} \geq \underline{t}$ such that the agent runs the project until \underline{t} with certainty and mixes between stopping and continuing at each $t \in [\underline{t}, \bar{t}]$ absent success. The agent continues forever following a success at any time, and he stops with certainty by \bar{t} if \bar{t} is finite and he has not succeeded by this time. The agent's stopping policy is such that the market's belief $\hat{\mu}_t^1$ evolves continuously at all $t > 0$. The threshold times satisfy $\underline{t} < \infty$ if and only if $R < -x/\mu_0$ and $\bar{t} < \infty$ 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 $\bar{t} \geq \underline{t}$. The agent continues working on the project until \underline{t} with certainty. From \underline{t} on, the agent implements a 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 $t \rightarrow \infty$. 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 $\hat{\mu}_t^1 = \mu_0$ for $t \leq \underline{t}$, $\hat{\mu}_t^1 \in (\mu_0, 1)$ for $t \in (\underline{t}, \bar{t})$, and $\hat{\mu}_t^1 = 1$ 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

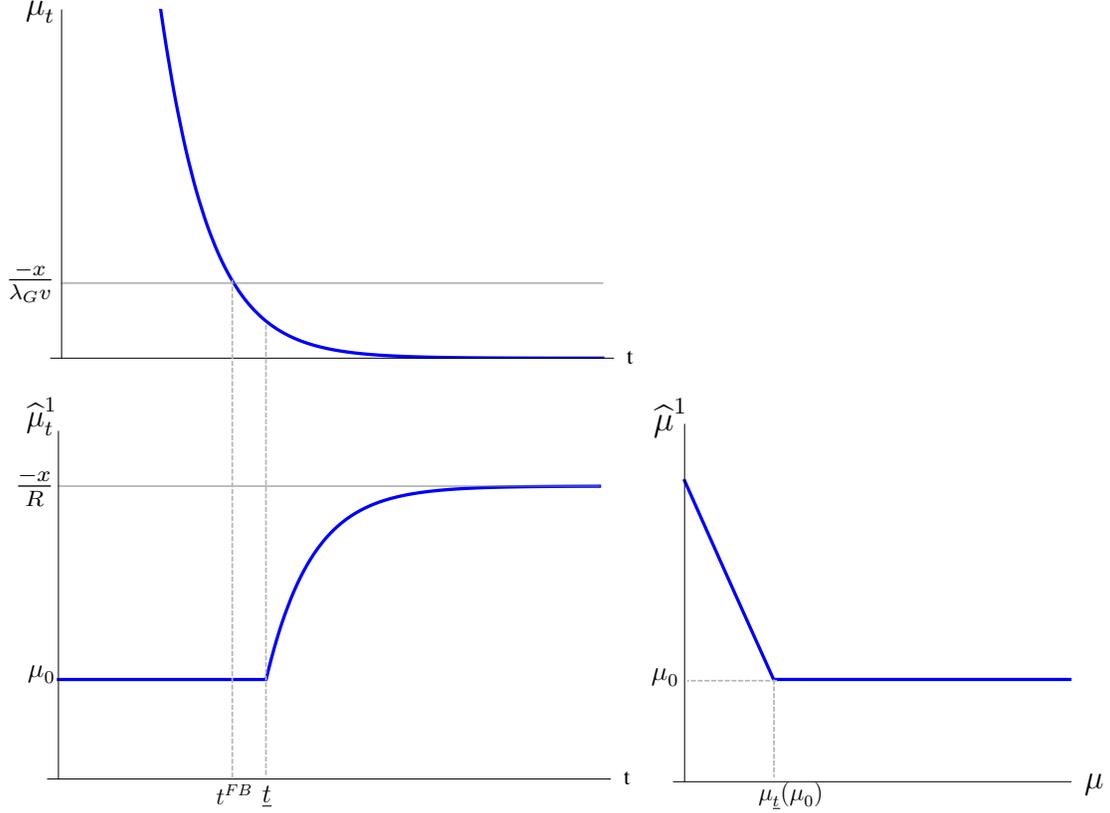


Figure 3: Beliefs and threshold times in the equilibrium of the good news setting. Parameters are $\mu_0 = 0.5$, $x = -0.8$, $\lambda_G = 6$, $R = 0.9$, and $r = 1$.

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, $\hat{\mu}^1$, 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 $\hat{\mu}^1 = \mu_0$; 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 $\hat{\mu}_t^1$ 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 $[t, t + dt]$:

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

where $V_t = v + R \int_t^\infty e^{-r(s-t)} \hat{\mu}_s^1 ds$ 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 pin down \underline{t} and \bar{t} and the market's belief $\hat{\mu}_t^1$. We relegate the details to the Appendix.

A suitable application of this good news setting may be to venture capital, where an entrepreneur learns about the quality of a project from the arrival of a breakthrough.²⁵ 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, $\hat{\mu}_t^1$, must be increasing over time.

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

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

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

²⁵In our model, such a breakthrough cannot be verifiably disclosed to the market. See [Section 4](#) for a discussion of verifiable news.

3.3 Expected quality and information

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

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

Proposition 5 (Expected quality and information in good news setting). *Suppose parameters $\{\mu_0, x, \lambda_G, R, r\}$ satisfy $-x/\mu_0 > R > -x$ (so the equilibrium features $\underline{t} \in (t^{FB}, \infty)$ 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 2](#), 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 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

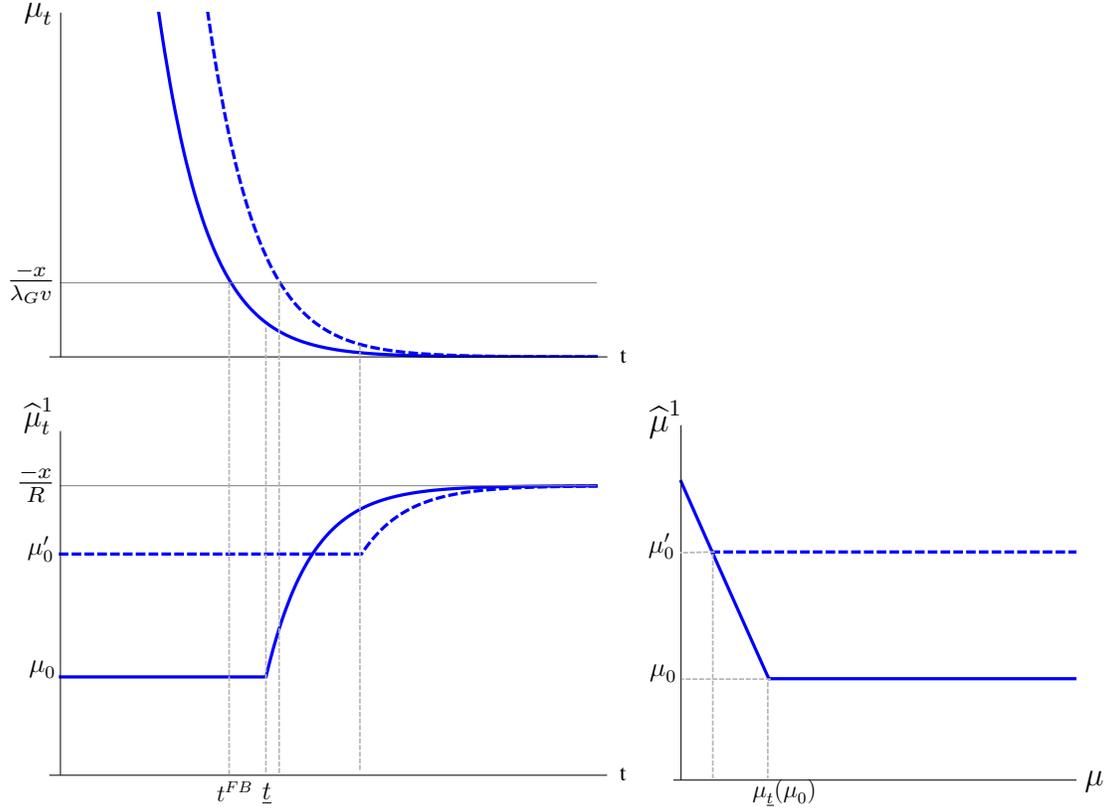


Figure 4: Effects of an increase in μ_0 in the equilibrium of the good news setting. Parameters are the same as in Figure 3, with $\mu'_0 = 0.75$.

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.

4 Discussion

Below we discuss some extensions and variations of the model. We consider each of these extensions separately, taking the model setup and assumptions as otherwise unchanged.

Information. We have studied the effects of releasing public information about project quality at time 0. What happens if this information arrives at a later time? We focus here on the bad news setting; an analogous analysis can be performed for the good news setting but it would take us too far afield.

Take the bad news setting with intermediate parameters as in [Proposition 3](#) and suppose a public signal of project quality may be released at a time $t > 0$. Depending on parameters, the agent may want to wait for the signal before starting the project, in which case the signal would have the same effects as when released at time 0 plus a delay. To rule out this possibility and focus the discussion, suppose the signal can only arrive while the agent is working, and it arrives via a Poisson process. That is, conditional on the agent working and the signal not having arrived by time t , the distribution of arrival times $s \geq t$ is independent of t . This ensures that the equilibrium will still have a threshold structure as in [Proposition 1](#).

Consider first a public signal that reveals the project's quality perfectly. At the time that the signal arrives, it makes the agent stop if the project is bad (and continue if the project is good), therefore increasing welfare. Now suppose the signal has not arrived by time t . If the agent has failed by t , he knows that a signal in the future would make him stop. The presence of a signal is then analogous to an increase in the discount rate,²⁶ which, as shown in [Proposition 1](#), has no implications for the agent's behavior (specifically, the equilibrium threshold time t^* in (7) is independent of r). If the agent has not failed by t , a signal may increase the market's belief that the project is good or make him stop; hence, the possibility

²⁶Note that stopping yields the agent a zero continuation payoff.

of a signal can only increase the agent's incentives to continue at t . It follows that a perfect signal affects the agent's behavior only when it arrives, and its effect on welfare is positive.

Suppose next that the public signal is sufficiently imperfect in the way defined in [Proposition 3](#). We begin by considering the effects of the signal at the time that it arrives. These effects depend on the arrival time. Take first a signal that arrives at a time $t > t^*$. If the signal realization is high (i.e. it increases the market's belief that the project is good), then its arrival has no effect on the behavior of the agent, who continues with the project. If the signal realization at t is low, then its arrival causes the agent to stop with positive probability if he has failed by t (and to continue otherwise). This follows from the fact that the agent who has failed is indifferent between continuing and stopping prior to the signal arrival, and, ceteris paribus, a low signal lowers the market's belief conditional on the agent continuing. Therefore, we obtain that the arrival of an imperfect signal at a time $t > t^*$ increases welfare.

Consider next an imperfect signal that arrives at a time $t < t^*$. Note that the agent has not failed by t and the agent and market's beliefs coincide at this time. The effects of the signal are thus analogous to those of a signal at time 0: a high signal realization brings the distortion forward in time (i.e. decreases t^*) while a low signal realization postpones the distortion (i.e. increases t^*), and the convexity of discounting implies that the distortion increases in expectation. Hence, in this case, the arrival of the imperfect signal lowers welfare.

Finally, consider the effects of an imperfect signal prior to its arrival. Note that regardless of his private history, the agent assigns positive probability to both low and high signal realizations. The agent's continuation payoff following a low realization is no smaller than that from stopping, and, as noted, the possibility that he may be forced to stop in the future does not affect the agent's behavior. Thus, the possibility of a low signal does not lower the agent's incentives to continue with the project. Furthermore, the possibility of a high signal increases the agent's continuation payoff and thus his incentives to continue with the project. It follows that the equilibrium threshold time t^* , and thus welfare, decline in anticipation of the signal.

In sum, the results of [Proposition 3](#) extend to information that arrives after time 0. This information is welfare-improving if perfect, but may be detrimental if sufficiently imperfect. The latter will be true for Poisson public signals with a relatively high arrival rate, which are more likely to arrive before the equilibrium threshold time t^* . It will also be true if r is large, so the negative welfare effects of a signal arriving at a time $t < t^*$ are less heavily discounted relative to the positive welfare effects of a signal arriving at a time $t > t^*$.

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, $\theta \in \{\underline{\theta}, \bar{\theta}\}$, and the state of the economy, $\omega \in \{\underline{\omega}, \bar{\omega}\}$. As in [Rajan \(1994\)](#), assume that the project is good if $\theta = \bar{\theta}$

and $\omega = \bar{\omega}$, 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 ($\theta = \bar{\theta}$) and $q_0 \in (0, 1)$ on a high state of the economy ($\omega = \bar{\omega}$). 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 $\hat{p}_t R$ (instead of $\hat{\mu}_t R$), where \hat{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 2](#), 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. 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 2](#) and [Section 3](#)). 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 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 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 notes that bank managers will have a higher incentive to hide bad loans during good times. We show that distortions will as a result be larger during good times, despite bad loans being less likely, and study the effects of releasing information about the quality of credit. By comparing

bad news and good news learning, our analysis reveals that dynamic aspects are important: the pattern of distortions in the banking industry owes to the fact that bank managers 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.

A main difference relative to our baseline model is that the value of the project depends on both the quality of the project and the agent’s actions. Thus, the market’s perception at any time will be a function of both its belief about quality and its (correct) conjecture of the agent’s equilibrium behavior. This will imply a reputation benefit term in the agent’s payoff function that is not linear in the market’s belief as in our baseline model, but rather some function $Rf(\hat{\mu}_t)$ where $f(\cdot)$ is increasing. Despite these differences, one can show that equilibrium constructions analogous to ours, and therefore our (qualitative) comparative statics, remain valid under this alternative specification. In particular, in the bad news setting, the equilibrium still takes the form of a threshold (albeit with a different threshold time t^*), and the non-linearity of the reputation benefit does not affect the convexity argument that is behind our information result in [Proposition 3](#).

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.

Verifiable news. We have assumed that the agent learns about the quality of the project from privately observed outcomes which are not verifiable. Suppose instead that the agent could verifiably disclose the arrival of news to the market (whereas the absence of news remains non-verifiable). Clearly, nothing would change in the bad news setting: the agent would have no incentive to disclose failure unless he decides to stop working on the project, at which point the market learns in equilibrium that the agent has failed. In the good news setting, on the other hand, the agent would choose to disclose a success as soon as it arrives. This setting thus becomes equivalent to one in which success is observable to the market, and the agent simply has an objective that differs from the first best due to his career concern. Absent private information, however, the agent's career concern has no effect on his behavior, as the agent's expected reputation gain from continuing at any point in time is zero (by the martingale property of beliefs). That is, the agent in equilibrium would follow the first best policy, implying no distortions and a positive value of information for all parameters. Consequently, we find that our main results on the pattern of distortions in the bad news setting, and its contrast with the good news setting, continue to hold when news are verifiable.

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.
- Baker, Malcolm**, “Career Concerns and Staged Investment: Evidence from the Venture Capital Industry,” 2000. Working Paper.
- Banks, Jeffrey S. and Joel Sobel**, “Equilibrium Selection in Signaling Games,” *Econometrica*, 1987, 55 (3), 647–661.
- Bar-Isaac, Heski**, “Reputation and Survival: Learning in a Dynamic Signalling Model,” *The Review of Economic Studies*, 2003, 70 (2), 231–251.
- Benmelech, Efraim, Eugene Kandel, and Pietro Veronesi**, “Stock-Based Compensation and CEO (Dis)Incentives,” *Quarterly Journal of Economics*, 2010, 125 (4), 1769–1820.
- Bernanke, Ben S.**, “Stress Testing Banks: What Have We Learned?,” 2013. Board of Governors of the Federal Reserve System, Speech.
- Bertrand, Marianne and Sendhil Mullainathan**, “Are CEOs Rewarded for Luck? The Ones Without Principals Are,” *Quarterly Journal of Economics*, 2001, 116 (3), 901–932.
- Board, Simon and Moritz Meyer-ter-Vehn**, “Reputation for Quality,” *Econometrica*, 2012, 81 (6), 2381–2462.
- Bobtcheff, Catherine and Raphaël Levy**, “More Haste, Less Speed: Signalling through Investment Timing,” *American Economic Journal: Microeconomics*, 2017, *Forthcoming*.
- Bonatti, Alessandro and Johannes Hörner**, “Career Concerns with Exponential Learning,” *Theoretical Economics*, January 2015, 12, 425–475.
- **and –**, “Learning to Disagree in a Game of Experimentation,” *Journal of Economic Theory*, 2017, 169, 234–269.

- Bond, Philip, Alex Edmans, and Itay Goldstein**, “The Real Effects of Financial Markets,” *Annual Review of Financial Economics*, 2012, 4 (1), 339–360.
- Caballero, Ricardo, Takeo Hoshi, and Anil K Kashyap**, *American Economic Review*, 2008, 98 (5), 1943–1977.
- Che, Yeon-Koo and Johannes Hörner**, “Optimal Design for Social Learning,” *Quarterly Journal of Economics*, 2017, *Forthcoming*.
- Daley, Brendan and Brett Green**, “Waiting for News in the Market for Lemons,” *Econometrica*, 2012, 80 (4), 1433–1504.
- Dang, Tri Vi, Gary Gorton, and Bengt Holmström**, “Ignorance, Debt and Financial Crises,” 2015. Working Paper.
- , – , and – , “The Information Sensitivity of a Security,” 2015. Working Paper.
- Dasgupta, Amil and Giorgia Piacentino**, “The Wall Street Walk when Blockholders Compete for Flows,” *Journal of Finance*, 2015, 70 (6).
- Ely, Jeffrey and Juuso Välimäki**, “Bad Reputation,” *Quarterly Journal of Economics*, 2003, 118 (3), 785–814.
- Frenkel, Sivan**, “Dynamic Asset Sales with a Feedback Effect,” 2017. Working Paper.
- Frick, Mira and Yuhta Ishii**, “Innovation Adoption by Forward-Looking Social Learners,” 2015. Working Paper.
- 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,” *Review of Economic Studies*, 2017, *Forthcoming*.
- Grenadier, Steven R. and Andrey Malenko**, “Real Options Signaling Games with Applications to Corporate Finance,” *The Review of Financial Studies*, 2011, 24 (12).
- , – , and **Ilya A. Strebulaev**, “Investment Busts, Reputation, and the Temptation to Blend in with the Crowd,” *Journal of Financial Economics*, 2014, 111, 137–157.

- Gul, Faruk and Wolfgang Pesendorfer**, “The War of Information,” *Review of Economic Studies*, 2012, *79*, 707–734.
- Hirtle, Beverly and Andreas Lehnert**, “Supervisory Stress Tests,” 2014. Federal Reserve Bank of New York, Staff Report No. 696.
- Holmström, Bengt**, “Managerial Incentive Problems: A Dynamic Perspective,” *Review of Economic Studies*, January 1999, *66* (1), 169–182.
- Hörner, Johannes and Nicolas Lambert**, “Motivational Ratings,” 2016. Working Paper.
- Keller, Godfrey and Sven Rady**, “Breakdowns,” *Theoretical Economics*, 2015, *10*, 175–202.
- , – , and **Martin Cripps**, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 2005, *73* (1), 39–68.
- Khromenkova, Daria**, “Collective Experimentation with Breakdowns and Breakthroughs,” 2015. Working Paper.
- Majumdar, Sumon and Sharun W. Mukand**, “Policy Gambles,” *American Economic Review*, 2004, *94* (4), 1207–1222.
- Makarov, Igor and Guillaume Plantin**, “Rewarding Trading Skills without Inducing Gambling,” *Journal of Finance*, 2015, *70* (3), 925–962.
- Morris, Stephen**, “Political Correctness,” *Journal of Political Economy*, 2001, *109* (2), 231–265.
- Ottaviani, Marco and Peter Sørensen**, “Professional Advice,” *Journal of Economic Theory*, 2006, *126*, 120–142.
- Oyer, Paul**, “Why Do Firms Use Incentives That Have No Incentive Effects?,” *Journal of Finance*, 2004, *59* (4), 1619–1649.
- Peek, Joe and Eric S. Rosengren**, “Unnatural Selection: Perverse Incentives and the Misallocation of Credit in Japan,” *American Economic Review*, 2005, *95* (4), 1144–1166.
- Pei, Harry Di**, “Reputation with Strategic Information Disclosure,” 2015. Working Paper.
- Prat, Andrea**, “The Wrong Kind of Transparency,” *American Economic Review*, 2005, *95* (3), 862–877.

- Prendergast, Canice and Lars Stole**, “Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning,” *Journal of Political Economy*, 1996, *104* (6), 1105–1134.
- Quah, John and Bruno Strulovici**, “Comparative Statics, Informativeness, and the Interval Dominance Order,” *Econometrica*, 2009, *77* (6), 1949–1992.
- Rajan, Raghuram**, “Why Bank Credit Policies Fluctuate: A Theory and Some Evidence,” *Quarterly Journal of Economics*, 1994, *109*, 399–441.
- Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille**, “Social Learning in One-Arm Bandit Problems,” *Econometrica*, 2007, *75* (6), 1591–1611.
- Scharfstein, David and Jeremy Stein**, “Herd Behavior and Investment,” *American Economic Review*, 1990, *80* (3), 465–479.
- Schularick, Moritz and Alan M. Taylor**, “Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008,” *American Economic Review*, 2012, *102* (2), 1029–1061.
- Sekine, Toshitaka, Keiichiro Kobayashi, and Yumi Saita**, “Forbearance Lending: The Case of Japanese Firms,” *Monetary and Economic Studies*, 2003, *21* (2), 69–92.
- Strebulaev, Ilya, Haoxiang Zhu, and Pavel Zryumov**, “Dynamic Information Asymmetry, Financing, and Investment Decisions,” 2016. Working Paper.
- The Economist**, “The Great Hole of China,” 2014. October 18.
- , “Deleveraging Delayed,” 2015. October 24.
- The Wall Street Journal**, “Skepticism on China’s Nonperforming Loans,” 2013. December 3.
- Thomas, Caroline**, “Career Concerns and Policy Intransigence – A Dynamic Signaling Model,” 2016. Working Paper.
- Townsend, Robert M.**, “Optimal Contracts and Competitive Markets with Costly State Verification,” *Journal of Economic Theory*, 1979, *21*, 265–293.
- Watson, Joel**, “Perfect Bayesian Equilibrium: General Definitions and Illustrations,” January 2016. Working Paper.

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

A Appendix

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

A.1 Proof of [Lemma 1](#)

We prove the lemma by proving two claims:

Claim 1: 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. $\hat{\mu}_{t'}^1 \geq \hat{\mu}_{t'}^0$.

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 $\hat{\mu}_{t'}^1 R/r$: 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 $\hat{\mu}_{t'}^1 R/r$ if the agent continues with the project at t' , and it is equal to $\hat{\mu}_{t'}^0 R/r \leq \hat{\mu}_{t'}^1 R/r$ 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's belief upon observing that the agent continues at t' is off the equilibrium path. However, note that by belief monotonicity, 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 $\hat{\mu}_{t'}^1 \geq \hat{\mu}_{t'}^0$. 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

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

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 $\hat{\mu}_t^0 = 0$ 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 $\hat{\mu}_t^1$ for all $t \geq 0$ and $\hat{\mu}_t^0 = 0$ for $t < t^*$. We now show that the market belief must satisfy $\hat{\mu}_t^0 = 0$ for all $t \geq t^*$. First, note that $\hat{\mu}_t^0$ 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 $\hat{\mu}_t^0 = 0$ for $t < t^*$, it follows that $\hat{\mu}_{t^*}^0 = 0$. 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 $\hat{\mu}_{t^*}^1$ for $t \geq t^*$. Thus, if $\hat{\mu}_{t'}^0 > 0$ for some $t' > t^*$, an agent who fails at t' strictly prefers to stop at that time, yielding another contradiction. The claim follows.

A.3 Proof of [Proposition 2](#)

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

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

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

The distortion relative to first best is

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

The derivative with respect to μ_0 is

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

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, equiva-

lently, 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

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

Substituting with (13) and rearranging terms,

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

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 $x - (1 - \mu_0)\lambda_B \leq 0$, violating [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 reduces to $(1 - \mu_0)r > 0$ and is thus always satisfied.

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

B Internet Appendix

This Internet Appendix contains the proofs of the results stated in [Section 3](#) of the paper.

B.1 Proof of [Proposition 4](#)

Existence. We begin by describing the equilibrium construction and showing existence. Recall that for each time $t \in [\underline{t}, \bar{t}]$, the agent's indifference condition (12) must be satisfied. Writing $\mu_{t+dt} = \mu_t + dt\dot{\mu}_t = \mu_t - \mu_t(1 - \mu_t)\lambda_G dt$ and dividing by dt both sides of the equation and taking dt to zero yields

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

where $V_t = v + R \int_t^\infty e^{-r(s-t)} \widehat{\mu}_s^1 ds$. Additionally, at each $t \in [\underline{t}, \bar{t}]$, the agent must be indifferent (absent success) between stopping and continuing until \bar{t} , where \bar{t} is possibly infinite. At \bar{t} , the agent continues forever if he has succeeded and stops otherwise. That is, for each $t \in [\underline{t}, \bar{t}]$,

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

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

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

Substituting with equation (20) yields

$$\begin{aligned} \dot{\mu}_t R &= r(\mu_r R - x - \mu_t \lambda_G) + \dot{\mu}_t \lambda_G \left(1 - e^{-r(\bar{t}-t)} \right) - rR \widehat{\mu}_t^1 \\ &\quad + e^{-r(\bar{t}-t)} \left\{ \begin{aligned} &\left[\dot{\mu}_t \left(1 - e^{-\lambda_G(\bar{t}-t)} \right) - \mu_t \lambda_G e^{-\lambda_G(\bar{t}-t)} \right] (x + \lambda_G + R) \\ &+ \left[\mu_t \lambda_G e^{-\lambda_G(\bar{t}-t)} - \dot{\mu}_t \left(1 - e^{-\lambda_G(\bar{t}-t)} \right) \right] \mu_{\bar{t}} R \end{aligned} \right\}. \end{aligned}$$

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

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

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 (21) reduce to

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

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

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 (23); 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 (23) 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 $\tau \in (0, \underline{t}]$.

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 (22) 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 (22) also shows that $\widehat{\mu}_t^1 \rightarrow -x/R$ as $t \rightarrow \infty$, where note $-x/R \in (0, 1)$.

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

$$\mu_{\underline{t}} = \frac{\lambda_G + r - \sqrt{\frac{\lambda_G^2(2r + \lambda_G - R) + \lambda_G r[4(\mu_0 R + x) + r - 2R] - r^2 R}{\lambda_G - R}}}{2\lambda_G}. \quad (24)$$

One can verify that $\mu_{\underline{t}}$ is strictly decreasing in R and $\mu_{\underline{t}} \rightarrow 0$ as $R \rightarrow -x/\mu_0$.

The equilibrium defines consistent market beliefs. We now 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

$$\int_0^\infty e^{-rt}(x + \mu_0 \lambda_G + \widehat{\mu}_t^1 R) dt. \quad (25)$$

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, (25) is greater than $\mu_0 R/r$, implying that it is optimal for the agent to start the project at time 0. Next, to show that it is optimal for the agent to continue at $t < \underline{t}$ absent success, it is sufficient to show that the left-hand side of (22) is smaller than its right-hand side for $t < \underline{t}$, or, equivalently,

$$\Psi_t := x + \mu_t(\lambda_G - R) + rR \int_t^\infty e^{-r(s-t)} \widehat{\mu}_s^1 ds \geq 0. \quad (26)$$

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 (26) yields that for $t < \underline{t}$,

$$\dot{\Psi}_t = -\mu_t(1 - \mu_t)\lambda_G(\lambda_G - R) - rR\mu_0 + r^2R \int_t^\infty e^{-r(s-t)} \widehat{\mu}_s^1 ds. \quad (27)$$

Using (26) and (27), note that $\Psi_{t'} = 0$, $\dot{\Psi}_{t'} < 0$, $\Psi_{t''} < 0$, and $\dot{\Psi}_{t''} = 0$ imply

$$-\mu_{t'}(1 - \mu_{t'})\lambda_G(\lambda_G - R) - rR\mu_0 - r[x + \mu_{t'}(\lambda_G - R)] < 0, \quad (28)$$

$$-\mu_{t''}(1 - \mu_{t''})\lambda_G(\lambda_G - R) - rR\mu_0 - r[x + \mu_{t''}(\lambda_G - R)] > 0. \quad (29)$$

However, as claimed above, there exists a unique solution $\mu_{\underline{t}} \in (0, \mu_0)$ to equation (23) when $\widehat{\mu}_{\underline{t}}^1 = \mu_0$, and therefore we cannot have $\mu_{t'}, \mu_{t''} \in (\mu_{\underline{t}}, \mu_0)$. Contradiction.

Finally, the equilibrium threshold time \underline{t} is uniquely pinned down by (9) and (24), and the claim that $\underline{t} > t^{FB}$ follows from the fact that $\mu_{\underline{t}} < \mu^{FB}$. This latter inequality can be verified by comparing $\mu_{\underline{t}}$ given in (24) and μ^{FB} given in (11), taking into account that [Assumption 1](#) and $R < -x/\mu_0$ imply $\lambda_G > R$.

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. Since $\widehat{\mu}_t^1 = 1$ for all $t \geq \bar{t}$, the agent's indifference

condition (19) at \bar{t} implies

$$R(1 - \mu_{\bar{t}}) + x + \mu_{\bar{t}}\lambda_G v = 0, \quad (30)$$

and hence²⁷

$$\mu_{\bar{t}} = \frac{-(x + R)r}{r(\lambda_G - R) + \lambda_G(\lambda_G + x)}. \quad (31)$$

The threshold time \bar{t} is uniquely pinned down by (9) and (31). Note that $\mu_{\bar{t}}$ is strictly decreasing in R , with $\mu_{\bar{t}} \rightarrow \mu^{FB}$ as $R \rightarrow 0$ and $\mu_{\bar{t}} \rightarrow 0$ as $R \rightarrow -x$.

Since μ_t is decreasing over time and $\lambda_G > R$ (by Assumption 1 and $R < -x$), it follows from equation (30) 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 (31) evaluated at $t > \bar{t}$ instead of \bar{t} is strictly negative. The market's belief of $\hat{\mu}_t^1 = 1$ for $t > \bar{t}$ is therefore consistent. Note that the market's belief upon observing that the agent stops at a time $t > \bar{t}$ is off the equilibrium path; we show here existence when this belief satisfies $\hat{\mu}_t^0 = \mu_t$ for all $t > 0$.

Finally, given that (31) uniquely pins down $\mu_{\bar{t}}$ and (together with (9)) \bar{t} , these can be substituted into equation (21) at $t = \underline{t}$ (with $\hat{\mu}_{\underline{t}}^1 = \mu_0$) to solve for $\mu_{\underline{t}}$ and \underline{t} . Combining (9), (21), and (31) yields

$$\begin{aligned} 0 = & r [x + \mu_{\underline{t}}\lambda_G + R(\mu_0 - \mu_{\underline{t}})] + \mu_{\underline{t}}(1 - \mu_{\underline{t}})\lambda_G(\lambda_G - R) \\ & + \mu_{\underline{t}}(1 - \mu_{\underline{t}})\lambda_G \left(\frac{-(x + R)r}{(\lambda_G + x)(r + \lambda_G)} \frac{1 - \mu_{\underline{t}}}{\mu_{\underline{t}}} \right)^{\frac{r}{\lambda_G}} \frac{\lambda_G(R + x)}{r + \lambda_G}, \end{aligned} \quad (32)$$

and one can verify that $\mu_{\underline{t}}$ coincides with (24) in the limit as $R \rightarrow -x$ and $\mu_{\underline{t}} \rightarrow \mu^{FB}$ as $R \rightarrow 0$. To show that $\underline{t} > t^{FB}$, or equivalently $\mu_{\underline{t}} < \mu^{FB}$, note that $\mu_{\underline{t}}$ is continuous in R for all $R \geq 0$, and by (24) we have $\mu_{\underline{t}} < \mu^{FB}$ for $R \geq -x$. Thus, suppose by contradiction that $\mu_{\underline{t}} \geq \mu^{FB}$ for some $R \in (0, -x)$. Then there must exist $R' > 0$ such that $\mu_{\underline{t}}(R') = \mu^{FB}$. More precisely, let $g(R)$ be the right-hand-side of (32) when $\mu_{\underline{t}} = \mu^{FB}$, as a function of R . Algebraic manipulations yield

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

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

²⁷Equation (31) can equivalently be derived from (21) at \bar{t} .

differentiating $g(R)$ gives

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

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

Lastly, consider the case in which $R > -x/\mu_0$. Suppose the market beliefs satisfy $\hat{\mu}_t^1 = \mu_0$ and $\hat{\mu}_t^0 = \mu_t$ 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 $\hat{\mu}_t^1 = \mu_0$ 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 $\hat{\mu}_t^0 = \mu_t$ 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 $\hat{\mu}_t^0 = \mu_t$ and $\hat{\mu}_s^1 = \mu_0$ for all $s \geq t$. It follows that the agent has strict incentives to continue at t if $\hat{\mu}_t^0 \leq \mu_t$ and $\hat{\mu}_s^1 \geq \mu_0$ for all $s \geq t$. By Claims 2 and 3 above, $\hat{\mu}_s^1 \geq \mu_0$ for all $s \geq 0$. Hence, the agent stopping at t in equilibrium requires $\hat{\mu}_t^0 > \mu_t$, 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 $\hat{\mu}_t^1 = \mu_0$ for all $t \geq 0$, and the agent must be willing to continue with the project rather than stop at all times. However, since $\mu_t \rightarrow 0$ as $t \rightarrow \infty$ (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 $t \rightarrow \infty$ does not exist.

Proof of Claim 6: Suppose by contradiction that such an equilibrium exists. Then as $t \rightarrow \infty$, absent success, the agent must weakly prefer to continue with the project rather than stop. Since $\mu_t \rightarrow 0$ as $t \rightarrow \infty$ (and the agent's expected payoff from stopping is weakly positive), this requires that for some $\hat{\mu}_\infty^1 \leq 1$, $\hat{\mu}_\infty^1 R + x \geq 0$. 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, $\hat{\mu}_t^1$, 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 $\hat{\mu}_t^1$ is discontinuous exists. Let \hat{t} be the earliest time at which this belief jumps. By Claims 2 and 3, $\hat{\mu}_t^1$ is always weakly increasing and can only jump up. Suppose the belief jumps at \hat{t} from $\hat{\mu}_{\hat{t}-}^1 = \hat{\mu}^{1-}$ to $\hat{\mu}_{\hat{t}+}^1 = \hat{\mu}^{1+} > \hat{\mu}^{1-}$. 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 $\hat{\mu}_{\hat{t}}^0 = \mu_{\hat{t}}$. 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 t 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 used to derive (19), taking dt to 0, this condition is

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

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

$$R(\hat{\mu}^{1-} - \hat{\mu}_{\hat{t}-}^0) + x + \mu_{\hat{t}}\lambda_G \left(V_{\hat{t}} - \frac{R}{r} \right) \geq 0. \quad (34)$$

However, $\hat{\mu}^{1+} > \hat{\mu}^{1-}$ and $\hat{\mu}_{\hat{t}-}^0 \geq \mu_{\hat{t}}$ imply that (33) and (34) 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 $\hat{\mu}_t^1 < 1$ 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 $\hat{\mu}_t^1$ is constant at some value, call it $\bar{\mu}$, for $t \geq \bar{t}$. The agent's indifference condition (19) at \bar{t} in this case yields

$$R\bar{\mu} + x + \mu_{\bar{t}} \left(\lambda_G - R + \lambda_G \frac{\lambda_G + x}{r} - \lambda_G (1 - \bar{\mu}) \frac{R}{r} \right) = 0. \quad (35)$$

Note that $\bar{\mu} < 1$ 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 (35) 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

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

Note that [Assumption 1](#) and $R < -x/\mu_0$ imply $\lambda_G > R$. Thus, the above inequality can hold only if $x + \bar{\mu}R < 0$. However, if the parenthesis in (35) is negative and $x + \bar{\mu}R < 0$, (35) 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 $\hat{\mu}_t^1 = 1$ for all $t \geq \bar{t}$, and hence equation (31) must hold at time \bar{t} . However, if $R > -x$, this equation yields $\mu_{\bar{t}} < 0$ (recall that $\lambda_G > R$ by Assumption 1 and $R < -x/\mu_0$), a 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 $\hat{\mu}_t^1$ 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} is given by (9) and (24), and equations (22) and (23) hold at each $t \geq \underline{t}$. The market's belief conditional on the agent not having stopped is $\hat{\mu}_t^1 = \mu_0$ for $t < \underline{t}$, and, by equation (23), this belief can be written as a function of μ_t independent of μ_0 for $t \geq \underline{t}$. Note also that, 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 (24).

The equilibrium therefore implies that, for any $\delta \geq 0$, $\hat{\mu}_{t^{FB}(\mu_0)+\delta}^1$ 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, $\hat{\mu}^1(\mu', \mu_0)$. The construction implies that if $\mu' > \mu_{\underline{t}(\mu'_0)}$, $\hat{\mu}^1(\mu', \mu_0)$ 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 $\hat{\mu}^1(\mu', \mu_0)$ is invariant to μ_0 up to $\mu_0 = \hat{\mu}^1(\mu', \mu'_0)$, and increases one-for-one with μ_0 for $\mu_0 > \hat{\mu}^1(\mu', \mu'_0)$. Figure 5 provides an illustration.

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

$$\frac{\partial \hat{\mu}_{t^{FB}(\mu_0)+\delta}^1}{\partial \mu_0} \geq 0, \quad \frac{\partial^2 \hat{\mu}_{t^{FB}(\mu_0)+\delta}^1}{\partial \mu_0^2} \geq 0. \quad (36)$$

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:

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

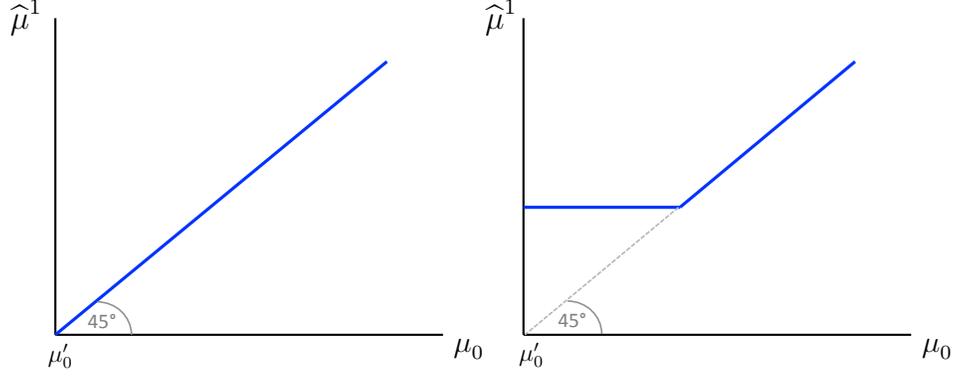


Figure 5: Market belief $\hat{\mu}^1$ as a function of μ_0 for a fixed posterior belief μ . The left graph corresponds to a fixed posterior belief $\mu > \mu_{t(\mu'_0)}$, where μ'_0 is the lowest prior considered in the figure. The right graph corresponds to a posterior belief $\mu < \mu_{t(\mu'_0)}$.

The market's belief satisfies

$$\hat{\mu}_{t^{FB}(\mu_0)+\delta}^1 = 1 - \eta_{t^{FB}(\mu_0)+\delta} + \eta_{t^{FB}(\mu_0)+\delta} \mu_{t^{FB}(\mu_0)+\delta}. \quad (37)$$

Since $\mu_{t^{FB}(\mu_0)+\delta}$ is independent of μ_0 , differentiating (37) yields

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

Combining this with (36), we obtain that for any $\delta \geq 0$,

$$\frac{\partial \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \leq 0, \quad \frac{\partial^2 \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq 0. \quad (38)$$

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

$$\begin{aligned} &\Pr(\text{succeeded \& cont})_{t^{FB}(\mu_0)+\delta} (\lambda_G + x) \\ &+ \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta} (\mu_{t^{FB}(\mu_0)+\delta} \lambda_G + x). \end{aligned} \quad (39)$$

Suppose for the purpose of contradiction that for some $\delta > 0$, (39) is decreasing in μ_0 , that is (using the fact that $\mu_{t^{FB}(\mu_0)+\delta}$ is independent of μ_0),

$$\begin{aligned} & \frac{\partial \Pr(\text{succeeded \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} (\lambda_G + x) \\ & + \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} (\mu_{t^{FB}(\mu_0)+\delta} \lambda_G + x) < 0. \end{aligned} \quad (40)$$

We can rewrite (40) as

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

Note that the derivative on the left-hand side is positive.²⁸ Moreover, Assumption 1 implies $(\lambda_G + x)/\lambda_G > 1 - \mu_0$. Hence, (41) 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

$$\frac{\partial \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0} \left(\mu_0 e^{-\lambda_G(t^{FB}(\mu_0)+\delta)} + 1 - \mu_0 \right) < \frac{\partial \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0}. \quad (42)$$

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, (42) implies

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

We now show that (43) contradicts (38), namely the fact that $\eta_{t^{FB}(\mu_0)+\delta}$ is decreasing in

²⁸To 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 μ_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 (43).

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

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

By (38), $\frac{\partial^2 \eta_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq 0$. Hence,

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

Equivalently,

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

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

$$\frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \leq \frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \eta_{t^{FB}(\mu_0)+\delta}. \quad (44)$$

If the derivative on the left-hand side of (44) 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 (44) 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:

$$\begin{aligned} & \frac{\partial^2 \Pr(\text{succeeded \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} (\lambda_G + x) \\ & + \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} (\mu_{t^{FB}(\mu_0)+\delta} \lambda_G + x) < 0. \end{aligned} \quad (45)$$

We can rewrite (45) as

$$\frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} (\lambda_G + x) < \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial^2 \mu_0} \lambda_G (1 - \mu_{t^{FB}(\mu_0)+\delta}). \quad (46)$$

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 (41)-(43) to show that (46) implies

$$\frac{\partial^2 \Pr(\text{cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2} \eta_{t^{FB}(\mu_0)+\delta} < \frac{\partial^2 \Pr(\text{did not succeed \& cont})_{t^{FB}(\mu_0)+\delta}}{\partial \mu_0^2}. \quad (47)$$

This inequality is in contradiction with (44).

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 $t \in \{0, dt, 2dt, \dots\}$, 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 $1 - (1 - \lambda_G dt)^{\frac{\underline{t}}{dt}}$.

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 (23). 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

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

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,

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

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

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

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

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

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.