SEQUENTIAL LEARNING

YAIR ANTLER[†], DANIEL BIRD[§], AND SANTIAGO OLIVEROS[⊥]

Abstract. Two players sequentially and privately examine a project of unknown quality. Launching the project requires mutual consent and the first player values the project more than the second player does. The combination of the conflict of interest and private learning leads to moral hazard. We show that an efficient equilibrium must take one of two forms as a function of the prior: either one player relinquishes control of the project, thereby rendering the collaboration moot, or the first player occasionally makes false claims about achieving positive findings. In the latter case, the players' relevant beliefs diverge as time progresses. In addition, we show that projects for which an initial examination failed to generate positive findings may be launched, and that projects known to be good by the first player may be delayed or even aborted.

1. Introduction

Modern economic interactions involve collaboration between agents with complementary skills. For instance, launching a new product may require

---

Date: Friday 16th August, 2019.
We thank seminar audiences at Hebrew University, Stanford GSB, UCLA, UCSB, University of Rochester, and SAET 2019 (Ischia). We thank Steve Callander, Yoram Halevy, Sinem Hidir, Hugo Hopenhayn, and Konrad Mierendorff. Yair Antler and Daniel Bird gratefully acknowledge the financial support of Israel Science Foundation grant No. 1917/19. Santiago Oliveros gratefully acknowledges the financial support of the Economic and Social Research Council, UK, grant ES/S01053X/1.
[†]Coller School of Management, Tel Aviv University, e-mail: yair.an@gmail.com
[§]Eitan Berglas School of Economics, Tel Aviv University, Israel, e-mail: dbird@post.tau.ac.il
[⊥]Department of Economics, University of Essex, United Kingdom, e-mail: soliveb@essex.ac.uk .

raising new capital, followed by research and development, and a marketing campaign; hiring a new employee may include an examination by the initiating division and final authorization by upper management; and standardization processes such as drug approval begin with private testing by the pharmaceutical company, followed by FDA scrutiny.

Typically, the agents involved in the collaboration act sequentially and have idiosyncratic goals and exposure to market forces. Moreover, as agents have different areas of expertise, it is difficult for them to share hard evidence about their findings. The combination of multiple players with different preferences who operate sequentially and the difficulty of providing hard evidence result in a moral hazard problem. Moreover, the absence of direct contractual relationships between these players exacerbates the moral hazard problem.

In this paper, we develop a model that incorporates these natural features, to better understand the implications of this moral hazard problem. We study how the players' characteristics determine both the manner in which they collaborate and the efficiency of the collaboration process.

In our model, two players, **F** and **S**, jointly decide whether to launch a project whose quality can be either good or bad. First, **F** examines the project and decides whether to abort it or pass it to **S**, who, upon receiving it, can also examine the project and decide whether to launch it or abort it. Each player can use a costly learning technology that produces conclusive evidence of the project's quality (breakthroughs) according to a Poisson process only if the project's quality is good. We assume that the outcome of the learning process is private and that **F** values a good project more than **S** does. Thus, **F** may have an incentive to pass the project *as if* a breakthrough occurred.

We show that in any equilibrium in pure strategies, at most one player examines the project. When the prior belief that the project is good (the prior) is low, **S** prefers to abort the project instead of examining it. When the prior is high **F** is too eager to launch the project and thus cannot be trusted to report his findings honestly. For intermediate priors, **S** is unwilling to examine the project and **F** cannot be trusted to do so. Hence, with intermediate priors, even though both players find aborting the project inferior to **F** examining it, the project is aborted without examination.

For these intermediate priors, in the unique Pareto-efficient equilibrium **F** randomizes between passing the project as if a breakthrough occurred and continuing to learn. As **F** occasionally fakes breakthroughs, when **S** receives the project he is skeptical about its quality and he may examine the project too. After some time in which **F** randomizes, she becomes less eager to launch the project and passes the project only if she observes a breakthrough. We refer to the moment at which **F** stops "faking" breakthroughs as τ^*. If **S** receives the project after τ^*, he infers that the project is good and launches it immediately.

The interaction prior to τ^* consists of two distinct phases: an earlier *verification* phase and a later *partial trust* phase.[1] In the verification phase, if **S** receives the project he launches it only after examining it and observing a breakthrough. As time progresses, **S** gradually increases the amount of time he will devote to examining the project should he receive it. In the partial trust phase, if **S** receives the project he randomizes between immediately launching it and further examining it. As time progresses, he breaks his indifference in a way that is more favorable to **F**.

Since **S** does not have commitment power, the amount of time he spends learning depends on his belief about the quality of the project. The higher his belief, the longer he will learn. Therefore, to sustain **S**'s equilibrium behavior, **S**'s belief conditional on receiving the project cannot decrease

[1]For some parameters, only one of the two phases exists. We provide necessary and sufficient conditions for this in the formal analysis.

over time. As **S**'s belief depends on **F**'s behavior, the frequency with which **F** fakes breakthroughs must decrease over time.

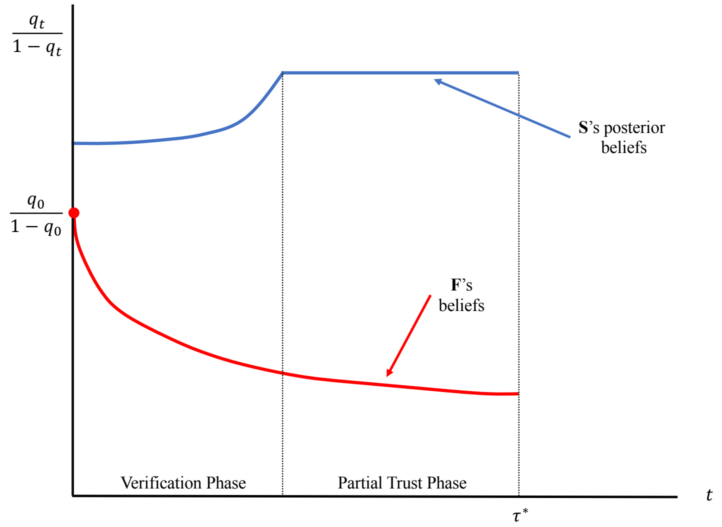

FIGURE 1. Evolution of the players' relevant beliefs in the mixing range.

Figure 1 illustrates the evolution of the likelihood ratios of the players' relevant beliefs in the mixing range of the equilibrium: **F**'s belief while examining the project without observing a breakthrough and **S**'s belief conditional on receiving the project. In the mixing range, these beliefs diverge. On the one hand, **F**'s belief about the project's quality decreases ("no news is bad news"). On the other hand, **S**'s belief conditional on receiving the project increases over time during the verification phase and remains constant in the partial trust phase. If **S** receives the project after time τ^* he infers that **F** has observed a breakthrough and his belief jumps to 1.

When **F** fakes breakthroughs, two types of inefficiencies may arise. In the verification phase, the project is launched only if **S** observes a breakthrough. This leads to delay and, if **F** observes a breakthrough but **S** does not, to abortion of a project that is known to be good. In the partial trust phase, **S** randomizes as well, which leads to approval of bad projects even when **F** examined the project and did not observe a breakthrough.

The rest of the paper is organized as follows. Section 2 covers the related literature and Section 3 presents the model. Section 4 lays the groundwork for the analysis by studying the single decision-maker's problem. In Section 5 we present the analysis of pure strategy equilibria while in Section 6 we present our main results on the existence, characterization, and efficiency of mixed strategy equilibria. Section 7 concludes. All proofs are relegated to the Appendix.

2. Literature Review

We study sequential learning when both players' consent is required to launch a project. The learning technology in our model is inspired by the exponential bandits framework (Keller, Rady, and Cripps (2005)), which the strategic experimentation literature uses to examine free-riding issues that arise in the presence of informational externalities.[2] In models of experimentation players obtain a stream of payoffs while they learn. In our model, to capitalize on their information, agents must terminate the learning process. As a result, optimistic players prefer to launch the project instead of collecting information.

Rosenberg, Solan, and Vieille (2007) and Murto and Välimäki (2011) study the effect of private information on learning and free riding. The main technical question in these papers, and in our model, is how beliefs evolve and whether they diverge or not. Our paper also studies exit decisions but, unlike in these papers, the main force behind our results is the combination of the private nature of the first mover's findings together with *the sequential*

[2] Décamps and Mariotti (2004) and Bonatti and Hörner (2011) introduce direct externalities into the players' payoffs.

interaction. The combination of these assumptions induces a moral hazard problem absent from the learning literature.

Sequentiality is important when a project's completion requires multiple stages. Green and Taylor (2016) study a principal-agent problem where intermediate success is unobservable and the principal provides funds based on reports. Moroni (2017) investigates incentives in teams in a two-stage project where effort is unobservable. Our model differs from these models in two main aspects. First, in their models the principal must design incentives for a specialized agent to learn. By contrast, we assume that no player can commit to a particular incentive scheme. Second, in their models the same agent (team) must complete two independent tasks that arrive sequentially. By contrast, we assume that different players work sequentially on the same task.

In our model, randomization enables the first player to partially transmit information in a credible manner, which induces the second player to collaborate.[3] Similar effects appear in other models using the exponential bandit framework but for different reasons. In Campbell, Ederer, and Spinnewijn (2014), a player who observes a breakthrough wishes to conceal it to incentivize the other player to keep on exerting effort as breakthroughs have decreasing returns. In Guo and Roesler (2016) and Dong (2018), a player who obtains a negative signal about the project wishes to conceal it to induce the other player to experiment. In all of these papers concealing bad information *encourages* (Dong (2018)) the other player to keep on exerting effort but this information can only be concealed if the informed player randomizes.[4] In our model, any new generated information must be

[3]In Kremer, Mansour, and Perry (2014) and Che and Horner (2018), the principal manages the information agents receive in order to induce the agents to acquire socially valuable information by experimenting. In contrast to our work, in these papers credible information transmission is sustained by the principal's commitment power.

[4]This is reminiscent of the "leading by example" effect in Hermalin (1998) and Komai, Stegeman, and Hermalin (2007).

disclosed by assumption, which effectively rules out any signalling motives or the possibility of extracting informational rents.

Finally, our paper is related to Bayesian persuasion models where the receiver has access to private information. In our paper, both the "sender" and the "receiver" have access to costly information. Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) assume that the receiver's information is exogenous while Matyskova (2018) assumes that the receiver's information is endogenous and costly. Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) find conditions for equivalence between mechanisms that condition on the receiver's information and those that do not, and Matyskova (2018) shows that there is no loss of generality in restricting attention to mechanisms in which the receiver does not collect information on the equilibrium path. The receiver's ability to endogenously generate private rents reduces the sender's payoff and may even reduce the receiver's payoff. In contrast to these papers, we assume that the available signals are restricted and that the receiver does not observe the sender's choice of information structure. We find that on the equilibrium path the receiver collects information and this information is privately and socially valuable.

3. The Model

Two players, **F** (she) and **S** (he), jointly decide whether or not to launch a project whose quality is either good or bad. They agree that the project should be launched if and only if it is of good quality and they share a common prior belief that the project is good, which we denote by q_0. If a good project is launched, each player i obtains a payoff of v^i, where $v^{\mathbf{F}} > v^{\mathbf{S}} > 0$. Their payoffs from aborting the project are 0, and we normalize each player's payoff from launching a bad project to -1.

Before making the decision, each player can *privately* examine the project. We assume that the players interact sequentially. The first mover, **F**, examines the project first and then decides whether to abort the project or pass it to the second mover, **S**. Upon receiving the project, **S** examines it and, in turn, decides whether or not to launch it.

We assume that time is continuous and that both players discount the future using a common discount factor of $r > 0$.[5] While a player is examining the project s/he incurs a flow cost of $c > 0$ and breakthroughs occur according to a Poisson process with intensity $\lambda > 0$ if the project is good and do not occur if the project is bad. Thus, a breakthrough reveals the quality of the project.

The moral hazard problem in our paper is a result of the combination of sequential private learning and the difference in players' payoffs. In some applications players may be heterogeneous in many aspects. In particular, they may have access to different learning technologies (λ, c) or have different discount factors (r). As will become clear later, our analysis depends on when and how much players learn. Since any changes in learning decisions can be obtained by changes in players' payoffs, whether or not the players are symmetric in other dimensions has no substantive impact on our results. In order to highlight the strategic trade-offs that arise in our model, we assume that the only heterogeneity is $v^{\mathbf{F}} > v^{\mathbf{S}}$.

3.1. **Strategies and Equilibrium.** Formally, **F** chooses a (potentially stochastic) stopping time $\tau^{\mathbf{F}}$ and, at $\tau^{\mathbf{F}}$, she decides whether to abort the project or pass it to **S**. We assume that if a breakthrough occurs prior to $\tau^{\mathbf{F}}$, **F**'s learning terminates and she passes the project to **S** immediately.[6] Since learning is private, when **S** receives the project he does not know whether **F** observed a breakthrough or not. Thus, his actions can depend only on the time at which he received the project. Without loss of generality, we restrict **S**'s strategies to those of the following form: with probability σ_t he chooses to learn until some stopping time $\tau_t^{\mathbf{S}} \geq t$, and with probability

[5] Our results remain valid if $r = 0$.

[6] This assumption abstracts away from signalling motives that emerge if **F** can hold onto the project after a breakthrough.

$1 - \sigma_t$ he launches the project immediately. If **S** chooses to learn, and observes a breakthrough, he launches the project immediately; otherwise, at $\tau_t^{\mathbf{S}}$, he aborts the project.[7]

We denote the CDF that corresponds to **F**'s stopping strategy τ by $G^{\mathbf{F}}(\tau)$ and denote by $G_t^{\mathbf{S}}(\tau)$ the CDF that corresponds to **S**'s stopping strategy when receiving the project at t. We refer to a stopping time and its CDF interchangeably. Since **F**'s learning is private, in equilibrium, **S** updates his belief based on **F**'s strategy. To ensure that **S**'s equilibrium belief is well defined at all points in time we impose the following restriction on **F**'s strategy.

Assumption 1. *$G^{\boldsymbol{F}}$ is continuously differentiable at all but a finite number of points.*

Note that this assumption does not rule out strategies under which **F** submits the project with positive probability at some (finitely many) points in time (i.e., atoms).[8]

We denote the density function of $G^{\mathbf{F}}$ (whenever it exists) by $g(\cdot)$ and the supremum of its support by $\omega(G^{\mathbf{F}}(\cdot)) = \inf\left\{t : G^{\mathbf{F}}(t) = 1\right\}$. With a slight abuse of notation, we denote this supremum simply by ω. We say that **F** *reports honestly* in the interval L if $g(t) = 0$, $\forall t \in L$. It follows that **F** does not report honestly at t if $g(t) \neq 0$ or if there is an atom at t.

To ensure that the outcome of the game is well defined for all of **F**'s permissible strategies, we restrict attention to **S**'s strategies that are Lebesgue measurable with respect to t.

Assumption 2. *The product measure $\sigma_t \times G_t^{\boldsymbol{S}}$ is Lebesgue measurable with respect to t.*

[7]In Section 4, we show that for any belief **S** may hold upon receiving the project, his optimal choice belongs to this class of strategies.

[8]Formally, we say that there is an atom at τ if $\lim_{t\uparrow\tau} G^{\mathbf{F}}(t) < G^{\mathbf{F}}(\tau)$.

Finally, if learning is prohibitively costly, no player will learn and the problem becomes trivial. We say that learning is non-redundant for a player with a value v if there exist priors for which that player uses the learning technology. It turns out that for our technology this is equivalent to assuming that $\lambda\frac{v}{v+1} > c$.

Assumption 3. *Learning is non-redundant for both players.*

We use Bayesian Nash equilibrium as the solution concept.[9] Since there might be multiple equilibria in this game, we often refine our analysis and focus on Pareto-efficient equilibria and refer to these equilibria as *efficient equilibria*. Moreover, we assume that if **S** receives the project after ω he will abort it immediately in order to avoid specifying redundant off-path behavior.

We denote by q_t the probability that the project is good conditional on no breakthrough occurring until time t. Note that **F**'s belief while she is learning is given by q_t. We denote **S**'s belief upon receiving the project at time t by $q_t^{\mathbf{S}}$. It will be useful to use the likelihood ratio of the project being good, $l_q = \frac{q}{1-q}$, instead of q.

4. The decision-maker's problem

In this section, we study the behavior of a decision maker (DM) whose prior belief is q_0 and the value she obtains from launching a good project is $v > 0$.[10] Before she makes a decision, the DM can examine the project by means of the technology described above. If the DM learns until $t \geq 0$ without observing a breakthrough, then the likelihood ratio that the project is good is given by $l_{q_t} = e^{-\lambda t} l_{q_0}$. Note that this likelihood ratio decreases over time, which justifies the assertion that no news is bad news.

---

[9]In dynamic games of incomplete information it is standard to adopt some form of perfection (e.g., a perfect Bayesian equilibrium), but in our game the two notions are outcome-equivalent. A detailed proof can be provided upon request.

[10]This analysis characterizes **S**'s best response in the collaboration game.

Clearly, if the DM decides to learn until time $t > 0$ she will adopt the project immediately after a breakthrough occurs, and will abort the project at t if a breakthrough has not occurred by then. The DM's expected utility under this strategy is

$$EU(t, q_0) = q_0 \int_0^t \lambda e^{-\lambda s} \left(e^{-rs} v - \int_0^s c e^{ru} du \right) ds - \left(q_0 e^{-\lambda t} + (1 - q_0) \right) \int_0^t c e^{ru} du. \tag{1}$$

The first term represents the net utility of a breakthrough when the project is good, which occurs with an instantaneous probability $\lambda e^{-\lambda s}$. The second term represents the event in which no breakthrough occurs prior to t. In this case, the DM incurs the cost of learning and aborts the project.

The DM will stop learning when she is indifferent between aborting the project immediately and aborting it in dt units of time unless a breakthrough occurs. The cost of learning for dt extra units of time is c while the benefit is $q_t \lambda v$. Thus, the DM will abort the project when her belief satisfies[11]

$$q_t \leq \underline{q}(v) = \frac{c}{\lambda v}.$$

A DM with belief q who decides to learn will learn until her belief decreases to $\underline{q}(v)$. Thus, the amount of time in which she will learn is given by $\frac{l_{\underline{q}(v)}}{l_q} = e^{-\lambda t}$. We can now define the DM's value from learning when her belief is q to be $EU^*(q) \equiv \max_{t \geq 0} EU(t, q)$.

For the DM to examine the project when $q_t > \underline{q}(v)$, it must be that $EU^*(q_t)$ is greater than her utility from launching the project immediately, $q_t v - (1 - q_t)$. The following proposition describes the DM's optimal behavior and establishes a few useful comparative statics.

[11] The following expression explains the precise form of Assumption 3. The DM prefers launching the project to aborting it if $q_t(v + 1) - 1 > 0 \Rightarrow q_t > \frac{1}{1+v}$. Thus, if under the optimal stopping time the DM stops learning when she prefers launching the project to aborting it ($\underline{q}(v) \geq \frac{1}{1+v}$), she will never learn.

Proposition 4.1. *Assume that learning is non-redundant for the DM. There exist two cutoffs* $0 < \underline{q}(v) < \overline{q}(v) < 1$ *such that it is optimal for the DM to abort the project if* $q_t \leq \underline{q}(v)$*, to examine the project if* $q_t \in (\underline{q}(v), \overline{q}(v)]$*, and to launch the project if* $\overline{q}(v) \leq q_t$*. Moreover,* $\underline{q}(v)$ *and* $\overline{q}(v)$ *converge monotonically to zero as* $v \to \infty$*.*

In Figure 2 we illustrate Proposition 4.1 when **F**'s and **S**'s learning regions are disjoint, where $\underline{q}^j = \underline{q}(v^j)$ and $\overline{q}^j = \overline{q}(v^j)$.

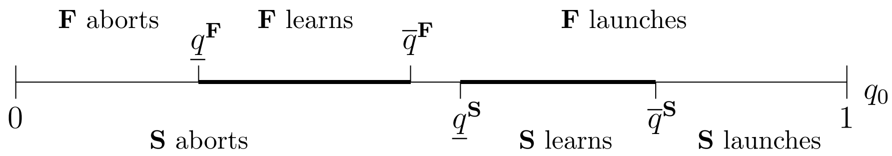

FIGURE 2. Learning regions for $\underline{q}^{\mathbf{S}} > \overline{q}^{\mathbf{F}}$.

In the following sections we study the interaction between the two players. It turns out that in the efficient equilibrium, **F**'s behavior may also resemble the DM's behavior. In the game, this requires that **S** agree with **F**'s choices. We say that ***F*** *behaves as DM* if the players use the following strategy profile: 1) **F** uses the optimal policy for $v^{\mathbf{F}}$ as described in Proposition 4.1, where she breaks her indifference at $\bar{q}^{\mathbf{F}}$ in favor of learning and 2) **S** launches the project immediately upon receiving it.

5. Large Conflict

Since $v^{\mathbf{S}} < v^{\mathbf{F}}$, Proposition 4.1 implies the following partial order of cutoff beliefs:

$$\underline{q}^{\mathbf{F}} < \{\underline{q}^{\mathbf{S}}, \overline{q}^{\mathbf{F}}\} < \overline{q}^{\mathbf{S}}.$$

We focus on the case where $\overline{q}^{\mathbf{F}} < \underline{q}^{\mathbf{S}}$ and refer to it as a *large conflict.* In this case, the players' learning regions in the decision problem do not intersect (see Figure 2). Proposition 4.1 implies that if the difference between $v^{\mathbf{F}}$ and $v^{\mathbf{S}}$ is sufficiently large, the conflict between the players is large.

When there is a large conflict between **F** and **S**, they never agree that examining the project is the best option. Moreover, the players take the same action in the decision problem only for extreme beliefs, namely, $q_0 < \underline{q}^{\mathbf{F}}$ or $\overline{q}^{\mathbf{S}} < q_0$. As we will see later, the large conflict has important implications for the players' interaction. In particular, when $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$, in every pure strategy equilibrium the project is aborted at $t = 0$.[12]

We now characterize the efficient pure strategy equilibria of the game.[13] Potentially, two types of equilibria may exist: equilibria in which at most one player learns and equilibria in which both players learn. We refer to the latter type as *dual-learning equilibria.* The next proposition shows that there are no such equilibria in pure strategies.

Proposition 5.1. *In the unique efficient equilibrium in pure strategies: i) If $q_0 \leq \overline{q}^{\mathbf{F}}$ **F** behaves as DM, and ii) if $q_0 > \overline{q}^{\mathbf{F}}$, **F** passes the project immediately to **S** who uses the optimal policy described in Proposition 4.1.*

Note that in the efficient equilibrium at most one player examines the project. The lack of dual learning implies that an appropriately chosen single decision maker can obtain the same outcome that is obtained in the efficient pure strategy equilibrium of the sequential learning game.

Proposition 5.1 relies on a simple but important observation about pure strategy equilibria: on the equilibrium path, **S** infers perfectly whether **F** observed a breakthrough or not and so the players' relevant beliefs are identical. By Proposition 4.1, a pure strategy equilibrium in which **S** examines the project cannot exist when $q_0 < \underline{q}^{\mathbf{S}}$. It is perhaps less intuitive that a

[12]The assumption that the conflict of interest is large enables us to convey the intuitions for our results clearly and focus on how the players overcome the moral hazard problem when $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$.

[13]The efficient equilibrium is unique up to the identity of the player who aborts the project if the project is aborted at $t = 0$.

dual-learning equilibrium cannot exist when $q_0 > \underline{q}^{\mathbf{S}}$. To see this, note that if **F** passes the project at $\tau^{\mathbf{F}}$, then **S** examines the project until his belief reaches $\underline{q}^{\mathbf{S}}$ unless a breakthrough occurs earlier. However, since $\underline{q}^{\mathbf{S}} > \overline{q}^{\mathbf{F}}$, **F** prefers to launch the project irrespective of the outcome of **S**'s learning, and so she would rather pass the project just before $\tau^{\mathbf{F}}$ and induce **S** to launch the project.

The above argument establishes that if $q_0 > \overline{q}^{\mathbf{F}}$ then **F** does not learn in a pure strategy equilibrium. Since **S** does not learn when $q_0 < \underline{q}^{\mathbf{S}}$, this implies that when $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}}]$ the project is aborted immediately in any pure strategy equilibrium. This collaboration failure is the most prominent manifestation of the moral hazard problem that we study.

Typically, moral hazard problems can be mitigated by contractual agreements that align the players' incentives. However, in our setting players cannot write contracts and so they must rely on other means instead. In particular, we show that strategic uncertainty alleviates the moral hazard problem.

6. Mixed Strategy Equilibrium

In this section we show that the efficient equilibrium is often a dual-learning equilibrium in which the players randomize. This equilibrium consists of two stages: an early stage in which **F** randomizes between examining the project and passing it without observing a breakthrough, and a later stage in which **F** behaves as DM. The transition between stages occurs at a time τ^* such that $q_{\tau^*} = \overline{q}^{\mathbf{F}}$.

First, we establish that in an efficient equilibrium **F** must behave as DM from $\tau^\star$ onward. This requires showing that the continuation values induced by this behavior at time τ^* are consistent with the continuation values that are required to sustain **F**'s mixing prior to τ^* and that this behavior is indeed efficient. In fact, we show that any profile of mixed strategies is

Pareto dominated by the profile of strategies in which **F** does not pass the project without observing a breakthrough up to ω and **S** best responds to this strategy. The problem with such a profile of strategies is that it may not be an equilibrium (as is the case when $q_0 > \overline{q}^{\mathbf{F}}$).

Second, we build on the fact that **F**'s randomization prior to τ^* requires her to be indifferent between passing the project at different times. This indifference condition connects **S**'s belief and his induced behavior at different points in time. In particular, it shows that as time progresses **S**'s belief about the quality of the project must gradually increase. In turn, this implies that **F** gradually reduces the probability with which she passes the project without a breakthrough.

Finally, we show by construction that if there is an equilibrium in mixed strategies in which **F** does not behave as DM from time τ^* onward, then there is another equilibrium that Pareto dominates the original one, in which **F** behaves as DM from time τ^* onward. We finish this section by providing a complete characterization of how players behave while **F** mixes and of the conditions for existence of an efficient mixed strategy equilibrium.

6.1. **The Merits of Honesty.** If **F** randomizes between passing the project without a breakthrough and examining it further, then **S** cannot infer that the project is of good quality upon receiving it. In particular, this may lead **S** to examine the project himself before launching it. In such instances, **F** can free-ride on **S**, as **F** can pass the project and let **S** incur the cost of examining it. It turns out that **F** prefers to report honestly and have **S** launch the project immediately after a breakthrough occurs rather than free-riding by faking breakthroughs.

Lemma 6.1. *Let $q_0 \in (\overline{q}^{\boldsymbol{F}}, \overline{q}^{\boldsymbol{S}}]$. Any equilibrium in which **F** does not report honestly is Pareto dominated by a profile of strategies in which **F** reports* *honestly until ω (the same as in the baseline equilibrium) and **S** best responds. Moreover, in any equilibrium in which $\omega > 0$ there exists $\tau^* > 0$ such that $q_{\tau^*} \leq \overline{q}^{\mathbf{F}}$ and $g(t) \neq 0$ at any $t < \tau^*$.*

To understand the second part of the result, recall that by Proposition 5.1 **F** cannot report honestly while $q_t > \overline{q}^{\mathbf{F}}$. The first part is more involved. Recall that if **F** randomizes she must be indifferent between choosing any stopping time in the support of her strategy and reporting honestly up to ω. If **F** observes a breakthrough, she would rather have **S** launch the project immediately. This is the case if **S** believes that **F** reports honestly, but may not be the case under the profile in which **F** mixes. This change is also beneficial for **S**. When **F** mixes, **S** may receive the project prior to ω for two different reasons: either **F** passed the project after observing a breakthrough or not. In the former case, **S** is better off launching the project immediately without examining it. In the latter case, **S** would rather have **F** continue learning than have her passing the project. In both cases, **S** gets his preferred action when **F** reports honestly.

The main lesson from Lemma 6.1 is that, when $q_0 \leq \overline{q}^{\mathbf{F}}$, **F** behaves as DM in the efficient equilibrium. Moreover, when $q_0 > \overline{q}^{\mathbf{F}}$ any equilibrium in which **F** learns is in mixed strategies. Hence, for the rest of this section we focus on the case where $q_0 > \overline{q}^{\mathbf{F}}$.

The fact that there is a continuation equilibrium in which **F** behaves as DM when $q_0 = \overline{q}^{\mathbf{F}}$ does not guarantee that it is part of a (mixed strategy) equilibrium when $q_0 > \overline{q}^{\mathbf{F}}$. The problem is that **S**'s behavior may induce continuation values that are not necessarily consistent with **F**'s indifference whenever $q_t > \overline{q}^{\mathbf{F}}$. Maintaining **F**'s indifference may require **F** to start behaving as a DM too late (i.e., when $q_t < \overline{q}^{\mathbf{F}}$) or to abort the project too early (i.e., when $q_t > \underline{q}^{\mathbf{F}}$). In the next subsection we study **F**'s incentives when she mixes, and determine what continuation values are required to incentivize her to do so.

6.2. **Beliefs and Incentives.** Examining the project is associated with a direct cost of c and an indirect delay cost due to discounting. In addition, **F**'s behavior is also affected by the changes in her beliefs while she examines the project. In this section we derive **S**'s behavior that maintains **F**'s indifference between continued learning and faking a breakthrough. We then use Bayes' law to infer **F**'s equilibrium behavior from **S**'s belief.

F's expected value from learning until some time $\tau < \tau^*$ while holding belief q_0 is given by

$$V_\tau^{\mathbf{F}} = q_0 \int_0^\tau \lambda e^{-\lambda s} \left[e^{-rs} W_s^B - \int_0^s ce^{-ru} du \right] ds + \left(q_0 e^{-\lambda\tau} + (1-q_0) \right) \left[e^{-r\tau} W_\tau^{NB} - \int_0^\tau ce^{-ru} du \right], \tag{2}$$

where W_s^B is her *expected continuation value after a breakthrough at* s and W_s^{NB} is her *expected continuation value without a breakthrough at* s. [14] The first term of (2) is the expected payoff from passing the project after a breakthrough prior to τ, while the second term is the expected payoff from passing the project at τ without observing a breakthrough.

As **F** is indifferent between all stopping times $\tau < \tau^*$, it follows that $V_\tau^{\mathbf{F}}$ is constant in τ, or, taking the derivative with respect to τ,

$$\lambda q_\tau \left[W_\tau^{NB} - W_\tau^B \right] + r W_\tau^{NB} + c = \frac{dW_\tau^{NB}}{d\tau}. \tag{3}$$

Note that (3) is not a differential equation as W_τ^{NB} and W_τ^B are jointly determined by **S**'s behavior at τ. Nevertheless, this equation can be used to pin down **S**'s equilibrium belief and behavior. In particular, we use it to prove Lemma 6.2, which establishes that at any point in time at which **F** randomizes, **S** must respond by learning with a strictly positive probability if he receives the project. Since at any point in time a breakthrough may cause **F** to pass the project, it must be that $q_t^{\mathbf{S}} \geq q_t$, and so Lemma 6.2 also implies that a mixed strategy equilibrium cannot exist if $q_0 > \overline{q}^{\mathbf{S}}$.

[14] $EU^*(q)$ is a particular case of the more general (2). To see this, set $W_s^B = v^{\mathbf{F}}$ and $W_\tau^{NB} = 0$, and let τ be defined implicitly as $l_{q_0} e^{-\lambda\tau} = l_{\underline{q}^{\mathbf{F}}}$. Integrating gives the result.

Lemma 6.2. *In any equilibrium, $q_t^{\mathbf{S}} \in (\underline{q}^{\mathbf{S}}, \overline{q}^{\mathbf{S}}]$ whenever $g(t) \neq 0$.*

Lemma 6.2 provides necessary conditions for equation (3) to hold. In order to derive the conditions that guarantee a solution, and hence describe the equilibrium behavior, we need to study **S**'s behavior and belief in more detail. Since **S** must respond optimally in equilibrium, it must be the case that if he receives the project at $t < \tau^*$ and his induced belief is $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ he will examine the project, whereas if his induced belief is $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ he may launch the project immediately with some probability. Thus (3) becomes a piecewise differential equation in W_t^B.

To see this formally, note that

$$W_t^B = \sigma_t v^{\mathbf{F}} P^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right) + (1-\sigma_t)v^{\mathbf{F}} \tag{4}$$
$$W_t^{NB} = q_t W_t^B - (1-q_t)(1-\sigma_t),$$

where $P^S(q)$ is the expected discounted arrival time of the first breakthrough when **S** starts learning (optimally) with belief q. Formally,

$$P^S(q) \equiv \begin{cases} \frac{\lambda}{r+\lambda}\left(1-\left(\frac{l_{\underline{q}^{\mathbf{S}}}}{l_q}\right)^{\frac{r+\lambda}{\lambda}}\right) & \text{if} \quad q > \underline{q}^{\mathbf{S}}, \\ 0 & \text{otherwise.} \end{cases}$$

If $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$, then **S** does not launch the project, $\sigma_t = 1$, and the conditions in (4) reduce to $W_t^B = v^{\mathbf{F}} P^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right)$ and $W_t^{NB} = q_t W_t^B$, and (3) becomes

$$r q_t W_t^B + c = q_t \frac{dW_t^B}{dt}. \tag{3b}$$

On the other hand, if $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ and **S** examines the project he will do so for the amount of time it takes his belief to drift down to $\underline{q}^{\mathbf{S}}$ if no breakthrough occurs. We refer to this amount of time as **S**'s *maximal learning*. Hence, the expectation of the discount factor at the first breakthrough is fixed at

$P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right)$. In this case, equation (3) becomes:

$$rq_tW_t^B + c = q_t\frac{dW_t^B}{dt}\frac{\Delta(q_t)}{v^{\mathbf{F}}(1 - P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right))}, \tag{3c}$$

where

$$\Delta(q) \equiv v^{\mathbf{F}} - \frac{1}{l_q} - v^{\mathbf{F}}P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right)$$

is the (scaled) difference between **F**'s payoff from launching the project immediately and his payoff from free-riding on **S**'s maximal learning. It is easy to see that $\Delta(q)$ is increasing in q and hence decreasing in t.

The following proposition describes the dynamics of **S**'s behavior while **F** is mixing. It shows that **S**'s behavior can be separated into (at most) two consecutive phases: a *verification* phase in which **S** examines the project when he receives it ($\sigma = 1$) and a subsequent *partial trust* phase in which he may launch the project immediately ($\sigma < 1$). Note that in the partial trust phase it must be the case that $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$. We denote the end of the verification phase by τ^{**}.

Proposition 6.3. *Let $q_0 \in \left(\overline{q}^{\boldsymbol{F}}, \overline{q}^{\boldsymbol{S}}\right)$. In any equilibrium with $\omega > 0$, **F** mixes continuously on $[0, \tau^*)$, and in this range $\frac{dW_t^B}{dt} > 0$ and $\dot{q}_t^{\boldsymbol{S}} \geq 0$. Moreover, in the verification phase $\dot{q}_t^{\boldsymbol{S}} > 0$, and in the partial trust phase $\dot{\sigma}_t < 0$.*

Since $P^{\mathbf{S}}(\cdot)$ is strictly increasing, this proposition implies that in the verification phase **S** increases his amount of learning over time ($\dot{P}^{\mathbf{S}} > 0$) while in the partial trust phase, if he learns, the amount of time he spends learning is constant over time.

In the verification phase, equation (3b) immediately shows that W_t^B is increasing. The same conclusion in the partial trust phase follows from equation (3c) after we establish that in this phase $\Delta(q_t) > 0$.[15] The economic

[15] When $\Delta(q_t) \leq 0$, **F** prefers free-riding on **S**'s maximal learning to **S** immediately launching the project. Hence, **F**'s value from faking a breakthrough (and inducing a convex combination of immediate approval and free-riding on **S**'s maximal learning) is

interpretation of **F**'s increasing continuation value after a breakthrough is that **S** increases the value of "finding" a breakthrough to compensate **F** for her "investment" in learning.

The intuitions for the proof of Proposition 6.3 highlight the reciprocal nature of the players' equilibrium behavior. Consider first the case in which **F**'s behavior induces $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$. In order for **F** to be continuously indifferent between learning and passing the project without a breakthrough, **S** must increase the time he spends examining the project upon receiving it as time progresses. Since **S** cannot commit to this behavior, **F**'s actions must lead **S**'s belief to increase over time. Hence, the frequency with which she passes the project without observing a breakthrough goes down as time progresses. To see this formally, note that by Bayes' law

$$\frac{g^{\mathbf{F}}(t)}{1 - G^{\mathbf{F}}(t)} = \lambda \frac{l_{q_t}}{l_{q_t^{\mathbf{S}}} - l_{q_t}}. \tag{5}$$

Consider now the case in which **F**'s behavior induces $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$. For this belief, **S**'s learning is constant, and since in this phase $\Delta(q_t) > 0$, to maintain **F**'s indifference **S** must increase the frequency with which he launches the project without examining it over time. As **S**'s belief remains constant, it may appear that **F** does not change her behavior in line with **S**'s changes. However, since true breakthroughs become less likely over time, to maintain **S**'s constant belief, **F** reciprocates by reducing the frequency of faking breakthroughs.

The above intuitions do not explain why there is a single transition between the two phases. To see this, note that **S** learns more in the partial trust phase than in the verification phase. Moreover, in the partial trust phase $\Delta(q_t) > 0$ and so **F** prefers launching the project immediately to free-riding on **S**'s maximal learning. Hence, were there a transition from a partial trust

---

greater than his value from behaving as DM. Since **F**'s continuation value in the partial trust phase is lower than her value from reporting honestly until ω and having **S** best respond to this strategy, it follows that **F** has a profitable deviation in the partial trust phase if $\Delta(q_t) \leq 0$.

phase to a verification phase, **F** would strictly prefer to pass the project at the end of the former phase.

While Assumption 1 does not rule out multiple discontinuities in **F**'s strategy, Proposition 6.3 implies that **S**'s belief may be discontinuous only at $t = \tau^*$. Thus, (5) also implies that **F** mixes continuously. As we will show in the next section, the discontinuity in **S**'s belief at τ^* is the natural reflection of the boundary conditions that follow from the continuity of **F**'s continuation values in the transition between mixing and honest reporting.

6.3. **Characterization of Equilibrium.** We now provide a characterization of the efficient mixed strategy equilibrium. We focus on priors $q_0 \in (\overline{q}^{\mathbf{F}}, \overline{q}^{\mathbf{S}}]$ as the efficient equilibrium is in pure strategies for all $q_0 \leq \overline{q}^{\mathbf{F}}$ and, by Lemma 6.2, there is no mixed strategy equilibrium when $q_0 > \overline{q}^{\mathbf{S}}$. The dynamics of the beliefs that follow from this characterization are illustrated in Figure (1).

First, we must determine the time $\tau^\star$ at which **F** stops randomizing and starts reporting honestly. By Lemma 6.1 the natural candidate for this transition point is the earliest time at which **F** behaving as DM can be sustained in equilibrium. We show that this is indeed the transition point in an efficient equilibrium by proving that efficiency requires that the continuation equilibrium when $q_t = \overline{q}^{\mathbf{F}}$ be efficient itself.[16] Second, we characterize the evolution of the players' behavior prior to $\tau^\star$.[17]

Proposition 6.4. *Let $q_0 \in (\overline{q}^{\boldsymbol{F}}, \overline{q}^{\boldsymbol{S}}]$. In any efficient equilibrium with $\omega > 0$, **F** mixes continuously in $[0, \tau^*)$, $q_{\tau^*} = \overline{q}^{\boldsymbol{F}}$, and **F** behaves as DM from time τ^* onward. Moreover, a partial trust phase exists if and only if*

[16]This step requires showing that the strategies that sustain this efficient equilibrium exist whenever there is an equilibrium (even an inefficient one) in mixed strategies. We show this in technical Lemma A.4, which will also be useful for proving the existence of mixed strategy equilibria in Section 6.4. We postpone the intuition for this result until that section.

[17]We ignore the non-generic case in which the solution to (3) is such that $\lim_{t\downarrow 0} q_t^{\mathbf{S}} \neq q_0$. In this case there may be an equilibrium with an atom at $t = 0$.

$\Delta(\overline{q}^{\mathbf{F}}) > 0$, *in which case* $\sigma_{\tau^*} = 0$ *and* $\sigma_{\tau^{**}} = 1$ *if a verification phase also exists.*

The sign of $\Delta(\overline{q}^{\mathbf{F}})$ determines **F**'s preferences over **S**'s actions if **F** passes the project at τ^*. If the sign is positive, **F** will prefer **S** to launch the project while if it is negative **F** will prefer to free-ride on **S**'s maximal learning. If **F** does not pass the project by τ^* she will behave as DM and obtain a value of $\overline{q}^{\mathbf{F}} v^{\mathbf{F}} - (1 - \overline{q}^{\mathbf{F}})$. Since **S**'s strategy must keep **F** indifferent between learning and passing the project just before τ^*, **S**'s action at τ^* must provide **F** with the (same) value $\overline{q}^{\mathbf{F}} v^{\mathbf{F}} - (1 - \overline{q}^{\mathbf{F}})$.

This indifference implies that the mixing range cannot end in a partial trust phase if $\Delta(\overline{q}^{\mathbf{F}}) < 0$. To see this, note that **F**'s value from passing the project just before τ^* is a convex combination of the value from free-riding on **S**'s maximal learning and the value from launching the project immediately. Since $\Delta(\overline{q}^{\mathbf{F}}) < 0$, any such convex combination exceeds **F**'s value from behaving as DM at τ^*, namely, $\overline{q}^{\mathbf{F}} v^{\mathbf{F}} - (1 - \overline{q}^{\mathbf{F}})$, and so **F** will strictly prefer to pass the project just prior to τ^*.

The mixing must end in a partial trust phase if $\Delta(\overline{q}^{\mathbf{F}}) > 0$. When $\Delta(\overline{q}^{\mathbf{F}}) > 0$ the only way **S** can provide **F** with the value she obtains from behaving as DM at τ^*, $\overline{q}^{\mathbf{F}} v^{\mathbf{F}} - (1 - \overline{q}^{\mathbf{F}})$, is by launching the project immediately if he receives it at τ^*. Hence **S**'s belief must equal $\overline{q}^{\mathbf{S}}$ for $\sigma_{\tau^*} = 0$ to be his optimal response. Furthermore, σ must decrease fast enough to maintain **F**'s incentives to learn in the partial trust phase. This implies that the length of the partial trust phase is bounded from above and so if q_0 is large relative to $\overline{q}^{\mathbf{F}}$ the interaction must start with a verification phase.[18]

The dynamics of the efficient mixed strategy equilibrium also highlight the reciprocal nature of the relationship between **S** and **F**. Although under the learning technology **F**'s learning time and **S**'s learning time are strategic substitutes, they become strategic complements in the efficient mixed

[18]If there is a verification phase before the partial trust phase, the intuition for why $\sigma_{\tau^{**}} = 1$ relies on a similar boundary argument.

strategy equilibrium. This strategic complementarity can be interpreted as an implicit agreement: as time progresses, **F** behaves more honestly, inducing **S** to increase the time he invests in the project upon receiving it, and **F** keeps on examining the project since **S** keeps on increasing the time he invests in the project. The mechanism that enables this type of "reciprocity" is **F**'s equilibrium ability to dilute the content of the information she acquires.

6.4. **Existence of Equilibrium.**

While our characterization is complete, we have not yet established whether a mixed strategy equilibrium exists or not. In Sections 6.2 and 6.3 we characterized **S**'s strategy and belief in the efficient equilibrium, conditional on the existence of an equilibrium in mixed strategies. In doing so we showed that Proposition 6.4 indeed describes the efficient equilibrium if and only if **S**'s belief induced by these strategies satisfies two conditions: it is consistent with Bayes' law and $q_t^{\mathbf{S}} \geq q_t$.

Thus, to establish existence, all we must do is to find a strategy for **F** that induces the belief $q_t^{\mathbf{S}}$ that solves (3). From Bayes' law it follows immediately that a necessary condition for the existence of such a strategy is $q_t^{\mathbf{S}} \geq q_t$. In fact, since **S**'s equilibrium belief pins down the hazard rate of $G^{\mathbf{F}}(\cdot)$ (see (5)) and from this hazard rate one can show that **F**'s strategy has an atom outside the mixing range, it follows that it is possible to find a strategy that supports any sequence of beliefs $q_t^{\mathbf{S}}$ for $t \in [0, \tau^*)$. Hence, the above condition is also sufficient.[19]

In the partial trust phase, $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ and so (if such a phase exists in the equilibrium candidate) the necessary and sufficient conditions for existence hold in that phase. Therefore, to establish that an equilibrium exists it is sufficient to show that $q_t^{\mathbf{S}} > q_t$ in the verification phase of the equilibrium candidate, i.e., when $t < \tau^{**}$. In fact, since $q_t^{\mathbf{S}}$ is increasing the *only*

[19]This claim is established formally in Lemma A.4 in the Appendix.

condition for the existence of **F**'s equilibrium strategy is that $q_0^{\mathbf{S}} > q_0$ when the equilibrium starts with a verification phase.

If the verification phase lasts until the end of the mixing range (i.e., if $\Delta(\overline{q}^{\mathbf{F}}) < 0$), then $\tau^{**} = \tau^*$ and $q_{\tau^{**}} = \overline{q}^{\mathbf{F}}$. If, on the other hand, the verification phase ends before the end of the mixing range, the differential equation (3c) implies that $q_{\tau^{**}}$ is independent of the prior belief q_0.[20] Therefore, (in both cases) equation (3b) pins down $q_0^{\mathbf{S}}$ independently of the prior and it is straightforward to check whether $q_0^{\mathbf{S}} > q_0$ or not. Moreover, (3b) gives rise to a simple and intuitive condition under which the mixed equilibrium exists.

Proposition 6.5. *Let $q_0 \in (\overline{q}^{\boldsymbol{F}}, \overline{q}^{\boldsymbol{S}})$. The mixed strategy equilibrium described by Proposition 6.4 exists if*

$$v^{\boldsymbol{F}} P^{\boldsymbol{S}}(q_0) + c \int_0^{\tau^{**}} \frac{e^{-ru}}{q_u} du < e^{-r\tau^{**}} \begin{cases} v^{\boldsymbol{F}} - \frac{1}{l_{\overline{q}^{\boldsymbol{F}}}} & \text{if } \Delta(\overline{q}^{\boldsymbol{F}}) < 0, \\ v^{\boldsymbol{F}} P^{\boldsymbol{S}}\left(\overline{q}^{\boldsymbol{S}}\right) & \text{if } \Delta(\overline{q}^{\boldsymbol{F}}) > 0. \end{cases} \tag{6}$$

Condition (6) can be interpreted as a cost-benefit analysis for **F**. The RHS is the expected value from passing the project at the end of the verification phase and letting **S** behave as the equilibrium suggests. When $\Delta(\overline{q}^{\mathbf{F}}) < 0$, then **S** adopts the project immediately, and, when $\Delta(\overline{q}^{\mathbf{F}}) > 0$, then **S** learns for the maximal amount of time. On the LHS, we have **F**'s opportunity cost: the value she obtains from passing the project at 0 and letting **S** behave as DM given the prior q_0, and the direct cost of learning until the end of the verification phase.

Proposition 6.5 implies that there is a unique critical belief $q^* > \overline{q}^{\mathbf{F}}$ such that the mixed equilibrium exists if and only if $q_0 < q^*$. This implication is illustrated in Figure 3, where we plot the likelihood ratios of the project's quality being good, l_{q_t}, under parametric assumptions for which $\Delta(\overline{q}^{\mathbf{F}}) > 0$. The blue line indicates the evolution of **S**'s belief (upon receiving the

[20]In the proof of Proposition 6.5 we show the exact determination of τ^{**} and $q_{\tau^{**}}$.

project) in a mixed strategy equilibrium and the red line indicates the evolution of **F**'s beliefs (i.e., q_t). By Proposition 6.4, the mixing ends in a partial trust phase when $\Delta(\overline{q}^{\mathbf{F}}) > 0$ and, given the prior q_0, there may be a verification phase from 0 to τ^{**}. A mixed strategy equilibrium exists as long as **F** can induce $q_0^{\mathbf{S}} > q_0$, which in Figure 3 occurs whenever **S**'s belief dynamics line is above **F**'s belief dynamics line. The intersection of these two lines determines q^*, the highest prior for which a mixed strategy equilibrium exists.

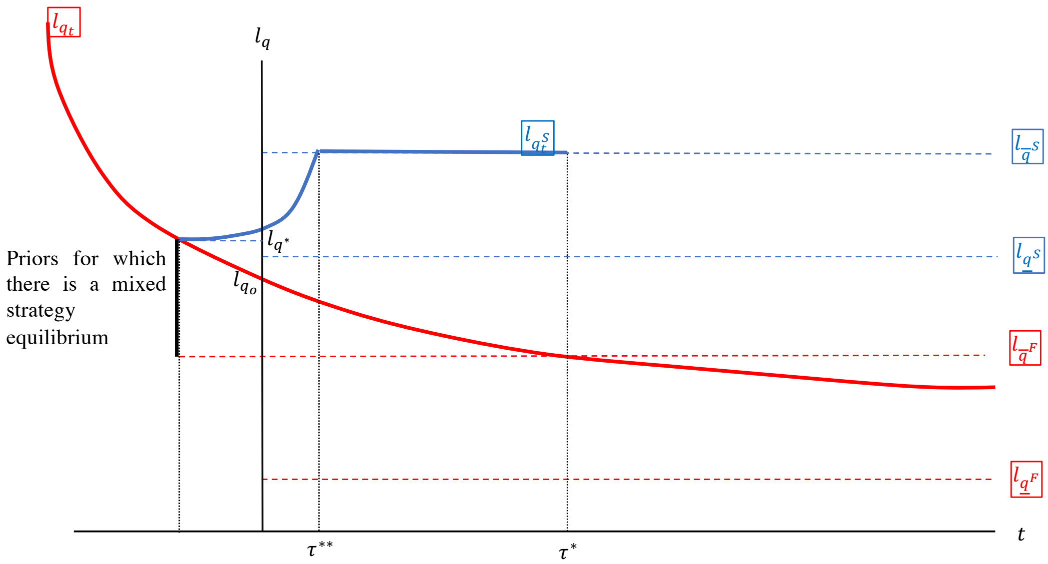

FIGURE 3. The players' relevant belief dynamics in the efficient mixed strategy equilibrium when $\Delta(\overline{q}^{\mathbf{F}}) > 0$ and $\tau^{**} \in (0, \tau^*)$.

Proposition 6.5 enables us to determine whether the collaboration failure in the maximal conflict range can be avoided. In particular, it is sufficient to check whether a mixed equilibrium exists when $q_0 = \underline{q}^{\mathbf{S}}$.

Corollary 6.6. *A mixed strategy equilibrium exists for all* $q_0 \in (\overline{q}^{\boldsymbol{F}}, \underline{q}^{\boldsymbol{S}})$ *if and only if*

$$c\int_0^{\tau^{**}} \frac{e^{-ru}}{q_u} < \begin{cases} e^{-r\tau^{**}}\left(v^{\boldsymbol{F}} - \dfrac{1}{l_{\overline{q}^{\boldsymbol{F}}}}\right) & \text{if } \Delta(\overline{q}^{\boldsymbol{F}}) < 0, \\ e^{-r\tau^{**}} v^{\boldsymbol{F}} P^{\boldsymbol{S}}\left(\overline{q}^{\boldsymbol{S}}\right) & \text{if } \Delta(\overline{q}^{\boldsymbol{F}}) > 0. \end{cases} \tag{7}$$

6.5. **Efficiency of Mixed Strategy Equilibrium.**

Although mixing allows players to increase learning time, there are two welfare costs associated with the mixed strategy equilibrium. The first inefficiency is the natural solution to the moral hazard problem in our model: excessive scrutiny of (good) projects. When **F** observes a breakthrough **S** does not necessarily launch the project immediately. Instead, he may examine the project for a while and approve it only if he observes a breakthrough himself. This leads to a delay in launching some projects or to aborting projects that **F** knows to be good. The second type of inefficiency is less intuitive: a project may be launched after its (costly) examination has uncovered only bad news. If **F** learns and then passes the project without observing a breakthrough in the partial trust phase, **S** may launch the project immediately even though the posterior probability that it is good is lower than what it initially was.

We now provide conditions under which the mixed strategy equilibrium we found is the unique Pareto efficient equilibrium. If the prior is in the range of maximal conflict it is clear that if this equilibrium exists, then it is the unique efficient equilibrium.[21] If $q_0 > \underline{q}^{\mathbf{S}}$, then there is a pure strategy equilibrium in which the project is examined, which **S** may prefer to the mixed strategy one. Nevertheless, the equilibrium in mixed strategies we found is always Pareto efficient as **F** prefers this equilibrium to any pure one. To see this, note that, since $q_0^{\mathbf{S}} > q_0$, passing the project at $t = 0$ in the mixed strategy equilibrium induces **S** to learn for a longer time than he

[21] When $\underline{q}^{\mathbf{S}} < \overline{q}^{\mathbf{F}}$ there is no region of maximal conflict; thus, our welfare analysis is less clear cut.

would learn in the pure strategy equilibrium, where **F** passes the project immediately.

The natural question that arises is whether the mixed strategy equilibrium is the unique Pareto efficient equilibrium when $q_0 > \underline{q}^{\mathbf{S}}$. When $q_0 > \underline{q}^{\mathbf{S}}$, in the efficient pure strategy equilibrium, **F** passes the project at $t = 0$ and **S** learns until his belief reaches $\underline{q}^{\mathbf{S}}$. Thus, when q_0 is sufficiently close to $\underline{q}^{\mathbf{S}}$ the project is launched only if a breakthrough occurs fairly quickly. By contrast, in the mixed strategy equilibrium, the project may be launched even if the first breakthrough occurs fairly late, which can offset the inefficiencies that may arise in the mixed strategy equilibrium. Thus, when the prior is close to $\underline{q}^{\mathbf{S}}$, the mixed strategy equilibrium outperforms the pure strategy equilibrium. The following proposition formalizes this intuition.

Proposition 6.7. *There exists some* $q^{**} \in \left(\overline{q}^{\boldsymbol{F}}, q^{*}\right)$ *such that for all* $q_0 \in \left(\overline{q}^{\boldsymbol{F}}, q^{**}\right)$ *the unique efficient equilibrium is in mixed strategies. Moreover, if condition* (7) *holds then* $\underline{q}^{\boldsymbol{S}} < q^{**}$.

To understand the logic behind this result, note that in both the efficient mixed strategy equilibrium and the efficient pure strategy equilibrium, the value of each player is continuous in the prior. Moreover, when $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$, in the latter case the value of each player is zero whereas in the former case, generically, it is strictly positive. The proof follows from a standard continuity argument and is omitted.

7. Concluding Remarks

We develop a model in which two players examine the quality of a joint project privately and sequentially. We assume that the first player obtains a higher value from launching a good project than the second player. The combination of these assumptions creates a moral hazard problem. The efficient equilibrium is in pure strategies only if the players' prior beliefs are extreme. For intermediate prior beliefs, the efficient equilibrium involves randomization and, unlike in pure strategy equilibria, both players examine the project.

The efficiency of the mixed strategy equilibrium hinges on the role of an action, "passing the project to **S**," as a mechanism that enables *partial* information transmission. Although the information about the quality of the project is perfectly transmitted in pure strategy equilibria, there is no room for collaboration in the region of maximal conflict $(\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$: the project is immediately aborted as no player can be trusted to learn. Mixing in that region allows **F** to dilute the information she transmits when she passes the project. This dilution of information aligns the players' incentives, thereby (partially) solving the moral hazard problem, which in turn benefits both players.

S's ability to also collect information makes **F** willing to collect information herself in this region. In fact, the value of **F**'s information follows from inducing **S** to examine the project. In particular, since **S** adjusts choices smoothly as a function of time, **F** keeps on investing in learning. If **S** were unable to collect information himself, **F** would not have any information to transmit. This is because **F**'s information has no value when **S** cannot reciprocate with respect to her learning time.

References

Bonatti, A., and J. Hörner (2011): "Collaborating," *American Economic Review*, 101(2), 632–63.

Campbell, A., F. Ederer, and J. Spinnewijn (2014): "Delay and deadlines: Freeriding and information revelation in partnerships," *American Economic Journal: Microeconomics*, 6(2), 163–204.

Che, Y.-K., and J. Horner (2018): "Optimal design for social learning," *Quarterly Journal of Economics*, forthcoming.

Décamps, J.-P., and T. Mariotti (2004): "Investment timing and learning externalities," *Journal of Economic Theory*, 118(1), 80–102.

Dong, M. (2018): "Strategic experimentation with asymmetric information," *Penn State University, Working Paper.*

Green, B., and C. R. Taylor (2016): "Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects," *American Economic Review*, 106(12), 3660–99.

Guo, Y., and A.-K. Roesler (2016): "Private learning and exit decisions in collaboration," *Northwestern, Working Paper.*

Hermalin, B. E. (1998): "Toward an economic theory of leadership: Leading by example," *American Economic Review*, 88(5), 1188–1206.

Keller, G., S. Rady, and M. Cripps (2005): "Strategic experimentation with exponential bandits," *Econometrica*, 73(1), 39–68.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017): "Persuasion of a privately informed receiver," *Econometrica*, 85(6), 1949–1964.

Komai, M., M. Stegeman, and B. E. Hermalin (2007): "Leadership and information," *American Economic Review*, 97(3), 944–947.

Kremer, I., Y. Mansour, and M. Perry (2014): "Implementing the wisdom of the crowd," *Journal of Political Economy*, 122(5), 988–1012.

Matyskova, L. (2018): "Bayesian persuasion with costly information acquisition," *CERGE-EI, Working Paper*, (614).

Moroni, S. (2017): "Experimentation in organizations," *University of Pittsburgh, Working Paaper.*

Murto, P., and J. Välimäki (2011): "Learning and information aggregation in an exit game," *The Review of Economic Studies*, 78(4), 1426–1461.

Rosenberg, D., E. Solan, and N. Vieille (2007): "Social learning in one-arm bandit problems," *Econometrica*, 75(6), 1591–1611.

Appendix A. Proofs

Proof of Proposition 4.1

Integrating and rearranging (1), we get

$$EU(t, q_0) = q_0 \left(v - \frac{c}{\lambda}\right) \frac{\lambda}{r+\lambda} \left(1 - e^{-(r+\lambda)t}\right) - (1 - q_0)\frac{c}{r} \left(1 - e^{-rt}\right).$$

Observe that $EU(t, q_0)$ is concave in t. Hence, the FOC $\frac{\partial EU(t,q_0)}{\partial t} = 0$ is necessary and sufficient for optimality:

$$0 = q_0 \left(v - \frac{c}{\lambda}\right) \lambda e^{-(r+\lambda)t} - (1 - q_0)ce^{-rt}$$

and since $l_{q_t} = l_{q_0} e^{-\lambda t}$ we obtain the cutoff beliefs $q_t = \underline{q}(v) = \frac{c}{\lambda v}$.
Learning is better than adopting the project if $M(q) \equiv \frac{EU^*(q)}{q} - \left(v^{\mathbf{F}} - \frac{1}{l_q}\right) \geq 0$ or

$$\frac{c}{r+\lambda}\left(1 - \left(\frac{l_{\underline{q}(v)}}{l_q}\right)^{\frac{\lambda}{r+\lambda}}\right) - \frac{l_{\underline{q}(v)}}{l_q}\frac{c}{r}\left(1 - \left(\frac{l_{\underline{q}(v)}}{l_q}\right)^{\frac{r}{\lambda}}\right) - l_{\underline{q}(v)}\left(v - \frac{1}{l_q}\right) \geq 0.$$

Note first that $\mathbf{M}(q)$ is continuous and decreasing in q. Second, it is positive at $\underline{q}(v)$ if $v < \frac{1}{l_{\underline{q}(v)}}$, which is true by Assumption 3. Third, when $q = 1$ launching the project is the unique best response and so $\lim_{q \to 1} M(q) < 0$. Thus, by the mean value theorem, there is a unique cutoff belief $\overline{q}(v) \in \left(\underline{q}(v), 1\right)$ for which $\mathbf{M}(\overline{q}(v)) = 0$. Hence, $q_t < \overline{q}(v)$ if and only if the DM prefers learning optimally to adopting the project.

We now establish the comparative statics. Trivially, $\underline{q}(v) = \frac{c}{\lambda v}$ converges monotonically to zero as $v \to \infty$. Since $\underline{q}(v)$ is continuous in v for any q the envelope theorem implies that

$$\frac{\partial EU^*(q)}{\partial v} = q\frac{\lambda}{r+\lambda}\left(1 - \left(\frac{l_{\underline{q}(v)}}{l_q}\right)\right)^{\frac{r+\lambda}{r}} < q,$$

which, in turn, implies that $M(q)$ is decreasing in v. Since $M(q)$ is decreasing in q it follows that $\overline{q}(v)$ is decreasing in v by the implicit function theorem. To complete the proof we must show that $\lim_{v\to\infty} \overline{q}(v) = 0$. Assume to the contrary that $\overline{q}(v)$ is bounded from below by some $\tilde{q} > 0$. Note that $EU^*(\tilde{q})$ is bounded from above by $\tilde{q}v\frac{\lambda}{r+\lambda}$; thus, for sufficiently high v we have that $EU^*(\tilde{q}) < \tilde{q}v - (1 - \tilde{q})$, a contradiction.

Proof of Proposition 5.1

In a pure strategy equilibrium Bayes' law implies that $q_t = q_t^{\mathbf{S}}$. To see this, note that if **F** submits the project at $t < \tau^{\mathbf{F}}$, then **S** must infer that a breakthrough occurred and update his belief to $q_t^{\mathbf{S}} = 1$. On the other hand, if **F** submits the project at $t = \tau^{\mathbf{F}}$, then **S** must infer that a breakthrough did not occur up to that point and therefore $q_t^{\mathbf{S}} = q_t$.

If $q_0 \leq \overline{q}^{\mathbf{F}}$, **F** behaving as DM is an equilibrium as **F** does not want to launch the project unless she observes a breakthrough. Moreover, since $q_0 < \underline{q}^{\mathbf{S}}$, it follows that **S** does not learn in any pure strategy equilibrium. Thus, this is the unique efficient equilibrium.

If $q_0 > \overline{q}^{\mathbf{F}}$, there is an equilibrium with $\tau^{\mathbf{F}} = 0$ in which **S** uses an optimal policy for a DM. Too see this, note that any deviation by **F** will lead to **S** aborting the project upon receiving it by our assumption on off-equilibrium beliefs. Next, we show that for these priors there is no equilibrium in pure strategies with $\tau^{\mathbf{F}} > 0$. Let $\tau^{\mathbf{F}} > 0$. There are two possible events: either a player observes a breakthrough before $\underline{q}^{\mathbf{S}}$ or no player does. In the first case, **F** is better off if the project is launched at $t = 0$, and in the second case, the project is aborted. If only **F** learns before the project is aborted,

$q_0 > \overline{q}^{\mathbf{F}}$ implies that she would rather launch at $t = 0$. If both players learn, the project is aborted when $q_t = \underline{q}^{\mathbf{S}}$, and so $\underline{q}^{\mathbf{S}} > \overline{q}^{\mathbf{F}}$ implies that **F** would rather launch the project than abort it. □

Proof of Lemma 6.1

We start this proof by establishing the following technical result.

Lemma A.1. *In any equilibrium, $\omega < \infty$. Moreover, if there is an atom at ω then $q_\omega = q_\omega^{\mathbf{S}}$, and otherwise $\lim_{t\to\omega}(q_t - q_t^{\mathbf{S}}) = 0$.*

Proof of Lemma A.1. First, we show that $\omega < \infty$. If $\omega = \infty$, for every stopping rule τ in the support of $G^{\mathbf{F}}(\cdot)$ there exists a stopping rule $\tau' > 2\tau$ that is also in the support. The continuation payoff at τ from the stopping time τ' is bounded from above by

$$q_\tau v^{\mathbf{F}} - (1-q_\tau)c\frac{1-e^{-r(\tau'-\tau)}}{r} < q_\tau v^{\mathbf{F}} - (1-q_\tau)c\frac{1-e^{-r\tau}}{r}.$$

Since $\lim_{\tau\to\infty} q_\tau = 0$, the RHS converges to $-\frac{c}{r}$, i.e., the continuation utility at sufficiently large τ is negative.

If ω is an atom of G, than **S**'s belief at ω is correct. Assume that ω is not an atom of G. Assumption 1 implies that there exists an interval L ending at ω in which $G(t)$ is continuously differentiable. Let

$$h(t) = \frac{g(t)}{1-G(t)} = \frac{g(t)}{\int_t^\omega g(s)ds}$$

denote the hazard ratio of $G(\cdot)$. By Bayes' law $q_t^{\mathbf{S}} = q_t\frac{\lambda+h(t)}{\lambda q_t+h(t)}$, which shows that $q_t^{\mathbf{S}}$ is continuous in t and thus $\lim_{t\to\omega}(q_t - q_t^{\mathbf{S}})$ exists.

Assume to the contrary that there exists $\epsilon > 0$ such that $q_t^{\mathbf{S}} - q_t > \epsilon$ for all $t \in L$. Since beliefs belong to a compact interval, $\frac{\lambda+h(t)}{\lambda q_t+h(t)} > 1+\delta$ for some $\delta > 0$. If $lim_{t\to\omega}g(t) > 0$, then $lim_{t\to\omega}h(t) = \infty$, and the proof is complete. If $g(\omega) = 0$, then due to the continuity of $g(\cdot)$ there exists a sequence $t_n \to \omega$ such that for every t_n, we have that $g(s) < g(t_n)$ for all

$s > t_n$. For this sequence,

$$h(t_n) > \frac{g(t_n)}{\int_{t_n}^{\omega} g(t_n)ds} = \frac{1}{\omega - t_n};$$

thus, $lim_{n\to\infty} h(t_n) = \infty$, which concludes the proof. □

Consider a mixed strategy equilibrium **E** in which either $g(\omega) > 0$ or there is an atom at ω (the proof when $g(\omega) = 0$ is analogous and hence omitted). **F**'s expected payoff in **E** is equal to her payoff when she uses the stopping time ω. By Lemma A.1, we have that $q_\omega = q_\omega^{\mathbf{S}}$ and thus **S**'s response at ω in **E** is also a best response in the profile in which **F** reports honestly until ω (which we denote by **E'**). In **E'** and when **F** uses the stopping rule ω in **E**, she passes the project only after a breakthrough. However, in the former case **S** immediately launches the project while in the latter case this may not occur. Thus, **F** prefers **E'** to **E**.

Now, we show that **S** strictly prefers the outcome under **E'** to the outcome under **E**. If the project is submitted after a breakthrough at some time $t < \omega$ or if it is submitted at ω, then **S** weakly prefers his payoff under **E'** to that under **E**. If the project is submitted without a breakthrough at some time $t < \omega$, since $q_0 < \overline{q}^{\mathbf{S}}$, then **S** strictly prefers to have **F** continue learning until ω and then proceed with the equilibrium behavior under **E**, which is exactly what occurs in **E'**.

Note that **E'** is a pure strategy profile in which **F** learns honestly until ω. By Proposition 5.1, for all $q > \overline{q}^{\mathbf{F}}$, **F**'s payoff in **E'** is strictly less than her payoff from launching the project immediately. Thus, for any t such that $q_t > \overline{q}^{\mathbf{F}}$, **F**'s continuation payoff from **E** is less than her payoff from launching the project immediately. Thus, for **E** to be an equilibrium, **S** cannot launch the project immediately if he receives it at any time t when $q_t > \overline{q}^{\mathbf{F}}$. Hence, when $q_t > \overline{q}^{\mathbf{F}}$, then **F** cannot report honestly. □

Proof of Lemma 6.2

As **F** must be indifferent between all of the stopping times in her strategy's support, the value function (2) is constant and differentiable at all stopping times $\tau < \tau^*$. Differentiating $V_\tau^{\mathbf{F}}$ w.r.t. to τ gives

$$\begin{aligned} 0 = q_0 \lambda e^{-\lambda\tau} &\left[e^{-r\tau} W_\tau^B - \int_0^\tau c e^{-ru} du \right] \\ &- \lambda q_0 e^{-\lambda\tau} \left[e^{-r\tau} W_\tau^{NB} - \int_0^\tau c e^{-ru} du \right] \\ &+ \left(q_0 e^{-\lambda\tau} + (1-q_0) \right) \left[e^{-r\tau} \frac{dW_\tau^{NB}}{d\tau} - r e^{-r\tau} W_\tau^{NB} - c e^{-r\tau} \right], \end{aligned}$$

which simplifies to equation (3). Note that equation (3) proves the differentiability (and hence the continuity) of W_τ^{NB} .

Assume to the contrary that $q_\tau^{\mathbf{S}} > \overline{q}^{\mathbf{S}}$. Since $q_\tau^{\mathbf{S}} > \overline{q}^{\mathbf{S}}$, **S** adopts the project at τ. Thus, in equilibrium, we must have

$$W_\tau^{NB} = q_\tau v^{\mathbf{F}} - (1-q_\tau) \qquad W_\tau^B = v^{\mathbf{F}} \qquad \frac{dW_\tau^{NB}}{d\tau} = \dot{q}_\tau (v^{\mathbf{F}} + 1).$$

Substituting into (3) and using the law of motion of beliefs $\dot{q}_\tau = -\lambda q_\tau (1 - q_\tau)$, we obtain

$$r \left(q_\tau v^{\mathbf{F}} - (1-q_\tau) \right) + c = 0,$$

which can only hold for a single τ (and not an interval). This is in contradiction to **F**'s indifference between all stopping times $\tau < \tau^*$.

Assume now that $q_\tau^{\mathbf{S}} \leq \underline{q}^{\mathbf{S}}$. It follows that **S** aborts the project at τ. Thus, (3) reduces to $c = 0$, which is also a contradiction. □

Proof of Proposition 6.3

By Assumption 1, $G^{\mathbf{F}}(\cdot)$ is continuously differentiable at all but a finite number of points. Since the number of nondifferentiability points of $G^{\mathbf{F}}(\cdot)$ is finite, we can partition the mixing range $[0, \tau^*)$ into a finite collection of intervals in which $g(t)$ is continuous. Note that within each such interval, **S**'s equilibrium belief is continuous. Let $\mathcal{L}$ denote the partition associated with the equilibrium **E**.

The proof proceeds as follows. First, we establish that the continuation value W_t^B is continuous and strictly increasing in $[0, \tau^*)$. Second, we use this result to show that $q_t^{\mathbf{S}}$ is continuous and increasing in t in this range. Finally, we use these results to derive the proposition.

Lemma A.2. *W_t^B is strictly increasing in any $L \in \mathcal{L}$.*

Proof of Lemma A.2. Consider an arbitrary $L \in \mathcal{L}$. Since $q_t^{\mathbf{S}}$ is continuous on L, this interval can be covered by a collection of intervals $\mathcal{L}'$ such that for every $L' \in \mathcal{L}'$, either: (1) $q_\tau^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ for every $t \in L'$, or (2) $q_\tau^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ for every $t \in L'$.

First, consider an interval $L' \in \mathcal{L}'$ such that $q_\tau^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ for every $t \in L'$. We obtain that the relevant differential equation is

$$c = \frac{dW_t^{NB}}{dt} - (r - \lambda(1 - q_t))\, W_t^{NB}. \tag{3b'}$$

Integrating we obtain that

$$e^{-rt}\frac{W_t^{NB}}{q_t} = e^{-r\bar{b}}\frac{W_{\bar{b}}^{NB}}{q_{\bar{b}}} - \int_t^{\bar{b}} c\frac{e^{-rs}}{q_s}ds,$$

where $\bar{b} \equiv \sup L'$. This implies that in terms of W^B,

$$e^{-rt}W_t^B = e^{-r\bar{b}}W_{\bar{b}}^B - \int_t^{\bar{b}} c\frac{e^{-rs}}{q_s}ds, \tag{8}$$

which implies that W_t^B is increasing over time.

Second, consider an interval $L' \subset \mathcal{L}'$ in which $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$. By Lemma 6.1 we have that in any equilibrium **F**'s continuation utility at any time $t \in L'$ is strictly less than her continuation utility from launching the project immediately. **F**'s expected utility from passing the project at t is a convex combination of launching the project and maximal learning. If $\Delta(q_t) < 0$, then **F** prefers any such convex combination to launching the project immediately, a contradiction. Hence $\Delta(q_t) > 0$ in this interval.

To complete the proof, note that (3c) implies that W_t^B is strictly increasing if $\Delta(q_t) > 0$. □

Lemma A.3. *$q_t^{\mathbf{S}}$ is continuous and increasing at all times $t \in [0, \tau^*)$.*

Proof of Lemma A.3. By the piecewise continuity of $g(\cdot)$ we have finitely many points of discontinuity of $q_t^{\mathbf{S}}$. Let $\tau \in (0, \tau^*)$ be a discontinuity point in $\mathbf{S}$'s belief. First, consider the case where $\lim_{t\uparrow\tau} q_t^{\mathbf{S}} > q_\tau^{\mathbf{S}}$. If $\lim_{t\uparrow\tau} q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$, then the definition of $W_t^B = v^{\mathbf{F}} P^{\mathbf{S}}(q_t^{\mathbf{S}})$ directly implies that W_t^B is not continuous. If $\lim_{t\uparrow\tau} q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$, by Lemma A.2 we have that $\Delta(q_t) > 0$, and so it follows that for any probability of adoption $1 - \sigma > 0$,

$$\text{(9)} \qquad \lim_{t\uparrow\tau}(1-\sigma)\big(q_t v^{\mathbf{F}} - (1-q_t)\big) + \sigma q_t P^{\mathbf{S}}(q_t^{\mathbf{S}}) > q_\tau P^{\mathbf{S}}(q_\tau^{\mathbf{S}})$$
$$\Rightarrow \lim_{t\uparrow\tau} W_t^{NB} > W_\tau^{NB},$$

which contradicts the continuity of $W_t^B = \frac{W_t^{NB}}{q_t}$. Now consider the case $\lim_{t\uparrow\tau} q_t^{\mathbf{S}} < q_\tau^{\mathbf{S}}$. If $q_\tau^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$, then by the definition of W_t^B it must be the case that W_t^B is discontinuous, which is not possible. But if $q_\tau^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$, then $\Delta(q_t) > 0$ implies that $\lim_{t\uparrow\tau} W_t^{NB} < W_\tau^{NB}$, and hence contradicts the continuity of $W_t^B = \frac{W_t^{NB}}{q_t}$. The case where $\lim_{t\downarrow\tau} q_t^{\mathbf{S}} \neq q_\tau^{\mathbf{S}}$ is analogous.

We now show that $q_t^{\mathbf{S}}$ is continuous at $t = 0$. Because $G^{\mathbf{F}}(\cdot)$ is right continuous we must focus only on an atom at $t = 0$. Because $q_0 \leq q_0^{\mathbf{S}}$, we must have that $q_0 < \lim_{t\downarrow\tau} q_t^{\mathbf{S}}$ and hence the previous argument can be applied.

To coomplete the proof, note that in any interval where $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$, by Lemma A.2, we have that $P^{\mathbf{S}}(q_t^{\mathbf{S}}) = \frac{W_t^B}{v^{\mathbf{F}}}$ is strictly increasing. This implies that $q_t^{\mathbf{S}}$ is strictly increasing in such an interval. The continuity of $q_t^{\mathbf{S}}$ together with the fact that $q_t^{\mathbf{S}} \in (\underline{q}^{\mathbf{S}}, \overline{q}^{\mathbf{S}}]$ for any $t < \tau^*$ implies that $q_t^{\mathbf{S}}$ is increasing in $[0, \tau^*)$. □

Since $q_t^{\mathbf{S}}$ is increasing for all $t < \tau^*$, it follows that there is at most one transition from the verification phase to the partial trust phase. That is,

τ^{**} is well defined. Moreover, by the previous argument in the verification phase $q_t^{\mathbf{S}}$ is strictly increasing. To complete the proof, note that in Lemma A.2 we showed that $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ only if $\Delta(q_t) > 0$. In this interval the relevant differential equation (3c) can be written in terms of σ_t to obtain

$$c + rq_t v^{\mathbf{F}} \left(1 - \sigma_t \left(1 - P^{\mathbf{S}} \left(\overline{q}^{\mathbf{S}}\right)\right)\right) = -\dot{\sigma}_t q_t \Delta(q_t). \tag{3c'}$$

Hence, in the partial trust phase $\dot{\sigma}_\tau < 0$. □

Proof of Proposition 6.4

First, we show that a necessary condition for an equilibrium to be efficient is that the continuation equilibrium at τ^* is undominated. To do so, we show that if there exists an equilibrium with the dominated continuation utilities at τ^*, there also exists an equilibrium with the dominating continuation utilities at τ^*. Then, we show that increasing both continuation utilities at τ^* increases the expected utility of both players.

Second, we use the fact that **F** behaving as DM is the unique efficient continuation equilibrium when $q_t = \overline{q}^{\mathbf{F}}$, which makes it sufficient to focus on $q_{\tau^*} = \overline{q}^{\mathbf{F}}$ to pin down the mixing behavior before τ^*.

To establish the first part of the proof we need the following technical lemma.

Lemma A.4. *The solution of equation* (3) *characterizes an equilibrium if and only if* $q_0^{\mathbf{S}} \geq q_0$ *and* $\sigma_0 \leq 1$.

Proof of Lemma A.4. Because the differential equation (3) assumes that **S** best responds to his belief, the solution of equation (3) describes an equilibrium if for every $t \in [0, \tau^*)$ we have that $\sigma_t \in [0, 1]$, $q_t^{\mathbf{S}} \in (\underline{q}^{\mathbf{S}}, \overline{q}^{\mathbf{S}}]$, and $q_t^{\mathbf{S}}$ is consistent with Bayes' law. The first two requirements follow directly from the construction of the differential equation while the last one follows from the fact that, in the mixing range, the hazard ratio of

$G^{\mathbf{F}}(\cdot)$ must be positive and:

$$\lambda \frac{l_{q_t}}{l_{q_t^{\mathbf{S}}} - l_{q_t}} = \frac{g^{\mathbf{F}}(t)}{1 - G^{\mathbf{F}}(t)}$$

(by (5)). This implies that $q_t < q_t^{\mathbf{S}}$ for all $t < \tau^*$ and since $q_t^{\mathbf{S}}$ is increasing it is in fact sufficient to verify that $q_0 \le q_0^{\mathbf{S}}$.

Finally, we also need to show that there is a strategy for $\mathbf{F}$ that sustains these beliefs. Integrating (5) we have that

$$1 - G^{\mathbf{F}}(t) = \left(1 - G^{\mathbf{F}}(\tau^*)\right) e^{-\lambda \int_t^{\tau^*} \left(\frac{l_{q_u}}{l_{q_u^{\mathbf{S}}} - l_{q_u}}\right) du}$$

and substituting into (5) we have that $\mathbf{F}$'s strategy is

$$g^{\mathbf{F}}(s) = \lambda \frac{l_{q_s}}{l_{q_s^{\mathbf{S}}} - l_{q_s}} \left(1 - G^{\mathbf{F}}(\tau^*)\right) e^{-\lambda \int_s^{\tau^*} \left(\frac{l_{q_u}}{l_{q_u^{\mathbf{S}}} - l_{q_u}}\right) du}. \tag{10}$$

Integrating $h(s) \equiv \frac{g^{\mathbf{F}}(s)}{1 - G^{\mathbf{F}}(s)}$ we have that $G^{\mathbf{F}}(\tau^*) = 1 - e^{-\int_0^{\tau^*} h(u) du} < 1$. Since $\mathbf{F}$ behaves as DM from time τ^* onward, her strategy has an atom at ω, and so her strategy can be completed by setting the weight of this atom to $1 - G^{\mathbf{F}}(\tau^*)$. □

Solving the piecewise differential equation (3) requires a boundary condition. We use $\mathbf{F}$'s continuation utility at τ^* as this boundary condition. Denote by $W_t^{NB}(U)$ the solution to (3) with boundary condition U at τ^*, denote by $q_t^{\mathbf{S}}(U)$ the belief associated with this solution, and denote by $\sigma_t(U)$ $\mathbf{S}$'s strategy in this equilibrium.

Recall that the particular solution to a first-order differential equation such as (3) is uniquely determined by its boundary condition, U. Moreover, it is pointwise monotonically increasing in this boundary condition and so $W_t^{NB}(U)$ is increasing in U for any $t < \tau^*$. Let $U^2 > U^1$ be two feasible continuation utilities (i.e., continuation utilities that can be supported by some continuation equilibria) at τ^*, by the previous observation that $W_0^{NB}(U^2) > W_0^{NB}(U^1)$.

First consider the case where $U^2 \leq q_{\tau^*} v^{\mathbf{F}} P^{\mathbf{S}}\left(\bar{q}^{\mathbf{S}}\right)$. In this case, Proposition 6.3 implies that for $i = 1, 2$, $q_t^{\mathbf{S}}(U^i) < \bar{q}^{\mathbf{S}}$ for all $t < \tau^*$ and so $\sigma_0(U^i) = 1$ and $W_0^{NB}(U^i) = q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0^{\mathbf{S}}(U^i)\right)$. Hence we have that $q_0^{\mathbf{S}}(U^1) < q_0^{\mathbf{S}}(U^2)$ by the monotonicity of $P^{\mathbf{S}}(\cdot)$ and $W_0^{NB}(U^2) > W_0^{NB}(U^1)$. Thus, the continuation utility U^2 can be supported by some equilibrium.

Next, consider the case where $U^2 > q_{\tau^*} v^{\mathbf{F}} P^{\mathbf{S}}\left(\bar{q}^{\mathbf{S}}\right)$; this continuation utility is feasible only if $\Delta(q_{\tau^*}) > 0$. Since $\Delta(\cdot)$ is monotonically increasing, it follows that $\Delta(q_0) > 0$, which implies that $W_0^{NB} \geq q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0^{\mathbf{S}}\right)$, and that this weak inequality holds with equality if $q_0^{\mathbf{S}} < \bar{q}^{\mathbf{S}}$. Thus, if $q_0^{\mathbf{S}}(U^1) > q_0^{\mathbf{S}}(U^2)$, we have that

$$W_0^{NB}(U^2) = q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0^{\mathbf{S}}(U^2)\right) < q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0^{\mathbf{S}}(U^1)\right) \leq W_0^{NB}(U^1),$$

in contradiction to the fact that $W_0^{NB}(U^2) > W_0^{NB}(U^1)$. Note, that if $q_0^{\mathbf{S}}(U^1) = q_0^{\mathbf{S}}(U^2) = \bar{q}^{\mathbf{S}}$, then $W_0^{NB}(U^1) < W_0^{NB}(U^2)$ implies that $\sigma_0(U^1) > \sigma_0(U^2)$. Hence the continuation utility U^2 can be supported by some equilibrium.

The previous argument establishes that increasing **F**'s continuation at τ^* strictly increases her value at 0. We now focus on the impact of increasing both players' continuation utilities at τ^* on **S**'s utility at time zero. There is a direct effect due to the change in **S**'s continuation utility at τ^* and an indirect effect due to the change in **F**'s equilibrium strategy induced by the change in her continuation utility at τ^*.

S's expected utility from an equilibrium of the type described in Proposition 6.3 is given by

(11)

$$V_0^{\mathbf{S}} = \int_0^{\tau^*} \left(q_0 \left(1 - G^{\mathbf{F}}(s)\right) \lambda e^{-\lambda s} + \left(q_0 e^{-\lambda s} + (1-q_0)\right) g^{\mathbf{F}}(s)\right) e^{-rs}$$
$$\left[q_s^{\mathbf{S}} \left(v^{\mathbf{S}} - \frac{c}{\lambda}\right) P^{\mathbf{S}}\left(q_s^{\mathbf{S}}\right) - (1-q_s^{\mathbf{S}})\frac{c}{r}\left(1 - \left(\frac{l_{\underline{q}^{\mathbf{S}}}}{l_{q_s^{\mathbf{S}}}}\right)^{\frac{r}{\lambda}}\right)\right] ds$$
$$+ e^{-r\tau^*}\left(1 - G^{\mathbf{F}}(\tau^*)\right)\left(q_0 e^{-\lambda\tau^*} + (1-q_0)\right) V_{\tau^*}^{\mathbf{S}},$$

where $V_{\tau^*}^{\mathbf{S}}$ is $\mathbf{S}$'s continuation utility at τ^*. It follows immediately that the direct effect of $V_{\tau^*}^{\mathbf{S}}$ is positive.

We now focus on the indirect effect. Using (10), we can write the value function (11) as

$$(11b) \quad \frac{V_0^{\mathbf{S}}}{(1-q_0)\left(1-G^{\mathbf{F}}(\tau^*)\right)} = e^{-r\tau^*}\frac{V_{\tau^*}^{\mathbf{S}}}{1-q_{\tau^*}} + \int_0^{\tau^*} \frac{d\left(e^{-\lambda\int_s^{\tau^*}\left(\frac{l_{q_u}}{l_{q_u^{\mathbf{S}}} - l_{q_u}}\right)du}\right)}{ds} e^{-rs}$$
$$+ \left[l_{q_s^{\mathbf{S}}}\left(v^{\mathbf{S}} - \frac{c}{\lambda}\right) P^{\mathbf{S}}\left(q_s^{\mathbf{S}}\right) - \frac{c}{r}\left(1 - \left(\frac{l_{\underline{q}^{\mathbf{S}}}}{l_{q_s^{\mathbf{S}}}}\right)^{\frac{r}{\lambda}}\right)\right] ds.$$

It is easy to see that the first and second parts of the integrand are increasing in $l_{q_u^{\mathbf{S}}}$ and $l_{q_s^{\mathbf{S}}}$, respectively. Hence, it is sufficient to show that $1 - G^{\mathbf{F}}(\tau^*)$ does not decrease when $\mathbf{F}$'s continuation utility increases. The fact that the indirect effect is also positive follows from Lemma A.4, where we showed that $G^{\mathbf{F}}(\tau^*) = 1 - e^{-\int_0^{\tau^*} h(u)du}$.

By Proposition 5.1 and Lemma 6.1, $\mathbf{F}$ behaving as DM from time τ^* onward is the unique Pareto efficient continuation equilibrium at time τ^*. Hence, by the previous arguments, this must be the continuation equilibrium at τ^* in any Pareto-efficient equilibrium of the type described in Proposition 6.3.

To complete the proof, we now determine the structure of the equilibrium in $[0, \tau^*]$ that must be played in order for there to exist a continuation equilibrium at τ^* in which $\mathbf{F}$ behaves as DM. Consider first the case where $\Delta(\overline{q}^{\mathbf{F}}) > 0$. If there is no partial trust phase, then

$$W_{\tau^*}^{NB} \leq \overline{q}^{\mathbf{F}} P(\overline{q}^{\mathbf{S}}) v^{\mathbf{F}} < \overline{q}^{\mathbf{F}} v^{\mathbf{F}} - (1 - \overline{q}^{\mathbf{F}}),$$

which contradicts the continuity of $\mathbf{F}$'s value function at τ^*. Next, consider the case where $\Delta(\overline{q}^{\mathbf{F}}) < 0$ and assume that there is a partial trust phase. By the continuity of $\Delta(q)$ this implies that there exists an interval in this phase for which $\Delta(q_\tau) < 0$, which contradicts Lemma A.2.

Note that, generically, in the equilibrium characterized above $\lim_{t\downarrow 0} q_t^{\mathbf{S}} \neq q_0$. Therefore, generically, $G^{\mathbf{F}}(\cdot)$ cannot have an atom at zero. □

Proof of Proposition 6.5

Recall that by Lemma A.4, $q_0^{\mathbf{S}} \geq q_0$ and $\sigma_0 \leq 1$ are necesarry and sufficient conditions for the existence of the mixed strategy equilibrium. We construct the equilibrium by focusing on the cases where $\Delta(\overline{q}^{\mathbf{F}}) > 0$ and $\Delta(\overline{q}^{\mathbf{F}}) < 0$.

Consider first the case where $\Delta(\overline{q}^{\mathbf{F}}) > 0$. The solution of the differential equation in the partial trust phase equation (3c′) is

$$\sigma_{\tau^*} - \sigma_t \times e^{\int_t^{\tau^*} a_s ds} = -\int_t^{\tau^*} \frac{\frac{c}{q_s} + rv^{\mathbf{F}}}{\Delta(q_s)} \times e^{\int_t^{\tau^*} a_u du} ds,$$

where $a_s = r\left(1 + \frac{1}{l_{q_s}\Delta(q_s)}\right)$. Integrating we have that

$$\int_t^{\tau^*} a_s ds = r\left(\tau^* - t - \frac{1}{\lambda}\ln\left(\frac{\Delta(q_{\tau^*})}{\Delta(q_t)}\right)\right)$$

and since in this case $\sigma_{\tau^*} = 0$, it follows that in the partial trust phase σ_t is given by

$$\sigma_t \times e^{-rt} \times (\Delta(q_t))^{\frac{1}{\lambda}} = \int_t^{\tau^*} \left(\frac{c}{q_s} + rv^{\mathbf{F}}\right) \times e^{-rs} \times (\Delta(q_s))^{\frac{1-\lambda}{\lambda}} ds. \tag{12}$$

If

$$(\Delta(q_0))^{\frac{1}{\lambda}} \geq \int_0^{\tau^*} \left(\frac{c}{q_s} + rv^{\mathbf{F}}\right) \times e^{-rs} \times (\Delta(q_s))^{\frac{1-\lambda}{\lambda}} ds,$$

then $\sigma_t < 1$ for all $t > 0$, and there is no verification phase. In this case $q_0^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ and the equilibrium exists.

When $\Delta(\overline{q}^{\mathbf{F}}) > 0$, a verification phase exists only if

$$(\Delta(q_0))^{\frac{1}{\lambda}} < \int_0^{\tau^*} \left(\frac{c}{q_s} + rv^{\mathbf{F}}\right) \times e^{-rs} \times (\Delta(q_s))^{\frac{1-\lambda}{\lambda}}\, ds,$$

in which case τ^{**} is given by setting $\sigma_t = 1$ in (12), which yields

$$e^{-r\tau^{**}} \times (\Delta(q_{\tau^{**}}))^{\frac{1}{\lambda}} = \int_{\tau^{**}}^{\tau^*} \left(\frac{c}{q_s} + rv^{\mathbf{F}}\right) \times e^{-rs} \times (\Delta(q_s))^{\frac{1-\lambda}{\lambda}}\, ds.$$

In this phase the relevant differential equation (8) reduces to

$$e^{-rt} P^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right) = e^{-r\tau^{**}} P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right) - \frac{c}{v^{\mathbf{F}}} \int_t^{\tau^{**}} \frac{e^{-ru}}{q_u} du. \tag{8b}$$

Recall that the equilibrium exists if $q_0^{\mathbf{S}} > q_0$, which is equivalent to:

$$P^{\mathbf{S}}\left(q_0\right) < e^{-r\tau^{**}} P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right) - \frac{c}{v^{\mathbf{F}}} \int_0^{\tau^{**}} \frac{e^{-ru}}{q_u} du.$$

Now consider the case where $\Delta(\overline{q}^{\mathbf{F}}) < 0$ and there is only a verification phase. The relevant differential equation (8) reduces to

$$e^{-rt} v^{\mathbf{F}} P^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right) = e^{-r\tau^{**}} \left(v^{\mathbf{F}} - \frac{1}{l_{\overline{q}^{\mathbf{F}}}}\right) - c \int_t^{\tau^{**}} \frac{e^{-ru}}{q_u} du, \tag{8c}$$

and the condition for existence reduces to

$$v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0\right) < e^{-r\tau^{*}} \left(v^{\mathbf{F}} - \frac{1}{l_{\overline{q}^{\mathbf{F}}}}\right) - c \int_0^{\tau^{*}} \frac{e^{-ru}}{q_u} du.$$

□