# SEQUENTIAL LEARNING

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

Abstract. Two players sequentially and privately examine a joint project. We assume that the player who moves first values the project more than the player who moves second, so private learning leads to a moral hazard problem. We show that for intermediate priors in the unique Pareto-efficient equilibrium the first mover makes false claims about achieving positive findings, which results in strategic uncertainty. In this equilibrium, the players' relevant beliefs diverge over time, projects for which an initial examination failed to uncover positive findings may be launched, and projects known to be good by the first player may be delayed or terminated.

JEL Codes: D83, D82, O31, D70.

Keywords: Sequential Learning, Strategic Learning, Strategic Uncertainty.

---

*Date*: Sunday 19th April, 2020.

We thank seminar audiences at Bocconi University, Cornell University, Collegio Carlo Alberto, Georgetown University, Hebrew University, Stanford GSB, UCLA, UCSB, University of Rochester, SAET 2019 (Ischia), and EEA-ESEM 2019 (Manchester). We thank Steve Callander, Yoram Halevy, Sinem Hidir, Hugo Hopenhayn, Navin Kartik, Konrad Mierendorff, Marco Ottaviani, Joel Sobel, Balazs Szentes, and Christoph Wolf. Yair Antler and Daniel Bird gratefully acknowledge the financial support of Israel Science Foundation grant 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 .

## 1. Introduction

Modern economic interactions involve collaboration between agents with complementary skills. Typically, the agents involved in the collaboration act sequentially and have idiosyncratic goals. For example, launching a new product requires research and development followed by quality assurance. Hiring a new employee may include an examination by the initiating division and final authorization by upper management. More prominently, standardization processes such as drug approval begin with private testing by the pharmaceutical company and end with FDA scrutiny.

Importantly, as agents have different expertise, it is difficult for them to share hard evidence about their findings. The combination of these features leads to a moral hazard problem where players who move early may try and misreport findings in order to influence players who move later. This moral hazard problem is often exacerbated by the lack of contractual relationships in such settings.

In this paper, we develop a model to study situations in which heterogeneous agents examine a common project privately and sequentially. In our model, two players, **F**(irst mover), she, and **S**(econd mover), he, jointly decide whether or not to launch a project whose quality can be either good or bad. First, **F** examines the project and decides whether to terminate it or submit it to **S** for approval. Then, upon receiving the project, **S** can further examine it and decide whether to launch it or terminate it. Each player has access to a costly learning technology that produces conclusive evidence of the project's quality (breakthroughs) according to a Poisson process if only if the project's quality is good. We assume that **F** values a good project more than **S** does, and therefore **F** may have an incentive to submit the project for approval *as if* a breakthrough has occurred.

Our main result is that efficient collaboration often relies on initial strategic uncertainty and gradually evolves into honesty. In particular, we show that strategic uncertainty is mutually beneficial for intermediate priors. For such priors, in the unique Pareto-efficient equilibrium, **F** initially randomizes between submitting the project as if a breakthrough has occurred and continuing to learn. As a result, **S** is suspicious when he receives the project and refrains from launching it without scrutiny. After some time, **F** stops randomizing and submits the project only after observing a breakthrough. As a result, **S** infers that a breakthrough has occurred and launches the project immediately upon receiving it.

Submitting the project as if a breakthrough has occurred saves on learning costs and time. Consequently, for **F** to be willing to postpone submitting the project, she must find it beneficial to continue learning. Since there are no transfers, this benefit must come from **S** treating **F**'s submissions more favorably in the future by increasing the amount of time he examines the project or by launching the project with higher probability. For this behavior to be optimal for **S**, the later he receives the project the more optimistic he must be that a breakthrough has occurred. Hence, **F** must behave more honestly as time goes by.

While the interaction is subject to strategic uncertainty it consists of (at most) two distinct phases: an earlier *verification* phase and a later *partial trust* phase. In the verification phase, if **S** receives the project, he launches it only after examining it and observing a breakthrough himself. 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, **S** gradually increases the probability of launching the project immediately.

Strategic uncertainty results in two types of inefficiencies. In the verification phase, the project is launched only if **S** observes a breakthrough. This leads to delay and, in instances where **F** observes a breakthrough but **S** does not, to termination of a project that is known to be good. In the partial trust phase, **S** launches the project upon receiving it with positive probability. This may lead to the approval of a bad project, even in instances where **F** examines the project without observing a breakthrough

The equilibrium with strategic uncertainty has additional notable features. First, players' relevant beliefs diverge over time while **F** randomizes: as no news is bad news, **F** becomes more pessimistic over time, while the later **S** receives the project, the more optimistic he becomes. Second, learning times become strategic complements in equilibrium: the more **F** learns, the more **S** learns upon receiving the project.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the model and Section 4 contains a preliminary analysis that lays the groundwork for the main result. In Section 5 we construct the efficient equilibrium with strategic uncertainty and in Section 6 we determine when that equilibrium is the unique Pareto efficient equilibrium. Section 7 concludes. All proofs are relegated to the Appendix.

## 2. Literature Review

The learning technology in our model is inspired by the strategic experimentation literature (e.g., Keller, Rady and Cripps, 2005), which examines free-riding issues that arise in the presence of informational externalities. 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.

Décamps and Mariotti (2004), Rosenberg, Solan and Vieille (2007), Bonatti and Hörner (2011), and Murto and Välimäki (2011) study the effect of asymmetric information on learning and free-riding. Although, in our model, **S** does not observe **F**'s breakthroughs, the main force driving our results is the *sequential* order in which players interact. The order of play induces a moral hazard problem that is absent from the learning literature.

Our paper is also related to the strand of literature that studies the approval of multistage projects. Green and Taylor (2016) study a principal-agent problem where the first breakthrough is unobservable and the principal provides funds based on soft information. Moroni (2019) and Wolf (2017) investigate the design of incentives for a two-stage project where effort is unobservable. Our model differs from these models in two main aspects. First, in their models, the principal designs incentives for a specialized agent to learn. By contrast, we assume that no player can commit to a particular incentive scheme.[1] Second, in their models, the same agent (team) completes 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 examine the project.[2] Similar randomization effects appear in some models using the the same learning technology, but the they appear for different reasons than in our sequential learning game. In Campbell, Ederer and Spinnewijn (2014), a player who observes a breakthrough wishes to conceal it to incentivize the other player to keep exerting effort. 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 each of these papers, concealing information *encourages* (Dong, 2018) the

[1]The design of dynamic incentives for a single stage has been studied by, among others, Bergemann and Hege (1998), Gerardi and Maestri (2012), Hörner and Samuelson (2013), and Halac, Kartik and Liu (2016) when transfers are available, and Guo (2016), McClellan (2017), and Henry and Ottaviani (2019) when transfers are not available.

[2]In Kremer, Mansour and Perry (2014) and Che and Hörner (2018), the principal strategically discloses information to induce current agents to acquire more information. Bimpikis, Ehsani and Mostagir (2019) study similar issues in a contest framework (see also Halac, Kartik and Liu, 2017 for the design of contests for experimentation). Unlike in our work, in these papers, credible information transmission is sustained by the principal's commitment power.

other player to keep exerting effort, but this information can only be concealed if the informed player randomizes.[3] By contrast, in our model, it is impossible to extract informational rents by signaling.

## 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 $q_0$. Each player's payoff from the project is 0 if the project is terminated and $-1$ if a bad project is launched. We denote player $i$'s payoff from launching a good project by $v^i$ and assume that $v^{\mathbf{F}} > v^{\mathbf{S}} > 0$.

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 terminate the project or submit it to the second mover, **S**, for approval. Upon receiving the project, **S** examines it and, in turn, decides whether or not to launch it.

The moral hazard problem we study stems both from the sequential private learning and from **F** being more eager than **S** to launch the project without examination. As will become clear later, other forms of heterogeneity besides $v^{\mathbf{F}} > v^{\mathbf{S}}$ are essentially equivalent. This is because **S**'s characteristics affect the analysis only through the mapping from his beliefs to his learning decisions. Therefore, any heterogeneity between the players arising from other characteristics (learning ability, discounting) can be analyzed by recovering **S**'s learning policy. To simplify the exposition and highlight the strategic trade-offs that arise in our model, we assume that the only heterogeneity is in the value from launching good projects.

[3]Cetemen, Hwang and Kaya (2019) study how this effect evolves in the presence of public feedback.

We assume that time is continuous and that both players discount the future at the common rate $r > 0$.[4] 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 that the project is good.

If learning is prohibitively costly, no player will learn and the problem becomes trivial. It turns out that both players are willing to learn for some beliefs if and only if $c < \lambda \frac{v^{\mathbf{S}}}{v^{\mathbf{S}}+1}$. Throughout the paper we assume that this inequality holds.[5]

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 terminate the project or to submit it to **S**. We assume that if a breakthrough occurs prior to $\tau^{\mathbf{F}}$, **F**'s learning terminates and she submits the project to **S** immediately.[6] Since learning is private, when **S** receives the project he does not know whether **F** has observed a breakthrough or not. Thus, his actions can depend only on the time at which he receives the project. Without loss of generality (see Section 4.1), we restrict **S**'s strategies, as a function of the time $t$ at which he receives the project, to those of the following form: with probability $\sigma_t$ he chooses to learn until some stopping time $\tau_t^{\mathbf{S}} \geq t$, and with probability $1 - \sigma_t$ he launches the project immediately. If **S** chooses to learn and observes a breakthrough, then he launches the project immediately; otherwise, he terminates the project at $\tau_t^{\mathbf{S}}$.

To ensure that the outcome of the game is well defined we restrict attention to strategies that are Lebesgue measurable with respect to time. Denote by $G^{\mathbf{F}}$ the measure that corresponds to **F**'s stopping rule.

**Assumption 1.** *$\sigma_t, \tau_t^{\mathbf{S}}$, and $G^{\mathbf{F}}$ are Lebesgue measurable with respect to $t$.*

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

[5]The specific form of this restriction is established in Proposition 4.1.

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

We use Bayesian Nash equilibrium as the solution concept.[7] Since there may be multiple equilibria in this game, we often refine our analysis by focusing on Pareto-efficient equilibria. We refer to these equilibria as *efficient equilibria*. Moreover, to avoid specifying redundant off-path behavior we assume that if **S** receives the project off the path of play he terminates it immediately.

Since $G^{\mathbf{F}}$ is Lebesgue measurable it is equivalent to a mixture of an absolutely continuous measure and a discrete measure. Let $g$ denote the the derivative (whenever it exists) of $G^{\mathbf{F}}$. We say that **F** *reports honestly* in the interval $L$ if $g$ exists and equals zero for every $t \in L$.

We abuse notation and refer to $G^{\mathbf{F}}$ and $g$ as the CDF and PDF of **F**'s strategy, respectively. Moreover, we denote the supremum of the support of $G^{\mathbf{F}}$ by $\omega(G^{\mathbf{F}}) = \inf \left\{ t : G^{\mathbf{F}}(t) = 1 \right\}$. When there is no risk of confusion, we denote this supremum by $\omega$. Finally, we denote the probability that the project is good conditional on no breakthrough occurring until time $t$ by $q_t$, the likelihood ratio of the project being good under beliefs $q$ by $l(q) = \frac{q}{1-q}$, and **S**'s belief upon receiving the project at time $t$ by $q_t^{\mathbf{S}}$. Note that **F**'s belief while she is learning is given by $q_t$.

## 4. Preliminary Analysis

4.1. **The Decision-maker's Problem.** The behavior of a decision maker (DM) who can learn about a project's quality before deciding whether or not to launch it plays an important role in the analysis of the sequential learning game. We now study the behavior of a DM whose prior belief is $q_0$ and who obtains a value of $v > 0$ from launching a good project.

Clearly, if the DM decides to learn until time $t > 0$ she will launch the project immediately after a breakthrough occurs, and will terminate the project at $t$ if a breakthrough has not occurred by then. Optimality requires

[7] The Pareto-efficient Nash equilibrium we identify is also a perfect Bayesian equilibrium.

that the DM stop learning when she is indifferent between terminating the project immediately and terminating it in $dt$ units of time, unless a breakthrough occurs. The cost of learning for $dt$ extra units of time is $c\,dt$ while the benefit is $q_t \lambda v\,dt$. Thus, the DM will terminate the project when

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

Since learning is costly, the DM may prefer to launch the project without examination. In particular, she will do so when her prior is sufficiently high. The following proposition not only describes the DM's optimal behavior and establishes some comparative statics, but also characterizes **S**'s best response in the sequential learning game.

**Proposition 4.1.** *Assume that $c < \lambda \frac{v}{v+1}$. There exist two cutoffs $0 < \underline{q}(v) < \overline{q}(v) < 1$ such that it is optimal for the DM to terminate 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 $q_t \geq \overline{q}(v)$. Moreover, $\underline{q}(v)$ and $\overline{q}(v)$ converge monotonically to zero as $v \to \infty$.*

We let $\underline{q}^j = \underline{q}(v^j)$ and $\overline{q}^j = \overline{q}(v^j)$ for $j \in \{\mathbf{F}, \mathbf{S}\}$. In Figure 1, we illustrate Proposition 4.1 when **F**'s and **S**'s learning regions are disjoint.

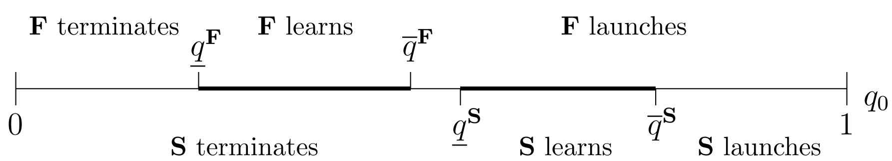


Figure 1. Learning regions for $\underline{q}^{\mathbf{S}} > \overline{q}^{\mathbf{F}}$.

It turns out that in the efficient equilibrium of the sequential learning game, **F**'s behavior occasionally resembles that of a DM. We say that ***F*** *takes over the project* when the strategies of both players replicate the outcomes of **F** behaving as decision maker. Formally, **F** taking over the project is the following strategy profile: 1) $\mathbf{S}$ launches the project immediately upon receiving it, and 2) $\mathbf{F}$ uses the optimal policy for $v^{\mathbf{F}}$ as described in Proposition 4.1, where she breaks her indifference at $\overline{q}^{\mathbf{F}}$ in favor of learning.

4.2. **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 refer to the case where $\overline{q}^{\mathbf{F}} < \underline{q}^{\mathbf{S}}$ as a *large conflict*. There is a large conflict if the difference between $v^{\mathbf{F}}$ and $v^{\mathbf{S}}$ is sufficiently large (Proposition 4.1). When there is a large conflict, the players' learning regions in the decision problem do not intersect (see Figure 1). Hence, they never agree that examining the project is the best option and this leads to an obvious moral hazard problem. In particular, as we establish in the following proposition, under pure strategies the project is terminated immediately when $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$ even though $\mathbf{F}$ strictly prefers learning to terminating the project. Furthermore, due to the moral hazard in the unique efficient equilibrium in pure strategies, at most one player examines the project.[8]

**Proposition 4.2.** *Assume that there is a large conflict. In the unique efficient equilibrium in pure strategies, $\boldsymbol{F}$ takes over the project at time zero if $q_0 \leq \overline{q}^{\boldsymbol{F}}$, and $\boldsymbol{F}$ submits the project immediately to $\boldsymbol{S}$ whose behavior is described in Proposition 4.1 if $\overline{q}^{\boldsymbol{F}} < q_0$.*

Proposition 4.2 relies on a simple but important observation about pure strategy equilibria: on the equilibrium path, $\mathbf{S}$ infers perfectly whether $\mathbf{F}$ has observed a breakthrough or not, and so the players' beliefs are identical. Thus, Proposition 4.1 implies that there is no pure strategy equilibrium in which $\mathbf{S}$ examines the project when $q_0 < \underline{q}^{\mathbf{S}}$. It is perhaps less intuitive

[8]The efficient equilibrium is unique but for the identity of the player who terminates the project when $\mathbf{S}$ does not learn.

that an equilibrium in which both players learn cannot exist when $q_0 > \underline{q}^{\mathbf{S}}$. To see this, note that if **F** submits the project at $\tau^{\mathbf{F}}$, then **S** examines the project until his belief reaches $\underline{q}^{\mathbf{S}}$, unless a breakthrough has occurred earlier. However, under the assumption of large conflict **F** prefers to launch the project at $\underline{q}^{\mathbf{S}}$ irrespective of the outcome of **S**'s learning, and so she would rather submit the project just before $\tau^{\mathbf{F}}$ and induce **S** to launch the project.

4.3. **The Merits of Honesty.** Proposition 4.2 implies that under the assumption of large conflict 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. However, this often results in a project being terminated too early. In particular, if $q_0 \in (\overline{q}^{\mathbf{F}}, \overline{q}^{\mathbf{S}})$, then **S** terminates the project immediately, even though **F** strictly prefers to continue learning.

The reason for this inefficiency is that **F**'s breakthroughs are not public. Thus, were **S** to trust **F** to learn, **F** would manipulate **S** into launching the project. Hence, in an equilibrium in which **F** learns, **S** cannot perfectly infer that a breakthrough occurred upon receiving the project. Thus, in an equilibrium in which **F** learns, she must use a mixed strategy, as otherwise **S** would be able to know upon receiving the project if a breakthrough occur or not.

Although **F**'s mixing creates learning opportunities it introduces a new inefficiency. Namely, the project may be delayed or terminated even after **F** observes a breakthrough. In the following lemma we show that this inefficiency makes strategic uncertainty second best at most.

**Lemma 4.3.** *Let $q_0 < \overline{q}^{\boldsymbol{S}}$. Any equilibrium in which **F** uses a mixed strategy $\hat{G}^{\boldsymbol{F}}$ is Pareto-dominated by a profile of strategies in which **F** reports honestly until $\omega(\hat{G}^{\boldsymbol{F}})$ and then submits the project to **S** who best responds.*

Intuitively, if **F** randomizes, then she must be indifferent between choosing any stopping time in the support of her strategy and reporting honestly up to $\omega$. If **F** observes a breakthrough, she would rather have **S** launch the project immediately. **S** will do so if he believes that **F** is reporting honestly, but he may not launch the project if he believes that **F** is mixing. **S** also benefits from **F** reporting honestly. When **F** uses a mixed strategy, she may submit the project at any time either after having observed a breakthrough or without having observed one. In the former case, **S** is better off launching the project immediately. In the latter case, **S** would rather have **F** continue learning than have her submit the project. In both cases, **S** gets his preferred action when **F** reports honestly.

Lemma 4.3 has two important consequences.

**Lemma 4.4.** *If $q_0 \leq \overline{q}^{\boldsymbol{F}}$ then any efficient equilibrium is in pure strategies. Furthermore, under the assumption of large conflict, there is no equilibrium where **F** reports honestly while $q_t > \overline{q}^{\boldsymbol{F}}$.*

## 5. Optimal Manipulation

Under the assumption of large conflict, the project is terminated immediately in the efficient pure strategy equilibrium if $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}}]$. This collaboration failure is the most prominent manifestation of the moral hazard we study and is due to the fact that **S** cannot trust **F** to report honestly given such priors. Typically, moral hazard problems can be mitigated by contractual agreements that align the players' incentives. However, in our setting, contracts are precluded and the moral hazard must be alleviated by other means.

In this section, we show that strategic uncertainty that arises from **F** using a mixed strategy helps to align the players' incentives. In particular, we focus on the case of large conflict where strategic uncertainty is necessary to induce **F** to learn whenever $q_t > \overline{q}^{\mathbf{F}}$ since the intuitions as of why it is beneficial are easier to convey. In fact, the way in which strategic uncertainty can help mitigate moral hazard is the same regardless of whether the conflict is large or small (Proposition 6.1), but we postpone the discussion of small conflict until later. In the remainder of this section we assume that $\overline{q}^{\mathbf{F}} < \underline{q}^{\mathbf{S}}$ (large conflict) and that $\overline{q}^{\mathbf{F}} < q_0$, as otherwise any efficient equilibrium would be in pure strategies.

5.1. **Characterization of Equilibrium.** In this section we construct the efficient equilibrium with strategic uncertainty. In this equilibrium **F** initially uses a mixed strategy and in the end takes over the project. As strategic uncertainty is costly to both players (Lemma 4.3), intuitively, mixing stops at $\tau^*$ such that $q_{\tau^*} = \overline{q}^{\mathbf{F}}$, which is the earliest point at which **F** can be trusted. This is the point where moral hazard ceases to be a problem.

The mixing region before $\tau^*$ is separated into at most two phases that differ in terms of how **S** best responds to receiving the project. In the first phase, **S** examines the project and launches it only after observing a breakthrough himself. In this phase **F** postpones submitting the project because the amount of time **S** spends verifying that the project is good is increases over time. We refer to this phase as the *verification* phase. In the second phase, **S** randomizes between examining the project and launching it without observing a breakthrough. In this phase **F** postpones submitting the project because the probability that **S** launches the project immediately relative to the probability of **S** examines it is increasing over time. We refer to this phase as the *partial-trust* phase.

We will now derive **F**'s indifference conditions formally. **F**'s expected value from using the stopping time $\tau$ is given by

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

where $W_s^B$ is her *expected continuation value from submitting the project after observing a breakthrough occur at* $s$ and $W_\tau^{NB}$ is her *expected continuation value from submitting the project without observing a breakthrough occur at* $\tau$. Note that $W_t^{NB}$ and $W_t^B$ are jointly determined by **S**'s behavior at $t$:

$$
\begin{aligned}
W_t^B &= \sigma_t v^{\mathbf{F}} P^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right) + (1-\sigma_t)v^{\mathbf{F}} \\
W_t^{NB} &= q_t W_t^B - (1-q_t)(1-\sigma_t),
\end{aligned} \tag{2}
$$

where $P^{\mathbf{S}}(q)$ is the expectation of the discount factor at the first breakthrough when **S** learns with belief $q$. Formally,

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

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

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

Note that (3) is not a differential equation as $W_\tau^{NB}$ and $W_\tau^B$ are jointly determined by **S**'s behavior at $\tau$.

In the verification phase $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$, **S** does not launch the project ($\sigma_t = 1$), and (3) becomes

$$
\frac{c}{q_t v^{\mathbf{F}}} = \frac{dP^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right)}{dt} - rP^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right). \tag{3b}
$$

On the other hand, in the partial-trust phase $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ and equation (3) becomes:

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

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 (i.e., **S**'s learning with beliefs $\overline{q}^{\mathbf{S}}$). Note that $\Delta(q)$ is increasing in $q$ and hence $\Delta(q_t)$ is decreasing in $t$.

We can now present the main result of this section, which is the formal description of the efficient equilibrium in which **F** learns when $q_0 > \overline{q}^{\mathbf{F}}$.[9]

**Proposition 5.1.** *Let $q_0 > \overline{q}^{\boldsymbol{F}}$ and assume that the conflict is large. If $\omega > 0$ in an efficient equilibrium, then there exists $\tau^{**} \leq \tau^{*}$ such that*

(1) *If $\boldsymbol{S}$ receives the project at $t \leq \tau^{**}$, he learns according to $P^{\boldsymbol{S}}\left(q_t^{\boldsymbol{S}}\right)$ given by (3b) (verification phase).*
(2) *If $\boldsymbol{S}$ receives the project at $t \in (\tau^{**}, \tau^{*})$, he launches the project with probability $1-\sigma_t$ given by (3c) and learns according to $P^{\boldsymbol{S}}\left(\overline{q}^{\boldsymbol{S}}\right)$ otherwise (partial-trust phase).*
(3) *$\boldsymbol{F}$ mixes at all $t < \tau^{*}$ such that $q_t^{\boldsymbol{S}}$ is consistent with Bayes' law*

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

*at $\tau^{*}$, $\boldsymbol{F}$ takes over the project.*

The fact that there are at most two phases during the mixing range follows from the limitations on **S**'s ability to incentivize **F** to postpone submitting the project without observing a breakthrough in equilibrium. In particular, the sign of $\Delta(q)$ determines whether or not it is possible for **S** to incentivize **F** to postpone submitting the project, by breaking the indifference condition in a more favorable way for her.

If $\Delta(q_t) < 0$, **S** can incentivize **F** to postpone submitting the project only by increasing his learning time. To see this, recall first that **F**'s continuation

[9] We provide conditions for the existence of this equilibrium in Section 5.2.

value at $t$ cannot exceed her value from reporting honestly until $\omega$ (Lemma 4.3). As we explained after Proposition 4.2, under the assumption of large conflict this continuation value is less than the value from launching the project immediately while $q_t > \overline{q}^{\mathbf{F}}$. Hence, were **S** to mix between launching the project and maximal learning while $\Delta(q_t) < 0$, **F** would submit the project right away. On the other hand, while $\Delta(q_t) > 0$, **S** can incentivize **F** to postpone submitting the project either by increasing his learning time or by breaking his indifference in **F**'s preferred way.

A key consequence of **S**'s limitations in incentivizing **F** to postpone submitting the project is that **S**'s beliefs do not decrease in the mixing range irrespective of the sign of $\Delta(q_t)$. This is trivial when **S** incentivizes **F** by increasing his learning. If **S** mixes between maximal learning and immediate launching, then **F** prefers any outcome of the randomization to any amount of learning that is consistent with $q_t^{\mathbf{S}} \leq \overline{q}^{\mathbf{S}}$. This is because **S** can respond with a mixed strategy only if $\Delta(q_t) > 0$. Hence, if $q_t^{\mathbf{S}}$ decrease from $\overline{q}^{\mathbf{S}}$ to a lower belief, it would be profitable for **F** to submit the project before the decrease.

Since **S**'s beliefs increase, the exact structure of the equilibrium depends on the incentives generated when **F** takes over the project at $\tau^*$. Again, the sign of $\Delta(q_{\tau^*})$ determines this structure. If $\Delta(q_{\tau^*}) < 0$, **S** incentivizes **F** by increasing his learning in the neighborhood of $\tau^*$ and so $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ as $t$ approaches $\tau^*$. Since $q_t^{\mathbf{S}}$ is increasing, there is no partial-trust phase and the verification phase ends at $\tau^*$. On the other hand, $\Delta(q_{\tau^*}) > 0$ implies that $\sigma_{\tau^*}$ must equal zero, and **F** takes over the project after a partial-trust phase.[10] Whether or not an initial verification phase exists in this case depends on the time it takes $q_t$ to reach $\overline{q}^{\mathbf{F}}$. Nevertheless, as the transition between phases must be smooth, a partial-trust phase that follows a verification phase must start with **S** launching the project

[10] We abstract away from the non-generic case where $\Delta(\overline{q}^{\mathbf{F}}) = 0$.

with probability 0, and so $\tau^{**}$ is uniquely defined (it is derived formally in Proposition 5.2).

Finally, from Lemma 4.3 it follows that **F** taking over the project at $\tau^*$ is the unique efficient continuation equilibrium in pure or mixed strategies. In general, increasing continuation values at $\tau^*$ may not be consistent with behavior prior to this point. However, in our case, increasing continuation values at $\tau^*$ decreases the need for strategic uncertainty prior to $\tau^*$. Intuitively, this means that there are mutually beneficial new strategies prior to $\tau^*$ in which **F** submits the project without observing a breakthrough with (weakly) lower frequency that sustains the higher continuation values.[11]

5.2. **Existence of Equilibrium.** To establish the existence of the equilibrium we characterized in Proposition 5.1, we must show that there is a strategy for **F** that induces the beliefs $q_t^{\mathbf{S}}$ defined in the proposition according to Bayes' law. Bayes' law requires that $q_t^{\mathbf{S}} \geq q_t$ for all $t$. Since $q_t^{\mathbf{S}}$ is increasing and $q_t$ is decreasing, it suffices to check that $\lim_{t\to 0} q_t^{\mathbf{S}} > q_0$.[12] Furthermore, Proposition 5.1 implies that $G^{\mathbf{F}}$, if it exists, has a hazard ratio that is decreasing in $(0, \tau^*)$, and hence $G^{\mathbf{F}}(\tau^*) < 1$. Thus, **F**'s strategy can be completed by assigning the rest of the mass to the atom at $\omega$, and so the only conditions for existence are $q_0 < \overline{q}^{\mathbf{S}}$ and $\lim_{t\to 0} q_t^{\mathbf{S}} > q_0$.

If $\tau^{**} = 0$, the randomization consist only of a partial trust phase and existence is trivial. However, when $\tau^{**} > 0$ the equilibrium starts with a verification phase and the condition $\lim_{t\to 0} q_t^{\mathbf{S}} > q_0$ may be violated. Hence, to establish conditions for existence we must first derive $\tau^{**}$.

[11]Increasing **F**'s continuation utility at $\tau^*$ implies that his utility from submitting the project without observing a breakthrough at any $t < \tau^*$ increases to keep her indifferent. This, in turn, requires either that $q_t^{\mathbf{S}}$ increases at all $t$ during the verification phase or that **S** launches the project with higher probability in the partial-trust phase. Therefore, **F** must submit the project without observing a breakthrough with a (weakly) lower frequency. As this behavior is possible if the original equilibrium was in mixed strategy, if **F**'s continuation utility increases at $\tau^*$ there is always a strategy for **F** that induces the right beliefs for **S**.

[12]We don't analyze the non-generic case when $\lim_{t\to 0} q_t^{\mathbf{S}} = q_0$, in which $G^{\mathbf{F}}$ may have an atom at zero.

If $\Delta(q_{\tau^*}) < 0$ only the verification phase exists and $\tau^{**} = \tau^*$. When $\Delta(q_{\tau^*}) > 0$, the existence of a verification phase depends on the size of $\tau^*$. As the transition between phases, if it occurs, must be smooth, the existence of two phases requires that $\sigma_{\tau^{**}} = 1$ and $\sigma_{\tau^*} = 0$. Since $q_{\tau^*} = \overline{q}^{\mathbf{F}}$, it follows that whether there is a verification phase when $\Delta(q_{\tau^*}) > 0$ depends on whether the mixing lasts for longer than it takes $\sigma$ to decrease from one to zero. Formally, there is a smooth transition from (3b) and (3c) at some positive time if $H(\tau^*, 0) < 0$, where

$$H(\tau^*, t) \equiv \int_t^{\tau^*} \left( \frac{c}{q_s} + rv^{\mathbf{F}} \right) \times e^{-rs} \times (\Delta(q_s))^{\frac{1-\lambda}{\lambda}} ds - (\Delta(q_t))^{\frac{1}{\lambda}}$$

in which case $\tau^{**} > 0$ and it is given by $H(\tau^*, \tau^{**}) = 0$. By contrast, if $H(\tau^*, 0) \geq 0$ there is no possibility of a smooth transition between (3b) and (3c) and $\tau^{**} = 0$.

**Proposition 5.2.** *The mixed strategy equilibrium characterized in Proposition 5.1 exists if and only if $q_0 < \overline{q}^{\boldsymbol{S}}$ and*

$$\text{(4)} \qquad v^{\boldsymbol{F}} P^{\boldsymbol{S}}(q_0) + c \int_0^{\tau^{**}} \frac{e^{-ru}}{q_u} du < e^{-r\tau^{**}} \begin{cases} v^{\boldsymbol{F}} - \frac{1}{l(\overline{q}^{\boldsymbol{F}})} & \textit{if } \Delta(\overline{q}^{\boldsymbol{F}}) < 0, \\ v^{\boldsymbol{F}} P^{\boldsymbol{S}} \left( \overline{q}^{\boldsymbol{S}} \right) & \textit{if } \Delta(\overline{q}^{\boldsymbol{F}}) > 0. \end{cases}$$

Condition (4) can be interpreted as a cost-benefit analysis of **F**. The RHS is the expected value from submitting the project at the end of the verification phase and letting **S** behave as the equilibrium suggests. The LHS is **F**'s opportunity cost: the value she obtains from submitting the project at $t = 0$ and letting **S** take over the project given the prior $q_0$, and the expected discounted cost of learning until the end of the verification phase. Thus, the closer $q_0$ is to $\overline{q}^{\mathbf{F}}$, the less it costs **F** to induce strategic uncertainty. In fact, condition (4) shows that there exists a threshold $q^* > \overline{q}^{\mathbf{F}}$ such that the candidate mixed strategy equilibrium exists if and only if $q_0 < q^*$.

The existence of a threshold prior $q^*$ is illustrated in Figure 2, where we plot the likelihood ratios, $l(q_t)$, of the project being of good quality when both phases exist. The upward-sloping line indicates the evolution of **S**'s belief (upon receiving the project) in a mixed strategy equilibrium and the downward-sloping line indicates the evolution of **F**'s beliefs while learning (i.e., $q_t$). A mixed strategy equilibrium exists as long as **F** can induce $q_0^{\mathbf{S}} > q_0$, which in Figure 2 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.

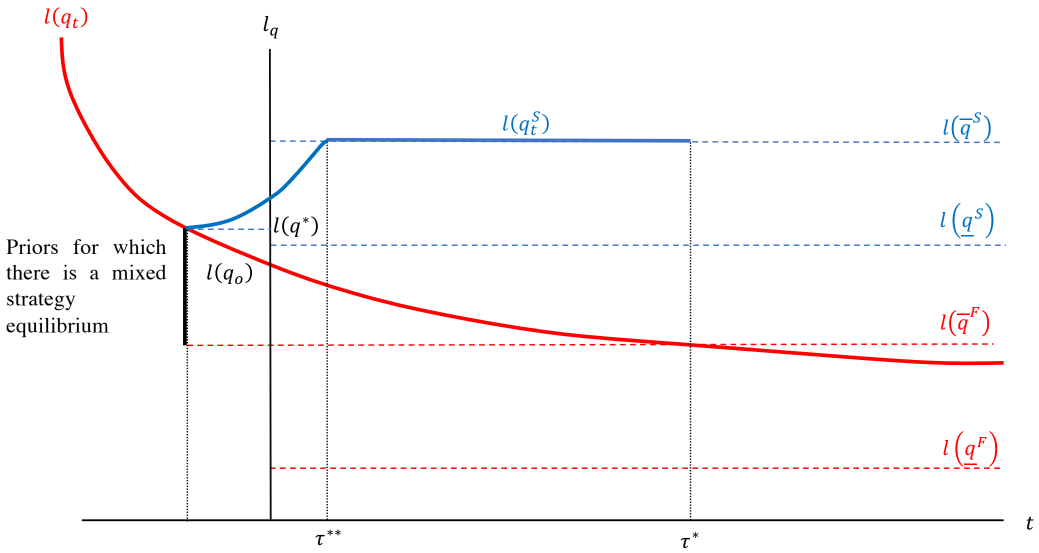


FIGURE 2. Dynamics of beliefs in the efficient mixed strategy equilibrium when $\Delta(\overline{q}^{\mathbf{F}}) > 0$ and $\tau^{**} \in (0, \tau^*)$.

Proposition 5.2 also enables us to determine whether the collaboration failure can be avoided for any $q_0 \in \left(\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}}\right)$. In particular, it is sufficient to check whether a mixed strategy equilibrium exists when $q_0 = \underline{q}^{\mathbf{S}}$.

**Corollary 5.3.** *A mixed strategy equilibrium exists for all* $q_0 \in (\overline{q}^{\boldsymbol{F}}, \underline{q}^{\boldsymbol{S}}]$ *if and only if condition* (4) *holds for* $q_0 = \underline{q}^{\boldsymbol{S}}$.

## 6. Efficiency

In this section we show that strategic uncertainty can be beneficial *regardless of the assumption of large conflict.* When $\overline{q}^{\mathbf{F}} < \underline{q}^{\mathbf{S}}$ the value of strategic uncertainty is self-evident: if $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$ in the unique equilibrium in pure strategies the project is terminated immediately, whereas strategic uncertainty enables the players to launch the project with positive probability. Since each player can unilaterally terminate the project, each player must be receiving a positive expected payoff in equilibrium. Moreover, even when $q_0 > \underline{q}^{\mathbf{S}}$ the probability of the project being launched in a pure strategy equilibrium is negligible if $q_0$ is close to $\underline{q}^{\mathbf{S}}$. Hence, continuity of the value functions in the mixed strategy equilibrium yields that strategic uncertainty is also beneficial for both players for some $q_0 > \underline{q}^{\mathbf{S}}$, if the mixed strategy equilibrium exists for those priors.

When the conflict is small it is unclear whether strategic uncertainty is efficient. First, it is not clear whether the equilibrium described in Proposition 5.1 is the efficient equilibrium in mixed strategies. Second, for every prior for which the mixed strategy equilibrium exists, **S** is willing to learn for some amount of time, which makes the comparison between pure and mixed strategy equilibria nontrivial. Moreover, for some priors, there may exist multiple equilibria in pure strategies in which both players learn.[13]

First, we establish that Proposition 5.1 characterizes the efficient mixed strategy equilibrium regardless of whether the conflict is large or small. To understand why this is the case, recall that the indifference condition (3) pins down the equilibrium behavior under strategic uncertainty when the continuation value is known. It turns out that strategic uncertainty is

[13] Such equilibria exist for some priors if $\underline{q}^{\mathbf{S}} < \frac{1}{1+v^{\mathbf{F}}}$. That is, **S** stops learning when **F** prefers terminating the project to launching it.

beneficial only insofar as it allows **F** to take over the project at $\tau^*$, which determines the efficient continuation value, and hence, the equilibrium behavior.

**Proposition 6.1.** *If $\boldsymbol{E}$ is an efficient equilibrium in which $\boldsymbol{F}$ uses a mixed strategy, then it is the equilibrium characterized in Proposition 5.1.*

Second, the mixed strategy equilibrium described in Proposition 5.1 is efficient whenever it exists and $q_0 > \overline{q}^{\mathbf{F}}$. This is because Proposition 5.2 implies that **F** is better off being under strategic uncertainty than submitting the project and free-riding on **S** learning optimally according to beliefs $q_0$. In fact, for some priors the mixed strategy equilibrium is the unique efficient equilibrium, as **S** also prefers the equilibrium with strategic uncertainty to his preferred equilibrium in pure strategies.

If the conflict is small and $q_0 > \overline{q}^{\mathbf{F}}$, then in any pure strategy equilibrium in which **F** learns, **S** must also learn **S** upon receiving the project at $\omega$. Therefore, **S**'s preferred pure strategy equilibrium has him free riding for some amount of time on **F** and then taking over the project himself. For $q_0$ close to $\overline{q}^{\mathbf{F}}$, the value of this free-riding is strictly smaller than what he would have gotten if the prior had been $\overline{q}^{\mathbf{F}}$ and **F** had taken over the project. Hence a standard continuity argument shows that **S** prefers strategic uncertainty to any pure strategy equilibrium given such priors. The following proposition, whose proof is omitted, states this result formally.

**Proposition 6.2.** *There exists $q^{**} > \overline{q}^{\boldsymbol{F}}$ such that for $q_0 \in (\overline{q}^{\boldsymbol{F}}, q^{**})$ the unique efficient equilibrium is the one characterized in Proposition 5.1.*

6.1. **The Economics of Optimal Manipulation.** Moral hazard problems are typically solved by leaving the informed party with informational rents. Under the assumption of large conflict and when the prior is in $(\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$ there are no rents to be appropriated in pure strategies as the project is terminated immediately. In this case, strategic uncertainty generates surplus that **F** can capture, thereby enabling both players to examine the project.

More generally, strategic uncertainty can be beneficial regardless of the size of the conflict or whether the players learn in pure strategy equilibria. The fundamental reason for this is that strategic uncertainty changes the nature of the players' learning times from strategic substitutes to strategic complements: the more **F** learns, the more **S** is willing to learn when he receives the project, which in turn incentivizes **F** to learn more.

Traditionally, the uninformed party designs incentives to mitigate the effect of moral hazard. In our sequential game **S** cannot design these incentives. Nevertheless, these incentives are present in a more subtle way. Recall that $W^B$ is **F**'s continuation for submitting the project after she observes a breakthrough and $W^{NB}$ is her continuation value without a breakthrough. In equilibrium $W^B - W^{NB}$ is increasing over time. Hence, it becomes more and more desirable for **F** to submit true breakthroughs instead of submitting the project without observing a breakthrough.

Strategic uncertainty also entails some costs and the size of these costs determines whether or not the efficient equilibrium involves manipulation. The first of these costs 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 termination of projects that **F** knows to be good. The second inefficiency is less intuitive: a project may be launched after its (costly) examination has not uncovered positive news. If **F** learns and then submits the project without observing a breakthrough in the partial-trust phase, **S** may launch the project immediately.

## 7. Concluding Remarks

Many collective decisions require a sequence of approvals to be implemented. Often each stage in the process is carried out by a different player with unique expertise who collects information before granting approval. Moreover, players' information is not often verifiable and cannot be credibly transmitted. Nevertheless, players participating late in the process can benefit from the information conveyed by observing other players' choices. However, this process of inference makes players in later stages susceptible to manipulation by players in earlier stages.

We study the moral hazard problem arising from this sequential collective decision-making process when players can collect information and disagree on the value of launching good projects in the absence of contractual agreements. This setup is appropriate for studying hiring in big companies, standardization processes, and modern product development.

We show that this moral hazard problem can be attenuated by strategic uncertainty and that this mixed strategy equilibrium resembles a quid-pro-quo type of implicit agreement. Players whose approval is needed early in the process collect information in such a way that players later in the process find it optimal to collect information themselves. This reciprocity is rooted in the dilution of information that follows from earlier players' random behavior.

Although mixing prevents later players from perfectly interpreting the actions they observe, it allows earlier players to transmit some information when making decisions. In the mixed strategy equilibrium information is obfuscated in such a way that incentives of earlier and later players align. In particular, earlier players behave more honestly over time and, as a result, the relevant beliefs of earlier and later players diverge, which in turn induces the later players to acquire more information. Thus, in the mixed strategy equilibrium, strategic uncertainty changes the nature of information acquisition from strategic substitutes to strategic complements.

The efficiency gains that strategic uncertainty brings hinge only on private learning and decreasing beliefs when no information arrives. This implies that our assumptions about homogeneous technologies that reveal conclusive evidence, and about projects requiring only one piece of positive evidence so there is perfect correlation between the dimensions that players care about, are not essential for our results. Therefore, our insights about the nature of efficient collaboration in sequential learning not only apply to unidimensional projects but also to projects of a more complex nature where multiple correlated dimensions need to be examined, and players have access to different information technologies.

## Appendix A. Proofs

### Proof of Proposition 4.1

A DM who decides to learn for $t$ units of time will launch the immediately after a breakthrough has occurred, and will terminate the project at $t$ if a breakthrough has not occurred by then. Therefore, her value from learning until time $t$ when her prior is $q_0$ is

$$\begin{aligned} EU(t,q_0) &= q_0\int_0^t \lambda e^{-\lambda s}\left(e^{-rs}v - \int_0^s ce^{ru}du\right)ds - \left(q_0e^{-\lambda t} + (1-q_0)\right)\int_0^t ce^{-ru}du \\ &= q_0\left(v-\frac{c}{\lambda}\right)\frac{\lambda}{r+\lambda}\left(1-e^{-(r+\lambda)t}\right) - (1-q_0)\frac{c}{r}\left(1-e^{-rt}\right). \end{aligned}$$

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 deriving the optimal learning time:

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

Since $l(q_t) = l(q_0)e^{-\lambda t}$ we obtain the cutoff beliefs $q_t = \underline{q}(v) = \frac{c}{\lambda v}$. The assumption that $c < \lambda\frac{v}{v+1}$ implies that the DM strictly prefers terminating the project to launching it at this threshold.

Learning is preferred to launching the project if $\mathbf{M}(q) \equiv \frac{EU^*(q)}{q} - \left(v^{\mathbf{F}} - \frac{1}{l(q)}\right) \geq 0$, where $EU^*(q)$ is the value of optimal learning with prior $q$, i.e.,

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

Note first that $\mathbf{M}(q)$ is continuous and decreasing in $q$. Second, it is positive at $\underline{q}(v)$ if $v < \frac{1}{l\left(\underline{q}(v)\right)}$, which holds since $c < \lambda\frac{v}{v+1}$. 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 $\mathbf{M}(\overline{q}(v)) = 0$. Since launching the project is better than terminating the project at $\overline{q}$, it follows that $\underline{q}(v) < \overline{q}(v)$.

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

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

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 4.2

In a pure strategy equilibrium, if $\mathbf{F}$ submits the project at $t < \tau^{\mathbf{F}}$, then $\mathbf{S}$ infers that a breakthrough has occurred and updates his belief to $q_t^{\mathbf{S}} = 1$. However, if $\mathbf{F}$ submits the project at $t = \tau^{\mathbf{F}}$, then $\mathbf{S}$ infers that a breakthrough has not occurred up to that point and therefore $q_t^{\mathbf{S}} = q_t$.

If $q_0 \leq \overline{q}^{\mathbf{F}}$, $\mathbf{F}$ taking over the project is an equilibrium as $\mathbf{F}$ 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 $q_0 > \overline{q}^{\mathbf{F}}$, then **F** submitting the project at $t = 0$ and **S** using an optimal policy for a DM is an equilibrium that is sustained by **S**'s off-path beliefs that lead to the project being terminated if submitted at any $t > 0$. Next, we show that for these priors there is no equilibrium in pure strategies with $\tau^{\mathbf{F}} > 0$. Since $q_0 > \overline{q}^{\mathbf{F}}$, there is no equilibrium in which only **F** learns, as **F** would rather submit the project at $t = 0$ and have it approved immediately. Consider an equilibrium in which both players learn. There are two possible events: either a breakthrough occurs before $q_t$ reaches $\underline{q}^{\mathbf{S}}$ or not. In the first case, **F** would rather deviate and submit the project at $t = 0$ (and have **S** respond by launching it). In the second case, the project is terminated when $q_t = \underline{q}^{\mathbf{S}} > \overline{q}^{\mathbf{F}}$, and it follows that **F** would also rather deviate by submitting the project at $t = 0$. □

## Proof of Lemma 4.3

We start this proof by establishing the following technical result.

**Lemma A.1.** *In equilibrium $\omega < \infty$. Moreover, if there is an atom at $\omega$, then $q_\omega = q^{\boldsymbol{S}}_\omega$; otherwise, there exists a sequence $t_n \to \omega$ such that* $\lim_{n\to\infty}(q^{\boldsymbol{S}}_{t_n} - q_{t_n}) = 0$.

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

By Bayes' law, if $t$ is an atom of $G^{\mathbf{F}}$ then $q_t^{\mathbf{S}} = q_t$. Hence, if either $\omega$ is an atom of $G^{\mathbf{F}}$ or there exists a sequence $t_n \to \omega$ such that every $t_n$ is an atom of $G^{\mathbf{F}}$, then it follows immediately that $q_\omega = q_\omega^{\mathbf{S}}$.

Otherwise, since Assumption 1 implies that $G^{\mathbf{F}}$ is non-singular, there exists $\tau < \omega$ such that $G^{\mathbf{F}}$ is absolutely continuous on $[\tau, \omega]$. Let

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

denote the hazard ratio of $G(\cdot)$ on $(\tau, \omega)$, and note that by Bayes' law $q_t^{\mathbf{S}} = q_t \frac{\lambda + h(t)}{\lambda q_t + h(t)}$ on this interval. If there exists a sequence $t_n \to \omega$ such that $\lim_{n\to\infty} g(t_n) > 0$, then for that sequence $\lim_{n\to\infty} h(t_n) = \infty$ and $\lim_{n\to\infty}(q_{t_n}^{\mathbf{S}} - q_{t_n}) = 0$. If there is no such sequence, and since $g(\omega) = 0$ and $g$ is left-continuous at $\omega$, there exists another 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 and $t_n$ in this sequence,

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

Hence, $\lim_{n\to\infty} h(t_n) = \infty$, and $\lim_{n\to\infty}(q_{t_n}^{\mathbf{S}} - q_{t_n}) = 0$. □

First, consider the case where **F** uses the stopping time $\omega$. If **F** submits the project at $\omega$, by Bayes law, we have that $q_\omega = q_\omega^{\mathbf{S}}$ and hence **S**'s response at $\omega$ in **E** is also a best response to the profile in which **F** reports honestly until $\omega$ (we denote this profile by **E'**). If **F** terminates the project at $\omega$ he gets a continuation utility of 0, which is less than his utility from having **S** learn with beliefs $q_\omega$. Hence, **F**'s continuation utility at $\omega$ in **E** is weakly lower than it is in **E'**. In **E'** or when **F** uses the stopping time $\omega$ in **E**, she submits the project before $\omega$ only after observing 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**.

Next, we show that **S** strictly prefers the outcome in **E'** to the outcome in **E**. If **F** observes a breakthrough at $t$ and submits the project at $t < \omega$ or submits the project without observing a breakthrough at $\omega$, then **S** weakly prefers his payoff in **E'** to that in **E**. If **F** submits the project at $t < \omega$ without observing a breakthrough, then **S** strictly prefers **F** continue learning until $\omega$ instead of submitting the project, since $q_0 < \overline{q}^{\mathbf{S}}$, and then continue with the equilibrium play in **E**. Note that this is exactly what occurs in **E'**.

If **F** does not use the stopping time $\omega$, by Lemma A.1 we have a sequence $t_n \to \omega$ for which $\lim_{n\to\infty}(q^{\mathbf{S}}_{t_n} - q_{t_n}) = 0$. Note that **F**'s continuation utilities in **E'** and **E** converge to one another along $t_n$, and so by the same arguments used above it follows that **E'** Pareto dominates **E**.[14] □

## Proof of Lemma 4.4

The first part of this lemma is immediate, and the second part follows from the following lemma.

**Lemma A.2.** *Under the assumption of large conflict, if $q_t > \overline{q}^{\mathbf{F}}$, then **F**'s utility from launching the project is strictly greater than his continuation utility in any equilibrium.*

*Proof.* By Lemma 4.3 **F**'s continuation utility in an arbitrary equilibrium **E** is weakly less than her continuation utility from learning honestly until $\omega$ and having **S** best respond. In the proof of Proposition 4.2 we established that if $q_t > \overline{q}^{\mathbf{F}}$, then **F**'s utility from launching the project is strictly greater than her utility from using any pure strategy and having **S** best respond. Hence, if $q_t > \overline{q}^{\mathbf{F}}$, then **F**'s continuation utility from **E** is strictly less than her utility from launching the project. □

## Proof of Proposition 5.1

By Lemma A.2, if $q_t > \overline{q}^{\mathbf{F}}$, then **F**'s continuation utility is less than it would be from launching the project immediately. It follows then that, if $q_t > \overline{q}^{\mathbf{F}}$, then **S** must learn with some probability if he receives the project. Moreover, if **S** launches the project with some positive probability upon

[14] Since $\omega$ is finite the strategy profile **E'** is well defined in this case.

receiving it, it must be the case that $\Delta(q_t) > 0$. Otherwise, if $\Delta(q_t) < 0$, then by submitting the project immediately, **F** would receive a higher utility than his utility form launching it immediately.

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

*Proof.* Recall that by equation (3), $W_t^{NB}$ must be differentiable and hence continuous. Let $\tau \in (0, \tau^*)$ be a discontinuity point in **S**'s belief. We establish that $q_t^{\mathbf{S}}$ is left-continuous at $\tau$; the proof that it is right-continuous is analogous. Note that since $q_t^{\mathbf{S}}$ is bounded there exists a sequence $t_n \to \tau$ for which $q_{t_n}^{\mathbf{S}}$ converges. Moreover, by Lemma A.2, $q_{t_n}^{\mathbf{S}} \leq \overline{q}^{\mathbf{S}}$ for every $n$ and $q_\tau^{\mathbf{S}} \leq \overline{q}^{\mathbf{S}}$.

First, consider the case where for such a convergent sequence $\lim_{n\to\infty} q_{t_n}^{\mathbf{S}} > q_\tau^{\mathbf{S}}$ (downward jump). If there exists $n_0$ such that $q_{t_n}^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ for every $n > n_0$, then the definition of $W_t^{NB} = q_t v^{\mathbf{F}} P^{\mathbf{S}}(q_t^{\mathbf{S}})$ directly implies that $W_t^{NB}$ is not continuous. Otherwise, if there is no such $n_0$ (for which $q_{t_n}^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$), there exists another sequence $t_{n_k} \to \tau$ such that $q_{t_{n_k}}^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ for which $\Delta(q_{t_{n_k}}) > 0$, and it follows that for any probability of launching $1 - \sigma_{t_{n_k}} \geq 0$,

$$
\begin{aligned}
(5)\qquad \lim_{k\to\infty} (1 - \sigma_{t_{n_k}}) \left(q_{t_{n_k}} v^{\mathbf{F}} - (1 - q_{t_{n_k}})\right) + \sigma_{t_{n_k}} q_{t_{n_k}} P^{\mathbf{S}}(\overline{q}^{\mathbf{S}}) > q_\tau P^{\mathbf{S}}(q_\tau^{\mathbf{S}}) \\
\Rightarrow \lim_{t_{n_k}\uparrow\tau} W_{t_{n_k}}^{NB} > W_\tau^{NB},
\end{aligned}
$$

which again contradicts the continuity of $W_t^{NB}$.

Now consider the case of a convergent sequence $\lim_{n\to\infty} q_{t_n}^{\mathbf{S}} < q_\tau^{\mathbf{S}}$ (upward jump). Assume without loss of generality that $q_{t_n}^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ for all $n$ such that $W_{t_n}^{NB} = v^{\mathbf{F}} q_{t_n} P^{\mathbf{S}}(q_{t_n}^{\mathbf{S}})$. If $q_\tau < \overline{q}^{\mathbf{S}}$, then $W_\tau^{NB} = v^{\mathbf{F}} q_\tau P^{\mathbf{S}}(q_\tau^{\mathbf{S}})$ and since $P^{\mathbf{S}}$ is strictly increasing it follows that $W_t^{NB}$ is discontinuous at $\tau$, a contradiction. If $q_\tau^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$, then $\Delta(q_t) > 0$ implies that $\lim_{n\to\infty} W_{t_n}^{NB} < W_\tau^{NB}$, and hence $W_t^{NB}$ is discontinuous at $\tau$.

We now show that $q_t^{\mathbf{S}}$ is continuous at $t = 0$. Because $G^{\mathbf{F}}$ is right continuous we must focus only on an atom at $t = 0$. In this case $q_0 = q_0^{\mathbf{S}}$

and there is a discontinuity at $t = 0$ if and only if $q_0 < \lim_{t \downarrow \tau} q_t^{\mathbf{S}}$. By the above argument about the continuity of the value functions can be applied we reach a contradiction.

To complete the proof, note that by equation (3b) we have that $P^{\mathbf{S}}(q_t^{\mathbf{S}})$ is strictly increasing in any interval when $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$, which implies that $q_t^{\mathbf{S}}$ is strictly increasing in such an interval. Since $q_t^{\mathbf{S}}$ is continuous and weakly less than $\overline{q}^{\mathbf{S}}$ for $t < \tau^*$, it follows that $q_t^{\mathbf{S}}$ is weakly increasing in $[0, \tau^*)$. □

Let $\tau^{**} = \min\{t : q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}\}$; then, points (1) and (2) of the proposition follow. Point (3) is established by Bayes' law.

Solving the piecewise differential equation (3) requires a boundary condition. By Lemma 4.4, as **F** mixes until (at least) $\tau^*$, we can use **F**'s continuation utility at $\tau^*$ as this boundary condition. Moreover, note that under the assumption of large conflict, the unique efficient equilibrium in pure or mixed strategies is that **F** takes over the project at $\tau^*$ (by Lemma 4.3). It remains to show that increasing continuation utilities at $\tau^*$ increases the players' utility at time zero, and that it is possible to support such play in $[0, \tau^*)$.

**Lemma A.4.** *The solution of equation* (3) *with boundary condition* $U$ *at* $\tau^*$ *characterizes an equilibrium if and only if* $q_0 < q_0^{\boldsymbol{S}}$ *and* $\sigma_0 \leq 1$.

*Proof of Lemma A.4.* Denote by $W_t^{NB}(U)$ the solution to (3) with boundary condition $U$ at $\tau^*$, denote by $q_t^{\mathbf{S}}(U)$ the belief associated with this solution, and denote by $\sigma_t(U)$ the probability **S** learns in this equilibrium.

Because the 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 are satisfied due to the construction of the differential equations. Recall that **F** is indifferent between all stopping times $t \in [0, \tau^*)$; thus, we are free to choose $G^{\mathbf{F}}$ so that Bayes law

$$\frac{g(t)}{1 - G^{\mathbf{F}}(t)} = \lambda \frac{l(q_t)}{l(q_t^{\mathbf{S}}) - l(q_t)} \tag{6}$$

holds. If $q_t^{\mathbf{S}} \leq q_t$, then (6) cannot hold for that $t$; however, if $q_t^{\mathbf{S}} > q_t$ for all $t \in [0, \tau^*)$, then there exists a $G^{\mathbf{F}}$ that satisfies (6). To see this, note that by integrating (6) 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 (6) we have that **F**'s strategy is

$$g^{\mathbf{F}}(s) = \lambda \frac{l(q_s)}{l(q_s^{\mathbf{S}}) - l(q_s)} \left(1 - G^{\mathbf{F}}(\tau^*)\right) e^{-\lambda \int_s^{\tau^*} \left(\frac{l(q_u)}{l(q_u^{\mathbf{S}}) - l(q_u)}\right) du}. \tag{7}$$

Integrating $h(s) \equiv \frac{g^{\mathbf{F}}(s)}{1 - G^{\mathbf{F}}(s)}$ we have that $G^{\mathbf{F}}(\tau^*) = 1 - e^{-\int_0^{\tau^*} h(u) du} < 1$. Thus, **F**'s strategy can be completed by assuming she plays the continuation strategy that provides her with payoff $U$ at $\tau^*$ with probability $1 - G^{\mathbf{F}}(\tau^*)$.

By Lemma A.3 $q_t^{\mathbf{S}}$ is increasing and continuous. Moreover, if $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$, then $\sigma_t = 1$. Since for all $t < \tau^*$ such that $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ it must be that $\Delta(q_t) > 0$, equation (3c) implies that $\dot{\sigma}_t < 0$ whenever $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$. Thus, it is sufficient to verify that $q_0 < q_0^{\mathbf{S}}$ and $\sigma_0 \leq 1$. □

The particular solution to a first-order differential equation such as (3) is uniquely determined by its boundary condition. 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 $\tau^*$; by the previous observation, $W_0^{NB}(U^2) > W_0^{NB}(U^1)$.

First, consider the case where $U^2 \leq q_{\tau^*} v^{\mathbf{F}} P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right)$. By Lemma A.2 it must be that $U^2 < q_{\tau^*} v^{\mathbf{F}} - (1 - q_{\tau^*})$ and at $\tau^*$ we have that $\sigma_{\tau^*}(U^i) = 1$ and so $W_{\tau^*}^{NB}(U^i) = q_{\tau^*} v^{\mathbf{F}} P^{\mathbf{S}}\left(q_{\tau^*}^{\mathbf{S}}(U^i)\right)$. To support this continuation equilibrium, $\tau^*$ must be approached in a verification phase, and it follows that for $i = 1, 2$, $q_t^{\mathbf{S}}(U^i) < \overline{q}^{\mathbf{S}}$ for all $t < \tau^*$ and so $W_0^{NB}(U^i) = q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0^{\mathbf{S}}(U^i)\right)$. Since $W_0^{NB}(U^2) > W_0^{NB}(U^1)$ and $P^{\mathbf{S}}(\cdot)$ is increasing we have that $q_0^{\mathbf{S}}(U^1) < q_0^{\mathbf{S}}(U^2)$. Thus, by Lemma A.4 the continuation utility $U^2$ can be supported by some equilibrium.

Next, consider the case where $U^2 > q_{\tau^*} v^{\mathbf{F}} P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right)$. Since this continuation utility is feasible only if $\Delta(q_{\tau^*}) > 0$ and $\Delta(\cdot)$ is monotonically increasing, it follows that $\Delta(q_0) > 0$, which implies that $W_0^{NB} \geq q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0^{\mathbf{S}}\right)$ (with a strict inequality if $q_0^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$). 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 $q_0^{\mathbf{S}}(U^1) = q_0^{\mathbf{S}}(U^2) = \overline{q}^{\mathbf{S}}$, then $W_0^{NB}(U^1) < W_0^{NB}(U^2)$ implies that $\sigma_0(U^1) > \sigma_0(U^2)$. Hence by Lemma A.4 the continuation utility $U^2$ can be supported by some equilibrium.

The previous argument establishes that increasing **F**'s continuation value at $\tau^*$ strictly increases her value at 0. We now focus on the impact of increasing both players' continuation utilities at $\tau^*$ on **S**'s utility at time zero. **S**'s expected utility from **E** is

$$\begin{aligned} V_0^{\mathbf{S}} = \int_0^{\tau^*} & \left(q_0\left(1 - G^{\mathbf{F}}(s)\right)\lambda e^{-\lambda s} + \left(q_0 e^{-\lambda s} + (1 - q_0)\right) g^{\mathbf{F}}(s)\right) e^{-rs} \\ & \left[q_s^{\mathbf{S}}\left(v^{\mathbf{S}} - \frac{c}{\lambda}\right) P^{\mathbf{S}}\left(q_s^{\mathbf{S}}\right) - (1 - q_s^{\mathbf{S}})\frac{c}{r}\left(1 - \left(\frac{l(\underline{q}^{\mathbf{S}})}{l(q_s^{\mathbf{S}})}\right)^{\frac{r}{\lambda}}\right)\right] ds \qquad (8) \\ & + e^{-r\tau^*}\left(1 - G^{\mathbf{F}}(\tau^*)\right)\left(q_0 e^{-\lambda\tau^*} + (1 - q_0)\right) V_{\tau^*}^{\mathbf{S}}, \end{aligned}$$

where $V^{\mathbf{S}}_{\tau^*}$ is $\mathbf{S}$'s continuation utility at $\tau^*$. Clearly, increasing $V^{\mathbf{S}}_{\tau^*}$ increases $V^{\mathbf{S}}_0$. Using (7), we can write the value function (8) as

$$(8b)\quad \frac{V^{\mathbf{S}}_0}{(1-q_0)\left(1-G^{\mathbf{F}}(\tau^*)\right)} = e^{-r\tau^*}\frac{V^{\mathbf{S}}_{\tau^*}}{1-q_{\tau^*}} + \int_0^{\tau^*} \frac{d\left(e^{-\lambda\int_s^{\tau^*}\left(\frac{l(q_u)}{l(q^{\mathbf{S}}_u)-l(q_u)}\right)du}\right)}{ds} e^{-rs} + \left[l(q^{\mathbf{S}}_s)\left(v^{\mathbf{S}}-\frac{c}{\lambda}\right)P^{\mathbf{S}}\left(q^{\mathbf{S}}_s\right)-\frac{c}{r}\left(1-\left(\frac{l(\underline{q}^{\mathbf{S}})}{l(q^{\mathbf{S}}_s)}\right)^{\frac{r}{\lambda}}\right)\right]ds.$$

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

## Proof of Proposition 5.2

Recall that we abstract away from the non-generic cases where $q_0 = \overline{q}^{\mathbf{S}}$ or $\Delta(\overline{q}^{\mathbf{F}}) = 0$.

Note that if $q_0 > \overline{q}^{\mathbf{S}}$ then, as discussed above, there is no strategy for $\mathbf{F}$ that induces beliefs $q^{\mathbf{S}}_t \leq \overline{q}^{\mathbf{S}}$ by Bayes' law. Hence, for the rest of the proof assume that $q_0 < \overline{q}^{\mathbf{S}}$. First, we characterize $\tau^{**}$ and $q_{\tau^{**}}$. When $\Delta(\overline{q}^{\mathbf{F}}) < 0$ there is no partial-trust phase and $\tau^{**} = \tau^*$; thus, we focus on 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, $\sigma_t$ is given by

$$\sigma_t \times e^{-rt} \times (\Delta(q_t))^{\frac{1}{\lambda}} = \int_t^{\tau^*} \left(\frac{c}{q_s} + rv^{\mathbf{F}}\right) \times e^{-rs} \times (\Delta(q_s))^{\frac{1-\lambda}{\lambda}} \, ds. \tag{9}$$

If

$$(\Delta(q_0))^{\frac{1}{\lambda}} \geq \int_0^{\tau^*} \left(\frac{c}{q_s} + rv^{\mathbf{F}}\right) \times e^{-rs} \times (\Delta(q_s))^{\frac{1-\lambda}{\lambda}} \, ds, \tag{10}$$

then $\sigma_t < 1$ for all $t > 0$, and there is no verification phase and so $\tau^{**} = 0$. Otherwise, if condition (10) does not hold, $\tau^{**}$ is given by setting $\sigma_t = 1$ in (9), which yields

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

Recall that by Lemma A.4, $q_0^{\mathbf{S}} > q_0$ and $\sigma_0 \leq 1$ are necessary and sufficient conditions for the existence of the mixed strategy equilibrium. First, consider the case where $\Delta(\overline{q}^{\mathbf{F}}) > 0$. If $\tau^{**} = 0$ existence is trivial; if $\tau^{**} > 0$, in the verification phase the differential equation (3b) reduces to

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

In this case, $q_0^{\mathbf{S}} > q_0$ if

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

Now consider the case where $\Delta(\overline{q}^{\mathbf{F}}) < 0$ and there is only a verification phase. The relevant differential equation (3c) reduces to

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

and the condition for existence reduces to

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

$\square$

### Proof of Proposition 6.1

To establish this result, we first show that mixed strategies cannot be part of an efficient equilibrium when $q_0 > \overline{q}^{\mathbf{S}}$ (Lemma A.5). We then show that, at some point in time moral hazard no longer hinders the interaction (this result requires using the technical Lemma A.6). Lemma A.7 shows that the interaction evolves in the same manner that we described in Proposition 5.1. Finally, we establish that in an efficient equilibrium where **F** uses a mixed strategy she must take over the project when she can be trusted to do so (Lemmas A.8 and A.9).

In this proof we often use the strategy profile where: 1) **F** reports honestly if $q_t > q_{\omega(\mathbf{E})}$ and submits the project otherwise, and 2) **S** best responds to **F**'s strategy. We refer to this profile as **E'**. Since submitting the project at any $t < \omega$ leads to the project being launched in **E'**, if it is profitable for **F** to deviate and submit the project at $t$ it is also profitable for her to submit the project at $t = 0$. Therefore, to verify that **E'** is an equilibrium it is sufficient to verify that **F**'s equilibrium payoff in **E'** at time zero is greater than her payoff from launching the project immediately.

**Lemma A.5.** *If $q_0 > \overline{q}^{\mathbf{S}}$, then **F** uses a pure strategy in any efficient equilibrium.*

*Proof.* Assume by way of contradiction that $q_0 > \overline{q}^{\mathbf{S}}$ and that **E** is an efficient equilibrium in which **F** uses a mixed strategy. By Bayes' law, $q_t^{\mathbf{S}} \geq q_t$ and so if **F** submits the project whenever $q_t > \overline{q}^{\mathbf{S}}$, then **S** will respond by launching the project immediately. Since the project will be launched whenever **F** submits it at any $t > 0$ such that $q_t > \overline{q}^{\mathbf{S}}$ or at time zero, **F** would rather submit the project at time zero in order to reduce her cost of learning. Hence, **F** does not submit the project without a breakthrough at any $t > 0$ whenever $q_t > \overline{q}^{\mathbf{S}}$ in **E**.

If $t = 0$ belongs to the support of $G^{\mathbf{F}}$, then there is an atom at $t = 0$. In this case, **F**'s equilibrium payoff at $t = 0$ is equal to the payoff from launching the project at $t = 0$. Note that the sum of **F**'s and **S**'s equilibrium payoffs can be no greater than the payoff of a DM with a value of $v^{\mathbf{F}} + v^{\mathbf{S}}$. Moreover, from Proposition 4.1 it follows that at any $q > \overline{q}^{\mathbf{S}}$ the value such a DM obtains from learning is strictly less than the DM's value from launching the project immediately. It follows that launching the project immediately Pareto dominates **E**. Therefore, $t = 0$ cannot be in the support of $G^{\mathbf{F}}$.

Since $t = 0$ is not in the support of **F**'s strategy, the support of **F**'s mixed strategy does not include any $t$ whenever $q_t > \overline{q}^{\mathbf{S}}$. Hence, from the same argument used to establish Lemma 4.3 it follows that **E'** Pareto dominates **E**. To see that **E'** is an equilibrium, note that in **E**, submitting the project at zero leads **S** to launch the project. Hence, **F**'s payoff in **E** is weakly greater than her payoff from launching the project immediately, which, in turn, implies that **F**'s payoff in **E'** is also weakly greater than his payoff from launching the project immediately. Thus, **E'** is an equilibrium that Pareto dominates **E**. □

By Lemma A.5, we focus on the case where $q_0 \leq \overline{q}^{\mathbf{S}}$ in the rest of this proof.

**Lemma A.6.** *There exists $t \leq \omega$ such that **F**'s continuation utility at $t$ in **E** is weakly greater than her payoff from launching the project immediately.*

*Proof.* Assume by way of contradiction that **F** strictly prefers launching the project to her continuation utility in **E** at all $t \leq \omega$. This implies that $q_t^{\mathbf{S}} \leq \overline{q}^{\mathbf{S}}$ for all $t \leq \omega$, and that $\Delta(q_t) \geq 0$ at all $t$ when $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$. Moreover, **F** must use a mixed strategy with a full support on $[0, \omega]$. First, consider the case where $\omega$ is not an atom of $G^{\mathbf{F}}$. By Lemma A.1 there exists $t_n \to \omega$ for which $\lim_{t_n \to \infty} q_{t_n}^{\mathbf{S}} = q_\omega$. By Bayes' law $q_{t_n}^{\mathbf{S}} \geq q_t$. It follows that there exists a subsequence $t_{n_k}$ along which $q_{t_{n_k}}^{\mathbf{S}}$ is decreasing. For every $t_{n_k}$ in the sequence $W_{t_{n_k}}^{NB} = q_{t_{n_k}} v^{\mathbf{F}} P^{\mathbf{S}}(q_{t_{n_k}}^{\mathbf{S}})$ and $W_{t_{n_k}}^{B} = v^{\mathbf{F}} P^{\mathbf{S}}(q_{t_{n_k}}^{\mathbf{S}})$. Hence both $W^{NB}$ and $W^{B}$ are decreasing in the subsequence. Since $W^{NB}$ is differentiable and hence continuous, it follows that **F**'s value function is decreasing around $\omega$, in contradiction to **F**'s indifference condition. Next, consider the case where $\omega$ is an atom. If there exists a decreasing sequence $q^{\mathbf{S}}_{t_n}$ for $t_n \to \omega$, the same argument used above establishes the claim. Otherwise, since $\Delta(q_t) \geq 0$ whenever $q^{\mathbf{S}}_t = \overline{q}^{\mathbf{S}}$ it follows that the value of submitting the project at $\omega$ where **S**'s beliefs are correct is strictly less than the value of submitting the project just prior to $\omega$. $\square$

Denote by $\psi$ the infimum of the set of $t \leq \omega$ such that **F**'s continuation utility at $t$ in **E** is weakly greater than her utility from launching the project (by Lemma A.6 this set is not empty). If $q^{\mathbf{S}}_t > \overline{q}^{\mathbf{S}}$ for some $t < \psi$, then **S** will launch the project upon receiving it and so **F**'s continuation utility will be at least the utility of launching the project. By the definition of $\psi$, **F**'s continuation utility is less than that, and so it must be that $q^{\mathbf{S}}_t \leq \overline{q}^{\mathbf{S}}$ for all $t < \psi$. This implies that **F** does not report honestly at any $t < \psi$. Moreover, if $q^{\mathbf{S}}_t = \overline{q}^{\mathbf{S}}$ for $t < \psi$, then it must be the case that $\Delta(q_t) \geq 0$ as otherwise by submitting the project **F** could obtain a utility that is greater than the utility obtained from launching the project.

If $\psi = 0$, then **E'** is an equilibrium. Moreover, by Lemma 4.3 **E'** Pareto dominates **E**. Hence, for the rest of the proof we focus on the case where $\psi > 0$. Note that since we are focusing on the case where $q_0 \leq \overline{q}^{\mathbf{S}}$, it follows that $q_\psi < \overline{q}^{\mathbf{S}}$.

**Lemma A.7.** *$q^{\mathbf{S}}_t$ is continuous and increasing at at all times $t \in [0, \psi)$ in **E**.*

*Proof.* The proof of this lemma is identical to that of Lemma A.3 with $\tau^*$ replaced by $\psi$. $\square$

**Lemma A.8.** *The project is terminated at $\omega$ in **E**.*

*Proof.* Assume by way of contradiction that **S** learns after receiving the project at $\omega$. By Lemma 4.3, **F**'s continuation utility at $\psi$ in **E** is weakly less than her continuation utility in **E'**. Moreover, **F**'s continuation payoff at $\psi$ in **E'** is weakly less than her payoff from submitting the project to **S** and having **S** best respond to beliefs $q_\psi$. Hence, **F**'s continuation utility at $\psi$ in **E** is no greater than $v^{\mathbf{F}} q_\psi P^{\mathbf{S}}(q_\psi)$.

By the definition of $W_t^{NB}$ and the fact that $\Delta(q_t) > 0$ whenever $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$, we have that $W_t^{NB} \geq v^{\mathbf{F}} q_t P^{\mathbf{S}}(q_t^{\mathbf{S}})$ for all $t \leq \psi$. Since Lemma A.7 implies that $q_0 < q_t^{\mathbf{S}}$ for all $t < \psi$, it follows that

$$\lim_{t \uparrow \psi} W_t^{NB} \geq v^{\mathbf{F}} q_\psi \lim_{t \uparrow \psi} P^{\mathbf{S}}(q_t^{\mathbf{S}}) > v^{\mathbf{F}} q_\psi P^{\mathbf{S}}(q_0) > v^{\mathbf{F}} q_\psi P^{\mathbf{S}}(q_\psi).$$

Hence, **F** would rather submit the project just prior to $\psi$ than at $\psi$. $\square$

Note that if **F** would rather launch the project than terminate it at $\omega$ (and receive a payoff of zero), then submitting the project just prior to $\omega$ is a profitable deviation regardless of **S**'s response. Hence, from Lemma A.8 it follows that $\psi < \omega$. Together with Lemma 4.3, this implies that **F**'s continuation utility at $t \leq \omega$ in **E** is weakly less than the value of optimal learning for a DM with a value of $v^{\mathbf{F}}$ and a prior of $q_t$. It follows that if $q_t > \overline{q}^{\mathbf{F}}$, **S** cannot launch the project immediately upon receiving it, which, in turn, implies that $\psi \geq \tau^*$.

The arguments used to establish Proposition 5.1 show that, if **E** is efficient, then the continuation equilibrium at $\tau^*$ must itself be efficient. Thus, to conclude the proof, we must show that **F** taking over the project at $\tau^*$ is the unique efficient equilibrium among all equilibria where the project is terminated at $\omega$.

**Lemma A.9.** *Assume that $q_0 = \overline{q}^{\boldsymbol{F}}$ and **F** must use a strategy in which she terminates the project at $\omega$. Then **F** taking over the project is the unique efficient equilibrium.*

*Proof.* Proposition 4.1 implies that **F** taking over the project is an equilibrium. Moreover, note that **S** considers launching the project immediately to be inferior to having **F** take over the project.

The argument used to establish Lemma 4.3 shows that **F**'s utility from any equilibrium in which she terminates the project at $\omega$, is weakly less than her utility from reporting honestly until $\omega$ and then terminating the project. Similarly, it shows that **S**'s utility from any equilibrium in which **F** terminates the project at $\omega$, is weakly less than the maximum between his utility from **F** reporting honestly until $\omega$ and then terminating the project, and his utility from launching the project immediately. □

## References

**Bergemann, Dirk, and Ulrich Hege.** 1998. "Venture capital financing, moral hazard, and learning." *Journal of Banking & Finance*, 22(6–8): 703–735.

**Bimpikis, Kostas, Shayan Ehsani, and Mohamed Mostagir.** 2019. "Designing dynamic contests." *Operations Research*, 67(2): 339–356.

**Bonatti, Alessandro, and Johannes Hörner.** 2011. "Collaborating." *American Economic Review*, 101(2): 632–663.

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

**Cetemen, Doruk, Ilwoo Hwang, and Ayça Kaya.** 2019. "Uncertainty-driven cooperation." *Available at SSRN 3000269.*

**Che, Yeon-Koo, and Johannes Hörner.** 2018. "Optimal design for social learning." *Quarterly Journal of Economics*, forthcoming.

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

**Dong, Miaomiao.** 2018. "Strategic experimentation with asymmetric information." *Penn State University, Working Paper.*

**Gerardi, Dino, and Lucas Maestri.** 2012. "A principal–agent model of sequential testing." *Theoretical Economics*, 7(3): 425–463.

**Green, Brett, and Curtis R Taylor.** 2016. "Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects." *American Economic Review*, 106(12): 3660–3699.

**Guo, Yingni.** 2016. "Dynamic delegation of experimentation." *American Economic Review*, 106(8): 1969–2008.

**Guo, Yingni, and Anne-Katrin Roesler.** 2016. "Private learning and exit decisions in collaboration." *Northwestern University, Working Paper.*

**Halac, Marina, Navin Kartik, and Qingmin Liu.** 2016. "Optimal contracts for experimentation." *The Review of Economic Studies*, 83(3): 1040–1091.

**Halac, Marina, Navin Kartik, and Qingmin Liu.** 2017. "Contests for experimentation." *Journal of Political Economy*, 125(5): 1523–1569.

**Henry, Emeric, and Marco Ottaviani.** 2019. "Research and the approval process: The organization of persuasion." *American Economic Review*, 109(3): 911–955.

**Hörner, Johannes, and Larry Samuelson.** 2013. "Incentives for experimenting agents." *The RAND Journal of Economics*, 44(4): 632–663.

**Keller, Godfrey, Sven Rady, and Martin Cripps.** 2005. "Strategic experimentation with exponential bandits." *Econometrica*, 73(1): 39–68.

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

**McClellan, Andrew.** 2017. "Experimentation and approval mechanisms." *New York University, Working Paper.*

**Moroni, Sofia.** 2019. "Experimentation in organizations." *University of Pittsburgh, Working Paper.*

**Murto, Pauli, and Juuso Välimäki.** 2011. “Learning and information aggregation in an exit game.” *The Review of Economic Studies*, 78(4): 1426–1461.

**Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille.** 2007. “Social learning in one-arm bandit problems.” *Econometrica*, 75(6): 1591–1611.

**Wolf, Christoph.** 2017. “Informative milestones in experimentation.” *University of Mannheim, Working Paper.*