SEQUENTIAL LEARNING

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

Abstract. We develop a model in which two players sequentially and publicly examine a project. In our model, the player who moves first can fabricate evidence to influence the second mover, which creates a moral hazard problem. We find that early strategic uncertainty can mitigate this problem. In particular, for intermediate prior beliefs about the project's quality, the unique Pareto-efficient equilibrium is in mixed strategies and consists of an early stage in which evidence may be fabricated and a later stage in which evidence is always authentic. Our findings shed light on the dynamics of R&D, quality assurance, and drug approval.
JEL Codes: D83, D82, O31, D70.
Keywords: Sequential Learning, Strategic Learning, Strategic Uncertainty.

1. Introduction

Economic ventures often require sequential interactions between players with different goals and abilities. For instance, launching a new product involves multiple stages of research and development inside the firm, newly developed software is often subject to a final stage of quality assurance to guarantee that it complies with industry standards, and drug approval processes begin with testing by the pharmaceutical company and end with

[†]Coller School of Management, Tel Aviv University, e-mail: yair.an@gmail.com
[§]Eitan Berglas School of Economics, Tel Aviv University, e-mail: dbird@tauex.tau.ac.il
[⊥]Department of Economics, University of Bristol, e-mail: s.oliveros@bristol.ac.uk
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, Alex Frug, Yoram Halevy, Sinem Hidir, Hugo Hopenhayn, Navin Kartik, Konrad Mierendorff, Zvika Neeman, Marco Ottaviani, Joel Sobel, Yossi Spiegel, Balazs Szentes, and Christoph Wolf. Yair Antler and Daniel Bird gratefully acknowledge the financial support of Israel Science Foundation grant 1917/19. Antler also thanks the Coller Foundation for financial support. Santiago Oliveros gratefully acknowledges the financial support of the Economic and Social Research Council, UK, grant ES/S01053X/1.

FDA scrutiny. The incentives and forces that shape these interactions are typically studied using a framework in which there is a clear separation between the players' roles. It is usually assumed that one player has access to information that she can then communicate to a second player who is in charge of making the final decision. However, in practice, the distinction between the players' abilities and responsibilities is blurrier, and learning is carried out in several stages.

When multiple players learn in sequence, the players who acquire information early in the sequence can misreport their findings to influence the choices of the players who act later in the sequence. For example, a pharmaceutical company may manipulate data to make a drug look more effective and improve the chances of it being approved by the regulating authority.[1] In a different context, a software engineer may claim that all minor bugs have been resolved to expedite the release of a new product. This moral hazard may lead players who learn late in the sequence to distrust the findings of early movers, which may result in excessive learning, welfare loss, and even in a holdup that jeopardizes the viability of the venture.

In this paper, we study the dynamics of efficient sequential learning in the realm of the above moral hazard. We develop a flexible continuous-time framework that enables us to study sequential interactions between two heterogeneous players who must decide whether to launch a project whose quality is initially unknown. In our model, at each point in time, the player who moves first (**F**, she) chooses between continuing to examine the project, or making an irreversible decision to either terminate it or submit it to the second mover (**S**, he). If **S** receives the project, he decides whether to launch the project or terminate it, and he can also acquire information before making this decision. We assume that **F** has access to a learning technology that generates a public signal according to an exponential distribution only if the project's quality is good (a "breakthrough"). In the main part of the paper, we assume that **S** looks for a breakthrough just like **F** does, which captures features relevant to R&D in large organizations. Later, exploiting the flexibility of our framework, we endow **S** with qualitatively different learning technologies and show that our findings extend to these settings. This allows us to accommodate diverse applications such as quality assurance, monitoring, and standardization processes.

[1] See, e.g., George and Buyse (2015) and Seife (2015).

The main tension in our model arises from our assumption that **F** can fabricate a breakthrough that is indistinguishable from an authentic one. Thus, when **S** receives the project, he cannot tell whether the breakthrough he observed was authentic or not. This moral hazard becomes a problem if there are prior beliefs about the quality of the project under which **F** wants to launch the project immediately (rather than examining it herself) but **S** wants to terminate the project (rather than examining it himself). We refer to the preferences and learning technologies that give raise to this disagreement as *large conflict*.[2]

When the conflict between the players is large, the implications of the moral hazard problem are stark. In particular, we show that in any pure strategy equilibrium, at most one of the players examines the project. Moreover, for the priors that manifest the large conflict, the project is terminated at $t = 0$ in every pure strategy equilibrium. The intuition for this collaboration failure is that, in pure strategy equilibria in which **F** learns, she is supposed to submit the project to **S** only after an authentic breakthrough. Hence, **S** interprets an early submission of the project as conclusive evidence that the project's quality is good. However, this creates an opportunity for **F** to submit the project early in order to manipulate **S** into launching it. We show that when the conflict between the players is large and the prior belief is such that **F** prefers to launch the project without examining it, she will use this opportunity and profitably deviate from any conjectured pure strategy equilibrium in which she is supposed to learn. In light of the futility of collaboration in pure strategies, it is natural to ask whether strategic uncertainty is beneficial for the players.

Our main result is that, for intermediate priors, in the unique Pareto-efficient equilibrium, **F** initially randomizes between fabricating a breakthrough (and submitting the project to **S**) and continuing to learn. As a result, **S** is suspicious when he receives the project and refrains from launching it without scrutinizing it first. If no breakthrough occurs while **F** learns, her belief drifts down and so, at some point in time τ^*, she no longer wants to launch the project without examining it further. Thus, if **F** does not submit the project in the initial part of the interaction, she finds it

[2]The same moral hazard problem arises in situations where **F** wants to launch the project but **S** wants to examine it (instead of terminating it). We analyze this case in Online Appendix B and show that our insights into efficient sequential collaboration remain valid when the conflict between the players is not large.

optimal to continue learning and refrains from fabricating a breakthrough after τ^*. As a result, if **S** receives the project after τ^*, he infers that it is because an authentic breakthrough occurred, and so he launches the project without scrutiny. We conclude that *initial* strategic uncertainty (i.e., an initial phase in which **F** may fabricate a breakthrough) may be required to support efficient sequential collaboration.

In the efficient equilibrium **F** learns even though she would rather launch the project without any costly examination. For **F** to be willing to learn in such a case, she must receive an additional benefit beyond the direct value of learning. In the absence of transfers, this benefit must result from **S**'s behavior, and, in the absence of commitment, **S**'s behavior and **F**'s behavior must be consistent with each other. We show that, in the efficient equilibrium, both players treat each other more favorably as time progresses. The exact manner in which this occurs depends on the learning technology that is available to **S**.

When **S** and **F** have qualitatively similar learning technologies, **S**'s response while **F** mixes can be decomposed into (at most) two 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 obtaining a breakthrough himself. Over time, the amount of time **S** devotes to examining the project gradually increases. In the partial-trust phase, if **S** receives the project, he randomizes between immediately launching it and further examining it. Over time, the probability of **S** launching the project immediately gradually increases.

In both phases, the probability that **S** eventually launches the project is increasing in the time that **F** spent examining the project. In that sense, **S**'s behavior becomes more favorable toward **F** over time. Importantly, for **S** to increase his investment in learning (which leads to a higher launching probability), he cannot become more pessimistic about the project's quality as time progresses. Thus, to sustain such belief dynamics, in the efficient equilibrium, as time progresses, **F** is less likely to fabricate a breakthrough.

Even though the mixed strategy equilibrium is the unique Pareto-efficient equilibrium, it exhibits two types of inefficiencies. First, **S** may terminate a project for which **F** obtained a breakthrough. Second, **S** may launch a project even after **F**'s costly information acquisition did not generate a breakthrough. We show that, while **F** mixes, the probability of terminating good projects goes down over time while the probability of launching bad projects goes up over time.

Our results hold for alternative specifications of **S**'s learning technology, which makes them applicable to a wider spectrum of applications. We study such additional applications in the final part of the paper. First, to capture the idea of quality assurance processes we assume that **S** looks for conclusive negative evidence about the project. Second, to capture the idea of standardization processes (e.g., drug approval) we assume that **S** gathers evidence on the nature of the breakthrough (i.e., whether it is fabricated or not). Finally, we consider the case where **S** cannot learn at all, which is the standard assumption in the literature on communication (with and without information acquisition). We show that despite some inevitable differences in **S**'s behavior, the nature of efficient collaboration for intermediate priors is the same in all of the above specifications. Moreover, the inefficiencies generated due to strategic uncertainly evolve in a similar fashion to the way they evolve in the baseline model. We conclude that efficient collaboration in sequential learning requires initial deceit and mistrust, and a gradual evolution toward honesty and trust.

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. In Section 5 we present our main results on the efficient equilibrium with strategic uncertainty in the baseline model. In Section 6 we apply our results to three different applications and Section 7 concludes. All proofs are relegated to Appendix A.

2. Literature Review

A large literature has focused on situations where the second mover has full decision rights over a project but cannot examine the project himself. The design of dynamic incentives in such cases 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 by Guo (2016), Henry and Ottaviani (2019), Bizzotto, Rüdiger and Vigier (2020), Escobar and Zhang (2020), and McClellan (2020) when they are not.[3] In contrast to these papers, we assume that neither player has commitment power, and focus on different forces that arise from the ability of the first mover to fabricate breakthroughs, the order of play, and the ability of the second player to learn.[4]

Our paper is also related to the strand of literature that studies multi-stage projects. Green and Taylor (2016) study a principal-agent problem where the first milestone is unobservable and the principal provides funds based on soft information. Wolf (2017) and Moroni (2019) investigate the design of incentives for a two-stage project where effort is unobservable. Our work differs from these papers in two key 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. 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.

Finally, 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.[5] Similar randomization effects appear in other models of strategic learning, but for different reasons than those in our sequential learning game. In Campbell, Ederer and Spinnewijn (2014), a player who observes a breakthrough wishes to conceal this information to incentivize the other player to keep exerting effort. Similarly, in Guo and Roesler (2016) and Dong (2018), a player who obtains a negative signal realization about the project wishes to conceal this information to induce the other player to experiment. In each of these papers, this information

[3]We discuss Escobar and Zhang (2020) in detail in Section 6.3.

[4]The literature on strategic experimentation (e.g., Décamps and Mariotti, 2004; Keller, Rady and Cripps, 2005; Rosenberg, Solan and Vieille, 2007; Bonatti and Hörner, 2011; Murto and Välimäki, 2011) also focuses on situations where players learn after observing (directly or indirectly) other players' learning choices. However, that literature studies the effect of asymmetric information on learning and free-riding when players move simultaneously.

[5]In Kremer, Mansour and Perry (2014) and Che and Hörner (2018), the principal strategically discloses partial 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.

can only be concealed if the informed player randomizes.[6] By contrast, in our model, the first mover cannot hold on to the project after an authentic breakthrough and the randomization helps to curtail the moral hazard problem.

3. The Model

Two players, **F** and **S**, jointly decide whether or not to launch a project whose unknown quality is either good or bad. 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 $v^i > 0$. We assume that the players share a common prior belief that the project is good, which we denote by q_0.

Before making the decision, the players may acquire information in a sequential manner. The first mover, **F**, publicly examines the project first. We model **F**'s acquisition of information as a continuous-time process and assume that the time at which the process starts and the time at which it ends are publicly observed. We normalize the time at which learning starts to 0 and assume that players do not discount the future.[7]

Throughout the paper we assume that if the project is good, then while **F** is examining the project a public signal ("breakthrough") arrives according to an exponential distribution, whereas if the project is bad the signal never arrives.[8] We denote **F**'s flow cost of learning by $c^{\mathbf{F}}$ and the arrival rate of the breakthrough while he is learning by $\lambda^{\mathbf{F}}$. In Proposition 4.1, we show that **F** is willing to learn for some beliefs if and only if $\frac{c^{\mathbf{F}}}{v^{\mathbf{F}}\lambda^{\mathbf{F}}} < \frac{1}{v^{\mathbf{F}}+1}$, and so we assume that this condition holds. Finally, we assume that once a breakthrough occurs **F**'s learning stops and the project must be submitted to **S**. This assumption rules out equilibria supported by **F**'s continued learning as a signaling device and we refer to it as *no money-burning*.

[6]Cetemen, Hwang and Kaya (2020) study how this effect evolves in the presence of public feedback.

[7]In a previous version of this paper (Antler, Bird and Oliveros, 2019), we show that all results hold when players do discount the future.

[8]Our results remain qualitatively unchanged if we allow **F** to obtain conclusive evidence that the project is bad from a public signal that arrives according to an exponential distribution, as long as the arrival rate of the positive signal is greater than that of the negative one.

While examining the project, **F** may decide to fabricate a breakthrough and submit the project to **S**. In our main analysis we make the simplifying assumption that fabricating a breakthrough is costless. Moreover, we assume that fabricated breakthroughs are indistinguishable from authentic ones. Hence, upon observing a breakthrough and receiving the project **S** may be interested in acquiring information himself.

To accommodate different types of applications, we consider three qualitatively different specifications for **S**'s learning technology. We postpone describing the learning technologies for the quality assurance and monitoring applications to Section 6 and, in this section, only describe **S**'s learning technology and strategy space in the baseline model, where **S** looks for positive news about the project just like **F** does.

We denote **S**'s flow cost of learning and the arrival rate of the breakthrough while he is learning by $c^{\mathbf{S}}$ and $\lambda^{\mathbf{S}}$, respectively, and assume that $\frac{c^{\mathbf{F}}}{v^{\mathbf{S}}\lambda^{\mathbf{S}}} < \frac{1}{v^{\mathbf{S}}+1}$. To highlight the moral hazard problem we further assume that $\frac{c^{\mathbf{S}}}{v^{\mathbf{S}}\lambda^{\mathbf{S}}}$ is sufficiently higher than $\frac{c^{\mathbf{F}}}{v^{\mathbf{F}}\lambda^{\mathbf{F}}}$. We refer to this situation as a *large conflict* and state this assumption formally in Section 4.2, where we also discuss its economic implications.

We wish to point out that some of our modeling assumptions are made mainly for expository purposes. In particular, this is the reason we assume that fabricating a breakthrough does not entail a direct or reputational cost, that **S** can never distinguish between authentic and fabricated breakthroughs, and that there is a large conflict between the players. In Section 5.4 we show that our insights hold if we relax these assumptions and discuss the implications of allowing for "money-burning" and private learning for our results.

3.1. **Strategies and Equilibrium.** Formally, **F** chooses a (potentially stochastic) stopping time $\tau^{\mathbf{F}}$ that represents the maximal time she will spend examining the project. The "no money-burning" assumption implies that if a breakthrough occurs prior to $\tau^{\mathbf{F}}$, **F** must submit the project to **S** immediately. At $\tau^{\mathbf{F}}$, **F** either fabricates a breakthrough *and* submits it to **S**, or submits it to **S** without fabricating a breakthrough, or terminates the project. As we shall see, in an efficient equilibrium, the latter two options are not used.

S's strategy is a function of the time t at which he receives the project and of whether he observed a breakthrough immediately before receiving it. We denote **S**'s strategy upon receiving the project after observing a breakthrough by $(\sigma_t, \tau_t^{\mathbf{S}})$. We let $1-\sigma_t$ denote the probability that **S** launches the project immediately, and $\tau_t^{\mathbf{S}} \geq t$ denote the stopping time at which **S**'s learning stops if he does not launch the project at t. If **S** chooses to learn and a breakthrough occurs, then he launches the project immediately; otherwise, he terminates the project at $\tau_t^{\mathbf{S}}$. Note that the choice $\tau_t^{\mathbf{S}} = t$ represents the choice to terminate the project upon receiving it at t. As pointed out above, in an efficient equilibrium **S** will always observe a breakthrough immediately before receiving the project. Since **S**'s off-path behavior is specified by our general assumptions about off-path behavior (see below), there is no need to define notation for his strategy in the case where he receives the project without observing a breakthrough.

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

In this paper, we characterize Pareto-efficient perfect Bayesian equilibria. We refer to these equilibria as *efficient equilibria*. To support efficient equilibria, it is without loss of generality to assume that **S** provides **F** with the lowest possible continuation payoff if **F** submits the project off the path of play. As **F** can guarantee herself a payoff of zero by terminating the project, we assume that if **S** receives the project off the path of play he terminates it. Note that this strategy is the best response to the belief that the project is bad.

Since $G^{\mathbf{F}}$ is Lebesgue measurable it is equivalent to a mixture of an absolutely continuous measure and a discrete measure. Let $g^{\mathbf{F}}$ denote the the derivative (whenever it exists) of $G^{\mathbf{F}}$. We denote the supremum of the support of $G^{\mathbf{F}}$ by $\omega(G^{\mathbf{F}}) = \inf \{t : G^{\mathbf{F}}([0,t]) = 1\}$. When there is no risk of confusion, we denote this supremum by ω. Finally, we denote by q_t the probability that the project is good conditional on **F** learning and no breakthrough occurring until time t, by $l(q) = \frac{q}{1-q}$ the likelihood ratio of the project being good under belief q, and by $q_t^{\mathbf{S}}$ **S**'s belief about the project's quality conditional upon receiving the project at time t.

4. Preliminary Analysis

4.1. **The Decision Maker's Problem.** We now study the behavior of a single decision maker (DM) who before deciding whether or not to launch a project can examine it using a *breakthrough* technology. We denote the DM's flow cost of learning by c and the arrival rate of the breakthrough by λ. Consider a DM who obtains a payoff of $v > 0$ from launching a good project and a payoff of -1 from launching a bad one. Clearly, if the DM decides to learn until (at most) time $t > 0$, she will launch the project immediately once a breakthrough occurs, and will terminate the project at t if a breakthrough has not occurred by then. Moreover, optimality requires 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 $c\,dt$ whereas the benefit is $q_t \lambda v\, dt$. Thus, the DM will terminate the project when

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

As learning is costly, the DM may prefer to launch the project without examination. In particular, she will do so if her prior is sufficiently high.

The next proposition characterizes the DM's optimal behavior, as well as **S**'s best response in the baseline version of the sequential learning game.

Proposition 4.1. *Assume that* $\frac{c}{\lambda} < \frac{v}{v+1}$. *There exist two cutoffs* $0 < \underline{q}(v, \frac{c}{\lambda}) < \overline{q}(v, \frac{c}{\lambda}) < 1$ *such that it is optimal for the DM to terminate the project if* $q_t \leq \underline{q}(v, \frac{c}{\lambda})$, *to examine the project if* $q_t \in (\underline{q}(v, \frac{c}{\lambda}), \overline{q}(v, \frac{c}{\lambda})]$, *and to launch the project if* $q_t \geq \overline{q}(v, \frac{c}{\lambda})$.

We let $\underline{q}^j = \underline{q}(v^j, \frac{c^j}{\lambda^j})$ and $\overline{q}^j = \overline{q}(v^j, \frac{c^j}{\lambda^j})$ 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.

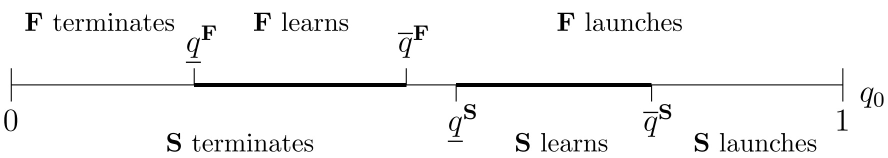

Figure 1. Learning regions for $\overline{q}^{\mathbf{F}} < \underline{q}^{\mathbf{S}}$ (*large conflict*).

As we shall see later, in an 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 a decision maker. Formally, **F** taking over the project is the following strategy profile: 1) **S** launches the project immediately upon receiving it, and 2) **F** uses the optimal policy for $(v^{\mathbf{F}}, \lambda^{\mathbf{F}}, c^{\mathbf{F}})$ as described in Proposition 4.1, where she breaks her indifference at $\overline{q}^{\mathbf{F}}$ in favor of learning.

4.2. **Large Conflict.** We say that there is a large conflict between the players if there exist beliefs for which, as a DM, **F** prefers launching the project without further investigation but **S** prefers terminating the project to examining or launching it. Formally, in our baseline model there exists a large conflict if $\overline{q}^{\mathbf{F}} < \underline{q}^{\mathbf{S}}$ (see Figure 1). Note that there is a large conflict if the difference between $v^{\mathbf{F}}$ and $v^{\mathbf{S}}$ is sufficiently large.[9] Thus, we expect large conflict to arise in contexts where **F**'s potential gain from the project is substantially greater than **S**'s potential gain.

The large conflict assumption has two implications. First, in the unique efficient equilibrium in pure strategies, at most one player examines the project.[10] Second, in every equilibrium in pure strategies, the project is terminated immediately when $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$, although both players strictly prefer that **F** examine the project rather than terminate it. **F** clearly prefers to learn about the project instead of terminating it at any $q_t > \underline{q}^{\mathbf{F}}$. While **S** would prefer to terminate the project as a single DM in this range of priors, he is better off having **F** incur the cost of learning and launch the project only after an authentic breakthrough occurs.

Proposition 4.2. *In the unique efficient equilibrium in pure strategies, **F** takes over the project at $t = 0$ if $q_0 \leq \overline{q}^{\mathbf{F}}$. Otherwise, at $t = 0$, she submits the project to **S**, whose behavior is as described in Proposition 4.1.*

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

[9] The assumption of large conflict, in terms of primitives, is equivalent to $\frac{1}{l\left(\frac{c^{\mathbf{F}}}{\lambda^{\mathbf{F}}}\right)} + \log\left(l\left(\frac{c^{\mathbf{F}}}{v^{\mathbf{F}}\lambda^{\mathbf{F}}}\right)\right) < l\left(\frac{c^{\mathbf{S}}}{v^{\mathbf{S}}\lambda^{\mathbf{S}}}\right) + \log\left(l\left(\frac{c^{\mathbf{S}}}{v^{\mathbf{S}}\lambda^{\mathbf{S}}}\right)\right)$.

[10] This equilibrium is unique except for the identity of the player who terminates the project when **S** does not learn.

players learn cannot exist when $q_0 \geq \underline{q}^{\mathbf{S}}$. To see this, note that if **F** submits the project at $\tau^{\mathbf{F}}$, then **S** examines the project until his belief reaches $\underline{q}^{\mathbf{S}}$ and terminates it at that point, unless a breakthrough has occurred earlier. However, under the assumption of large conflict, even if no breakthrough occurs while **S** learns and beliefs reach $\underline{q}^{\mathbf{S}}$, **F** would still prefer to launch the project. Hence, she would rather fabricate a breakthrough and submit the project before $\tau^{\mathbf{F}}$, misleading **S** to infer that an authentic breakthrough has occurred and thereby manipulating him into launching the project.

4.3. **The Merits of Honesty.** Proposition 4.2 implies that an appropriately chosen single decision maker obtains the same outcome that is obtained in the efficient pure strategy equilibrium of the sequential learning game. For priors in $(\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$, this outcome is Pareto inferior to what the players could obtain if they were to trust each other: the project is terminated at $t = 0$ although both players would be better off if **F** were to examine the project. The reason for this inefficiency is that were **S** to trust **F** to learn when $q_0 > \overline{q}^{\mathbf{F}}$, **F** would fabricate a breakthrough and manipulate **S** into launching the project. This suggests that, for **F** to learn in equilibrium, **S** cannot trust her blindly and launch the project whenever it is submitted. In other words, **S** must have some doubt as to whether a breakthrough is authentic or not. Inducing this uncertainty in **S** requires **F** to use a mixed strategy.

While **F**'s randomization may sustain equilibria where the players examine the project, it also introduces a new type of inefficiency: if **S** is uncertain as to whether a breakthrough is authentic or not, he may terminate the project even after an authentic breakthrough has occurred. Hence, even when equilibria with strategic uncertainty can mitigate the moral hazard problem, they do not lead to first-best outcomes. The next result formalizes this argument.

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 $\tau^{\boldsymbol{F}} = \omega(\hat{G}^{\boldsymbol{F}})$ and **S** best responds to that strategy.*

Lemma 4.3 states that any equilibrium in which there is strategic uncertainty is Pareto-dominated by a profile of pure strategies in which **F** learns honestly until $\omega(\hat{G}^{\mathbf{F}})$ and **S** best responds. If **F** did not have the ability to fabricate breakthroughs the latter profile would be an equilibrium. However, when **F** can fabricate breakthroughs, this profile of pure strategies is not necessarily an equilibrium. We can conclude that both players are worse off due to the moral hazard problem.

Lemma 4.3 has two important implications. First, if $q_0 \leq \overline{q}^{\mathbf{F}}$, the efficient equilibrium is in pure strategies: for such priors **F** can be trusted to take over the project (Proposition 4.2), which renders strategic uncertainty unnecessary. Second, in equilibrium, **S** cannot launch the project with probability one upon receiving it if $q_t > \overline{q}^{\mathbf{F}}$. The reason for this is that for such beliefs, **F** strictly prefers launching the project to examining it. However, by Lemma 4.3, **F**'s payoff from examining the project and launching it once a breakthrough occurs is weakly greater than her payoff in a mixed strategy equilibrium. Since **F**'s equilibrium payoff is at least her payoff from submitting the project, transitivity implies that, in equilibrium, **S** cannot launch the project with probability one upon receiving it if $q_t > \overline{q}^{\mathbf{F}}$.

Corollary 4.4. *If $q_0 \leq \overline{q}^{\boldsymbol{F}}$, then any efficient equilibrium is in pure strategies. Furthermore, for any t such that $q_t > \overline{q}^{\boldsymbol{F}}$ there is no equilibrium in which $g^{\boldsymbol{F}}(t)$ exists and is equal to zero. That is, if $q_t > \overline{q}^{\boldsymbol{F}}$, then **F** must be indifferent between fabricating a breakthrough and continuing to learn.*

5. Optimal Manipulation

Under the assumption of large conflict, the project is terminated immediately in every 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 problem we are studying and is due to the fact that, for such priors, **S** is unwilling to examine the project himself and cannot trust **F** to refrain from fabricating breakthroughs. 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 mitigated by other means.

In this section, we show that strategic uncertainty helps align the players' incentives. In the remainder of this section, we assume that $q_0 > \overline{q}^{\mathbf{F}}$; that is we focus on priors for which strategic uncertainty is necessary to induce **F** to learn.

5.1. **Characterization of Equilibrium.** We now construct an equilibrium with strategic uncertainty and we later show that, if it exists, then it is the unique Pareto-efficient equilibrium. In this equilibrium, **F** initially mixes continuously between continuing to learn and fabricating a breakthrough, and eventually takes over the project. As strategic uncertainty is costly to both players (Lemma 4.3), mixing stops at τ^* such that $q_{\tau^*} = \overline{q}^{\mathbf{F}}$, which is the earliest point at which **F** can be trusted. Note that the assumption that learning is public implies that τ^* is common knowledge.

The mixing region before τ^* can be decomposed into at most two phases that differ in terms of how **S** responds to receiving the project. In the first phase, **S** examines the project and launches it only after a breakthrough occurs while he is learning. In this phase, **F** may refrain from fabricating breakthroughs because the amount of time **S** spends verifying that the project is good 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 immediately. In this phase, **F** may refrain from fabricating breakthroughs because the probability that **S** launches the project immediately increases 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 τ is given by

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

where W_s^B is her *expected continuation value from submitting the project after an authentic breakthrough occurs at* s and W_τ^{NB} is her *expected continuation value from submitting the project after fabricating a breakthrough at* τ. Note that W_t^{NB} and W_t^B are jointly determined by **S**'s behavior at t:

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

where $P^{\mathbf{S}}(q)$ is the probability that, if the project is good, a breakthrough occurs if **S** begins learning (optimally) with belief q. It can be shown that $P^{\mathbf{S}}(q) \equiv 1 - \frac{l(\underline{q}^{\mathbf{S}})}{l(q)}$ if $q \geq \underline{q}^{\mathbf{S}}$, and 0 otherwise.

As $\mathbf{F}$ must be indifferent between all stopping times $\tau < \tau^*$, it follows that $V^{\mathbf{F}}(\tau)$ is constant in τ for all $\tau < \tau^*$. Taking the derivative of $V^{\mathbf{F}}(\tau)$ and equating it to zero yields

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

Note that (3) is a partial differential equation. Nevertheless, $\mathbf{S}$'s equilibrium behavior makes (3) a piecewise differential equation in W^B.

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

$$\frac{c^{\mathbf{F}}}{q_t} = \frac{dW_t^B}{dt}. \tag{3b}$$

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

$$\left(1 + \frac{1}{l(q_t)\Delta(q_t)}\right)\frac{c^{\mathbf{F}}}{q_t} = \frac{dW_t^B}{dt}, \tag{3c}$$

where

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

is the (scaled) difference between $\mathbf{F}$'s payoff from launching the project immediately and her payoff from free-riding on $\mathbf{S}$'s *maximal learning* (i.e., $\mathbf{S}$'s learning with belief $\overline{q}^{\mathbf{S}}$). 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 paper, which is the formal description of the efficient equilibrium in which $\mathbf{F}$ learns when $q_0 > \overline{q}^{\mathbf{F}}$. We postpone the discussion about existence of this equilibrium to Section 5.2.

Proposition 5.1. *Let $q_0 > \overline{q}^{\mathbf{F}}$. If in an efficient equilibrium $\omega > 0$, then there exists $\tau^{**} \leq \tau^*$ such that*

(a) *If $\mathbf{S}$ receives the project at $t \leq \tau^{**}$, he learns according to $P^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right)$ as given in (3b) (verification phase).*

(b) *If $\mathbf{S}$ receives the project at $t \in (\tau^{**}, \tau^*)$, he launches the project with probability $1 - \sigma_t$ as given in (3c), and otherwise learns according to $P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right)$ (partial-trust phase).*

(c) *$\mathbf{F}$ mixes at all $t < \tau^*$ such that $q_t^{\mathbf{S}}$ is consistent with Bayes' law*

$$\frac{g^{\mathbf{F}}(t)}{1 - G^{\mathbf{F}}(t)} = \lambda^{\mathbf{F}} \frac{l(q_t)}{l(q_t^{\mathbf{S}}) - l(q_t)}, \tag{4}$$

*and, at τ^*, **F** takes over the project.*

The evolution of beliefs in the equilibrium characterized in Proposition 5.1 is illustrated in Figure 2 for parameters under which both the verification phase and the partial-trust phase exist, i.e., $\tau^{**} \in (0, \tau^*)$. In particular, the figure shows that the players' *relevant* beliefs diverge. The red downward-sloping line indicates the evolution of the (public) belief conditional on no breakthrough occurring while **F** is learning and the blue upward-sloping line indicates the evolution of **S**'s belief upon receiving the project required to maintain **F**'s indifference. Between 0 and τ^{**}, the verification phase, **S**'s belief is strictly increasing, and between τ^{**} and τ^*, the partial-trust phase, **S**'s belief is constant at $\overline{q}^{\mathbf{S}}$.

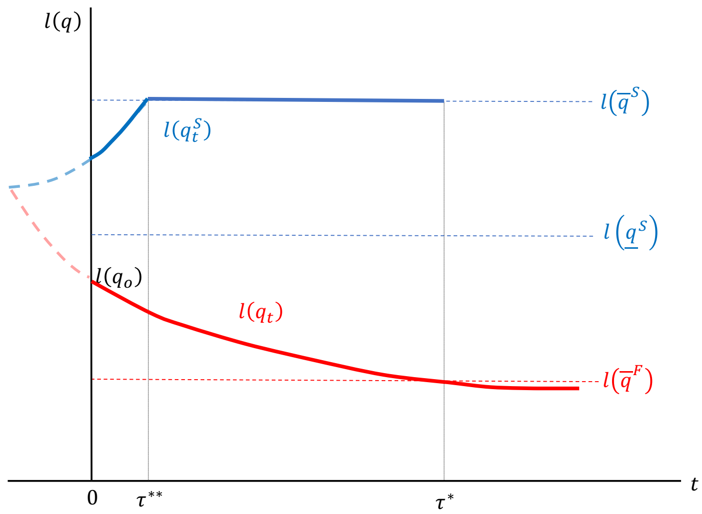

FIGURE 2. Dynamics of beliefs in the efficient mixed strategy equilibrium when $\tau^{**} \in (0, \tau^*)$.

Proposition 5.1 is silent about the exact value of τ^{**}. Note that if $\tau^{**} = \tau^*$, then a partial-trust phase does not exist, whereas if $\tau^{**} = 0$, then a verification phase does not exist. Both phases exist if and only if $\tau^{**} \in (0, \tau^*)$. To determine which phases exist, it is important to understand whether **S** can incentivize **F** to continue learning by breaking his indifference in a way that is more favorable to her. It turns out that this depends on the sign of $\Delta(q)$.

If $\Delta(q_t) < 0$, then the only way **S** can incentivize **F** to postpone submitting the project is by increasing his learning time. To see this, recall first that **F**'s continuation value at t cannot exceed her value from learning honestly until ω and having **S** best respond to that strategy (Lemma 4.3). As we established in the proof of Corollary 4.4, under the assumption of large conflict, the latter continuation value is less than the value from launching the project immediately if $q_t > \overline{q}^{\mathbf{F}}$. Hence, since $\Delta(q_t) < 0$, were **S** to mix between maximal learning and immediate launch of the project, **F**'s continuation value at t would be greater than the value she obtains from immediately launching the project, which is a contradiction. On the other hand, if $\Delta(q_t) > 0$, **S** can incentivize **F** to postpone submitting the project either by gradually increasing his learning time over time or by gradually increasing the probability of launching the project over time.

A key consequence of **S**'s limitations in incentivizing **F** to refrain from fabricating breakthroughs is that **S**'s belief does not decrease in the mixing region irrespective of the sign of $\Delta(q_t)$. If **S** incentivizes **F** by increasing his learning time this is immediate. If **S** mixes between maximal learning and immediate launching, then **F** strictly prefers any such lottery 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 $q_t^{\mathbf{S}}$ decreases from $\overline{q}^{\mathbf{S}}$ to a lower belief, it would be profitable for **F** to submit the project before $q_t^{\mathbf{S}}$ decreases.

Since **S**'s belief is weakly increasing, the exact structure of the equilibrium depends on **F**'s preferences when she takes over the project at τ^* (i.e., the sign of $\Delta(q_{\tau^*})$), which determine the boundary condition of the differential equation (3). If $\Delta(q_{\tau^*}) < 0$, **S** cannot incentivize **F** by increasing the probability of launching the project in the neighborhood of τ^* and so $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ as t approaches τ^*. Since $q_t^{\mathbf{S}}$ is increasing, there is no partial-trust phase ($\tau^* = \tau^{**}$) and the verification phase ends at τ^*. On the other hand, $\Delta(q_{\tau^*}) > 0$ implies that in the neighborhood of τ^*, **S** launches the project with a probability that approaches one, and so **F** takes over the project after a partial-trust phase.[11] It follows that a partial-trust phase exists ($\tau^{**} < \tau^*$) if and only if $\Delta(q_{\tau^*}) > 0$.

If $\Delta(q_{\tau^*}) > 0$, then an initial verification phase exists if and only if $\tau^{**} > 0$. As the transition between phases must be smooth, the existence of both phases requires that $\sigma_{\tau^*} = 0$ and, at the same time, there is some $\tau > 0$

[11] We abstract away from the non-generic case where $\Delta(q_{\tau^*}) = 0$.

such that $\sigma_\tau = 1$: this time τ is essentially τ^{**}. The differential equation (3c) implies that σ_t is decreasing, and so whether there is a verification phase when $\Delta(q_{\tau^*}) > 0$ depends on whether **F**'s mixing takes longer than it takes σ to decrease from one to zero.

More formally, in a partial-trust phase, the indifference condition (3) yields that

$$\sigma_t = \int_t^{\tau^*} \frac{c^{\mathbf{F}}}{q_s\,\Delta(q_s)} ds \tag{5}$$

at any[12] $t \in (\tau^{**}, \tau^*)$. The existence of a verification phase depends on the value of this integral when $t = 0$. If the integral in (5) evaluated at $t = 0$ is less than 1, it follows that $\sigma_t < 1$ at any $t \geq 0$ and there cannot be a verification phase. On the other hand, if the integral in (5) evaluated at $t = 0$ is greater than 1, σ_t must be equal to 1 at some point $\tau^{**} \in (0, \tau^*)$ and both phases exist. The next lemma formalizes this discussion.

Lemma 5.2. *In the equilibrium described in Proposition (5.1),*

(1) *if $\Delta(\overline{q}^{\boldsymbol{F}}) < 0$, there is only a verification phase, i.e., $\tau^{**} = \tau^*$;*
(2) *if $\Delta(\overline{q}^{\boldsymbol{F}}) > 0$, there is a partial-trust phase, i.e., $\tau^{**} < \tau^*$, and if*

$$\int_0^{\tau^*} \frac{c^{\boldsymbol{F}}}{q_s\,\Delta(q_s)} ds > 1,$$

there is also a verification phase, i.e., $0 < \tau^{**}$.

To see why in the efficient equilibrium **F** takes over the project at τ^*, recall that by Lemma 4.3 this is is the unique efficient continuation equilibrium in pure or mixed strategies. In general, increasing continuation values at τ^* may not be consistent with equilibrium behavior prior to this point as the indifference condition (3) may not hold. However, in our case, the equilibrium values follow a pairwise continuous first-order differential equation, and so increasing continuation values at τ^* increases **F**'s value of learning at all $t \in (0, \tau^*)$. Intuitively, increasing the continuation values at τ^* decreases the need for strategic uncertainty prior to τ^*, and so **F** fabricates breakthroughs with a (weakly) lower frequency prior to τ^*. This lower probability of faking a breakthrough while **F** mixes sustains the higher continuation value of **F** and also increases **S**'s equilibrium payoff.

[12]This equivalence is derived in the proof of Proposition 5.3.

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 $\mathbf{F}$ that induces the beliefs $q_t^{\mathbf{S}}$ according to Bayes' law, i.e., a strategy and belief that satisfy (4). Consistency with Bayes' law requires that $q_t^{\mathbf{S}} > q_t$ for all t. Since $q_t^{\mathbf{S}}$ is increasing and q_t is decreasing, it suffices to check that[13] $q_0^{\mathbf{S}} > q_0$. 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, $\mathbf{F}$'s strategy can be completed by assigning the rest of the mass to the atom at ω, and so the only condition for existence is $q_0^{\mathbf{S}} > q_0$.

If $\tau^{**} = 0$, then the mixing region consists of only a partial-trust phase, $q_0^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$, and existence boils down to the simple condition $q_0 < \overline{q}^{\mathbf{S}}$. However, if $\tau^{**} > 0$, the equilibrium starts with a verification phase in which $q_t^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ and so the condition $q_0^{\mathbf{S}} > q_0$ is harder to verify. The following proposition shows that existence of the equilibrium reduces to a cost-benefit analysis for $\mathbf{F}$.

Proposition 5.3. *The mixed strategy equilibrium characterized in Proposition 5.1 exists if and only if* $q_0 < \overline{q}^{\boldsymbol{S}}$ *and*

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

The RHS of (6) is $\mathbf{F}$'s expected value (scaled by $\frac{1}{q_0}$) from submitting the project at the end of the verification phase and letting $\mathbf{S}$ behave as the equilibrium suggests. The LHS is the sum of $\mathbf{F}$'s expected cost of learning until the end of the verification phase and her opportunity cost, i.e., her expected value from submitting the project at $t = 0$ and letting $\mathbf{S}$ take over the project given the prior q_0. Thus, the closer q_0 is to $\overline{q}^{\mathbf{F}}$, the less it costs $\mathbf{F}$ to induce strategic uncertainty. In fact, condition (6) implies that there exists a threshold $q^* > \overline{q}^{\mathbf{F}}$ such that the mixed strategy equilibrium exists if and only if $q_0 < q^*$.

Proposition 5.3 also enables us to determine whether the collaboration failure can be avoided for any $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}}]$. In particular, it is sufficient to check whether (6) holds at $q_0 = \underline{q}^{\mathbf{S}}$.

[13]We don't analyze the non-generic case of $q_0^{\mathbf{S}} = q_0$ in which $G^{\mathbf{F}}$ may have an atom at zero.

Corollary 5.4. *There exists $q^* \in (\overline{q}^{\boldsymbol{F}}, \overline{q}^{\boldsymbol{S}})$ such that the mixed strategy equilibrium exists if and only if $q_0 \in (\overline{q}^{\boldsymbol{F}}, q^*)$. Moreover, $\underline{q}^{\boldsymbol{S}} < q^*$ if*

$$c^{\boldsymbol{F}}\tau^{**} + \frac{c^{\boldsymbol{F}}}{\lambda}\frac{1-e^{-\lambda\tau^{**}}}{l(q_{\tau^{**}})} < v^{\boldsymbol{F}}P^{\boldsymbol{S}}\left(\overline{q}^{\boldsymbol{S}}\right) + \min\left\{0, \Delta(\overline{q}^{\boldsymbol{F}})\right\}. \tag{6b}$$

The existence of q^* is illustrated in Figure 2 by the intersection of the dashed parts of the two lines. In particular, in Figure 2, we assume that (6b) holds and so the equilibrium with strategic uncertainty exists for some $q_0 > \underline{q}^{\mathbf{S}}$ under which **S** is willing to learn.

5.3. **Efficiency.** In this section we show that the mixed strategy equilibrium characterized in Proposition 5.1 is the unique Pareto-efficient equilibrium whenever it exists. If $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}}]$, then the project is terminated at $t = 0$ in every pure strategy equilibrium (Proposition 4.2). Hence, if the mixed strategy equilibrium exists for such a prior, then it is the unique Pareto-efficient equilibrium. We now focus on the case where $q_0 > \underline{q}^{\mathbf{S}}$.

In the mixed strategy equilibrium, submitting the project at $t = 0$ is an optimal strategy for **F**. If there is an initial verification phase, this implies that **F**'s equilibrium payoff is $q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0^{\mathbf{S}}\right)$. Otherwise, it must be the case that $\Delta(q_0) > 0$, and so **F**'s equilibrium payoff is higher than $q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0^{\mathbf{S}}\right)$. On the other hand, in the efficient pure strategy equilibrium, **F**'s payoff is $q_0 v^{\mathbf{F}} P^{\mathbf{S}}\left(q_0\right)$ (Proposition 4.2). As $q_0^{\mathbf{S}} > q_0$ in a mixed strategy equilibrium, it follows that, if the mixed strategy equilibrium exists, then it provides **F** with a higher expected payoff than she obtains in any other equilibrium.

From **S**'s perspective, **F**'s behavior in the equilibrium characterized in Proposition 5.1 induces a distribution $H_{G^{\mathbf{F}}}$ of beliefs with which he may receive the project that has a support of $\left\{q_t^{\mathbf{S}}\right\}_{t\in[0,\omega]}$. The expected belief at which he receives the project must equal the prior belief by Bayes' law ($E(q_t^{\mathbf{S}}) = q_0$). Moreover, since there is no discounting in our model, the time at which **S** receives the project has no direct impact on his payoff (besides the information this time conveys). Therefore, **S**'s expected payoff in the mixed strategy equilibrium is equal to his payoff as a DM with prior q_0 who receives a signal about the quality of the project that induces a posterior belief distribution $H_{G^{\mathbf{F}}}$. Since **S**'s optimal strategy for every belief $q_t^{\mathbf{S}}$ is different from his unique optimal strategy for belief q_0, this signal generates (strictly) valuable information. In the efficient pure strategy equilibrium characterized in Proposition 4.2, **S**'s payoff is equal to his payoff as a DM with a belief q_0. Hence, we can conclude that, if the mixed strategy equilibrium exists, then **S** obtains a higher expected payoff in that equilibrium than in any pure strategy equilibrium.

The following proposition, whose proof is omitted, formalizes the above arguments.

Proposition 5.5. *Let $q_0 > \overline{q}^{\boldsymbol{F}}$. If the equilibrium characterized in Proposition 5.1 exists, then it is the unique Pareto-efficient equilibrium.*

The cost of strategic uncertainty. While strategic uncertainty can mitigate the moral hazard problem, it does not lead to first best outcomes. In particular, it leads to two types of inefficiency. The first inefficiency is the natural solution to the moral hazard problem in our model: excessive scrutiny of (good) projects. When an authentic breakthrough occurs while **F** is learning, **S** does not necessarily launch the project immediately. Instead, he may examine the project for a while and approve it only if a second breakthrough occurs while he is learning. This can lead to the termination of a project that **F** knows to be good. The second inefficiency is less intuitive: a project may be launched after its (costly) examination by **F** has not uncovered positive news: if **F** fabricates a breakthrough in the partial-trust phase, **S** may respond by launching the project without scrutiny.

These two inefficiencies are related to two possible errors in collective decision making. Let $\Psi_1(t)$ denote the probability that a *good project is terminated*, conditional on it being submitted at time $t < \tau^*$, and let $\Psi_2(t)$ denote the probability that a *bad project is launched*, conditional on it being submitted at time $t < \tau^*$. In the verification phase, **S** launches the project only after observing a breakthrough while he is learning, and so $\Psi_1(t) = 1 - P^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right)$ and $\Psi_2(t) = 0$. In the partial-trust phase, **S** launches the project immediately with probability σ_t and examines it according to $P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right)$ otherwise, and so $\Psi_1(t) = (1-\sigma_t)(1-P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right))$ and $\Psi_2(t) = \sigma_t$. As both $P^{\mathbf{S}}\left(q_t^{\mathbf{S}}\right)$ and σ_t are increasing in t, and the partial-trust phase comes after the verification phase, it follows that the probability of making the first type of error goes down over time whereas the probability of making the second type of error goes up over time.

5.4. **Discussion of Assumptions.**

5.4.1. *Limitations on Fabricating Breakthroughs.* In practice, fabrication can be costly due to, for example, reputational or moral concerns. Throughout the analysis we abstracted away from these considerations in order to ease the exposition. We now explain why such considerations have no qualitative effect on our main results.

Costly Fabrication— To capture the idea that fabrication is costly, assume that **F** must incur a fixed cost of $\kappa > 0$ to fabricate a breakthrough. In this case, if **S** launches the project immediately upon receiving it, it is possible to show that there exists $\overline{q}^{\mathbf{F}}(\kappa) > \overline{q}^{\mathbf{F}}$ such that for all $q_t \leq \overline{q}^{\mathbf{F}}(\kappa)$ **F** prefers not to fabricate breakthroughs. Thus, in equilibrium, **F** can be trusted to take over the project once her beliefs reach $\overline{q}^{\mathbf{F}}(\kappa)$. That is, a positive cost of fabrication enables **F** to take over the project at a higher belief than $\overline{q}^{\mathbf{F}}$.

In the mixing region, the fixed cost of fabrication linearly affects **F**'s continuation value from faking a breakthrough, and so the dynamics of her indifference condition are given by

$$\frac{dW_t^{NB}}{dt} = \lambda^{\mathbf{F}} q_t \left[W_t^{NB} - \kappa - W_t^B \right] + c^{\mathbf{F}}.$$

That is, the indifference condition remains unaltered if we define $\hat{W}_t^{NB}(\kappa) = W_t^{NB} - \kappa$. Moreover, as $\lim_{\kappa \to 0} \overline{q}^{\mathbf{F}}(\kappa) = \overline{q}^{\mathbf{F}}$ and $\lim_{\kappa \to 0} \hat{W}_t^{NB}(\kappa) = W_t^{NB}$, equilibrium behavior is robust to allowing for $\kappa > 0$ as long as κ is not too large.

Adding a positive cost of fabrication increases the players' equilibrium payoff. First, breakthroughs that occur at any $q_t \in (\overline{q}^{\mathbf{F}}, \overline{q}^{\mathbf{F}}(\kappa))$ lead to the project being launched when $\kappa > 0$, which need not be the case when $\kappa = 0$. Second, the cost of continuing learning is effectively reduced with respect to faking a breakthrough as an authentic breakthrough saves the cost of fabrication. Intuitively, the cost of fabrication acts as a partial commitment device for **F** to refrain from fabricating breakthroughs and partially mitigates the mutually detrimental moral hazard problem.

Inability to Fabricate— To capture the idea that moral concerns may prevent some players from fabricating breakthroughs, assume that there are two types of first mover: an "honest" type who is incapable of fabricating breakthroughs and a "dishonest" type for whom fabricating breakthroughs is costless. Moreover, assume that **F**'s type is her private information. In our equilibrium, with probability $1 - G^{\mathbf{F}}(\tau^*)$, **F**'s strategy does not call on her to fabricate a breakthrough at any point in time. Hence, as long as the fraction of honest types is less than $1 - G^{\mathbf{F}}(\tau^*)$ the equilibrium we defined exists. Moreover, since the dishonest types must still be indifferent between fabricating a breakthrough and continuing to learn before τ^*, the evolution of $q_t^{\mathbf{S}}$ in $[0, \tau^*)$ is given by Proposition 5.1. Hence, this equilibrium remains the unique Pareto-efficient equilibrium.[14]

5.4.2. *Exogenous Detection of Fraud.* In various applications, **S** may immediately detect that a breakthrough was fabricated with some probability.[15] However, if the probability of detection is small, then this does not affect our main results (there is no moral hazard if the detection probability is high). The only difference is that, as in the case where fabrication is costly, **F** can be trusted to learn while her belief is (slightly) above $\overline{q}^{\mathbf{F}}$. Intuitively, if **F** is caught fabricating a breakthrough, then **S** may terminate the project. Hence, if q_t is close enough to $\overline{q}^{\mathbf{F}}$, **F** strictly prefers learning honestly until her belief reaches $\underline{q}^{\mathbf{F}}$ to fabricating a breakthrough that may lead to the project's termination. It follows that in the efficient mixed equilibrium, **F** will take over the project before τ^*.

5.4.3. *Money-burning.* Suppose that **F** can "burn money," that is, she can keep learning after a breakthrough occurs to signal its authenticity. Consider the intuitive (perfect Bayesian) equilibrium in which **S**'s beliefs are such that he terminates the project if he receives it before τ^*, and launches it if he receives it after τ^*. Taking this behavior into account, **F** never submits the project prior to τ^*. Effectively, in this equilibrium **F** invests $c\tau^*$ in gaining credibility, which leads **S** to launch the project at τ^* if a breakthrough occurs before τ^* (and immediately upon the occurrence of one after τ^*).

Note that to sustain the players' equilibrium behavior **S** must believe that a breakthrough he observed before τ^* is fabricated with a high probability. However, **F**'s equilibrium strategy is such that she does not fabricate breakthroughs, while authentic breakthroughs may occur at any point in

[14]Since $G^{\mathbf{F}}(\tau^*)$ converges monotonically to zero as $q_0 \downarrow \overline{q}^{\mathbf{F}}$ (see proof of Lemma A.4), for any fraction of honest types (that is strictly less than one), our equilibrium exists and is the unique efficient equilibrium if q_0 is near enough to $\overline{q}^{\mathbf{F}}$.
[15]In Section 6.2 we analyze endogenous detection of fraud.

time. It follows that such off-path beliefs do not satisfy the logic of standard refinements such as sequential equilibrium.[16]

5.4.4. *Small Conflict.* We now consider the implications of the large conflict assumption on the characterization, existence, and efficiency results of this section. The characterization result of Proposition 5.1 and the existence result of Proposition 5.3 do not depend on the large conflict assumption. Though the proof of Proposition 5.1 becomes more involved, both results hold under the assumption of *small conflict*: $\underline{q}^{\mathbf{F}} < \underline{q}^{\mathbf{S}} < \overline{q}^{\mathbf{F}} < \overline{q}^{\mathbf{S}}$. The main difference between the cases of small and large conflict is in the comparison between the mixed strategy equilibrium and the efficient pure strategy equilibrium. In particular, when the conflict between the players is small and the prior is in $(\overline{q}^{\mathbf{F}}, \overline{q}^{\mathbf{S}})$, there exist pure strategy equilibria in which both players learn that do not exist when the conflict between the players is large. Thus, the comparison in Section 5.3 may not suffice to establish the efficiency result of Proposition 5.5.

The additional pure strategy equilibria under the small conflict assumption share a similar structure: **F** learns until some time $\tau^{\mathbf{F}}$, and then submits the project to **S** who then learns until the public belief drifts down to $\underline{q}^{\mathbf{S}}$. Among these equilibria **F** prefers the one in which $\tau^{\mathbf{F}} = 0$, as she does not incur any learning cost in this equilibrium. This equilibrium is the unique efficient pure strategy equilibrium when the conflict is large. Hence, the comparison made in Section 5.3 establishes that the mixed strategy equilibrium is **F**'s preferred equilibrium under the small conflict assumption, and so it is Pareto-efficient regardless of the size of the conflict. Unlike **F**, **S** prefers the pure strategy equilibrium with the maximal $\tau^{\mathbf{F}}$. In this equilibrium, **F**'s learning provides an informative signal to **S** in a similar manner to the way her learning does in the mixed strategy equilibrium. In order to determine which equilibrium is **S**'s preferred one, we need to compare his value from the two signals. In Online Appendix B, we show that there are parameters such that the mixed strategy equilibrium is the unique Pareto-efficient equilibrium. In particular, in line with the intuition for large conflict, if the maximal $\tau^{\mathbf{F}}$ is small, the equilibrium with strategic uncertainty is also preferred by **S**, which makes it the unique Pareto-efficient equilibrium.

[16]We provide this heuristic argument in lieu of a formal argument since we are not aware of refinements in continuous-time games that capture this idea.

5.4.5. *Private Learning.* In our model, **F**'s belief drifts down over time conditional on no breakthrough occurring. In particular, this feature implies that there exists an exogenous point in time (which we refer to as τ^*) at which **F** can be trusted to take over the project. Our analysis hinges on the fact that this point in time is commonly known. By contrast, if **F** can privately stop learning about the project (and in so doing avoid the cost of learning), then her belief need not drift down over time. Hence, a point in time where **S** can be trusted to take over the project need not even exist.

In fact, under private learning, our sequential learning model resembles a "cheap talk" game in which **F** can induce any equilibrium belief at zero cost. The reason for this is that **F** can hold on to the project without learning and submit it at any time she wishes at no cost. Under our large conflict assumption, this would lead to a collaboration failure when $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{\mathbf{S}})$, similar to the one that occurs in the pure strategy equilibrium where the project is terminated at $t = 0$.

6. Applications

6.1. **Quality Assurance.** One of the last stages in product development is quality assurance, namely, the process of making sure that there are no faults that prevent a company from bringing its product to market. In this section, we capture this idea by modifying **S**'s learning technology. In particular, we assume that **S** looks for breakdowns instead of for breakthroughs.[17] We show that, as in the baseline model, for intermediate priors, initial strategic uncertainty is necessary to sustain efficient collaboration.[18]

Assume that **S** has access to a learning technology that generates a signal according to an exponential distribution with intensity $\gamma > 0$ only in the *bad* state and entails a (sufficiently low) flow cost of c. The problem of a decision maker who looks for negative evidence is a mirror image of the problem of a DM who looks for positive evidence: instead of the DM's belief

[17]See Keller and Rady (2015) for a comparison of good news and bad news models in the experimentation literature and Bonatti and Hörner (2017) for an analysis of free-riding when the learning technology delivers negative evidence.

[18]Bonatti and Hörner (2017) also find dispersion of beliefs and mixed strategy equilibria in the case of hidden actions. Their setting differs from ours: learning is simultaneous, there is independence of payoffs across players, and the strategic interaction is due to informational externalities. Importantly, in their setting, strategic uncertainty is a reflection of coordination failure, whereas in our setting it is a coordination device.

drifting down while she examines the product, it drifts up, and, instead of the DM launching the product after receiving a signal, she terminates it. The DM's behavior is characterized by two cutoff beliefs, $\underline{q}^{bd} < \overline{q}^{bd}$, such that she finds it optimal to learn only if $q_t \in [\underline{q}^{bd}, \overline{q}^{bd})$, in which case she learns until her belief drifts up to (at most) $\overline{q}^{bd}$ and then launches the project.[19] Note that at $\underline{q}^{bd}$ the DM is indifferent between terminating the project and examining it.

Consider now the sequential interaction between **F** and **S**. We maintain the large conflict between the players by assuming that there exist priors for which **F** prefers launching the project to examining it, while **S** prefers terminating the project to examining it: $\overline{q}^{\mathbf{F}} < \underline{q}^{bd}$. Moreover, to highlight the moral hazard problem, we focus on priors $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{bd})$; priors for which the project is terminated in every pure strategy equilibrium.

In contrast to the baseline model, when **S** is endowed with a breakdown learning technology there cannot be a verification phase. To see this, note first that if **S** learns with probability 1 after receiving the project, then **F**'s payoff from submitting the project to **S** is greater than her payoff from launching the project. This is because if the project is good, **S** will not observe a breakdown and will launch the project, whereas if the project is bad, with some probability **S** will observe a breakdown and terminate the project. However, by arguments similar to the ones used in the baseline model (Lemma 4.3), **F**'s equilibrium payoff cannot be greater than her payoff from learning until ω and having **S** best respond to this strategy. If $q_t > \overline{q}^{\mathbf{F}}$, the latter payoff is less than **F**'s payoff from launching the project. It follows that **S** must mix between terminating and launching the project.

To keep **F** indifferent between learning and fabricating a breakthrough at every $t < \tau^*$ (the definition of τ^* is unchanged), **S** must mix between terminating the project and examining it, should he receive it at t. Since **S** is indifferent between these two actions only if $q_t^{\mathbf{S}} = \underline{q}^{bd}$, his beliefs must be $\underline{q}^{bd}$ at all $t < \tau^*$. As **S** randomizes between examining and terminating the project, we refer to this phase as the *partial-mistrust phase*.

Let η_t denote the probability that **S** examines the project if he receives it at time $t < \tau^*$, and $1 - \eta_t$ denote the probability with which he terminates

[19] It can be shown that, if the cost of learning is reasonable, $l(\overline{q}^{bd}) = \frac{\gamma - c}{c}$, and that $l(\underline{q}^{bd}) = \frac{\gamma - c}{c} e^{-\gamma\hat{\tau}}$ for $\hat{\tau}$ given implicitly by $l(\overline{q}^{bd})(v - c\hat{\tau}) = 1 + c\hat{\tau} + e^{\gamma\hat{\tau}} \left[\int_0^{\hat{\tau}} \gamma e^{-\gamma s}\, c\, s\, ds\right]$.

the project. The probability that a bad project is launched if **S** examines it is $\frac{l(\underline{q}^{bd})}{l(\overline{q}^{bd})}$, whereas the probability that a good project is launched if **S** examines it is one. Thus, **F**'s continuation values from submitting the project to **S** with and without an authentic breakthrough are

$$(2b)\qquad W_t^B = \eta_t v^{\mathbf{F}}$$

$$W_t^{NB} = q_t W_t^B - \eta_t(1-q_t)\frac{l(\underline{q}^{bd})}{l(\overline{q}^{bd})}.$$

F's indifference between continued learning and submitting the project with a fabricated breakthrough at any time $t < \tau^*$ implies that (3) is now a differential equation that determines η_t, namely,

$$c^{\mathbf{F}} = \frac{d\eta_t}{dt}\left(q_t v^{\mathbf{F}} - (1-q_t)\frac{l(\underline{q}^{bd})}{l(\overline{q}^{bd})}\right).$$

This equation, together with a boundary condition for η_{τ^*} that makes **F** indifferent between taking over the project and submitting it to **S** at τ^*, characterizes the efficient equilibrium with strategic uncertainty.[20]

Proposition 6.1. *Assume that **S** searches for breakdowns and that $q_0 \in (\overline{q}^{\boldsymbol{F}}, \underline{q}^{bd})$. In an efficient equilibrium with $\omega > 0$: 1) **F** mixes at all $t < \tau^*$ so that $q_t^{\boldsymbol{S}} = \underline{q}^{bd}$, 2) **F** takes over the project at τ^*, and 3) if **S** receives the project at $t < \tau^*$ he examines it with probability η_t and terminates it with probability $1-\eta_t$, where*

$$(7)\qquad \eta_t = \eta_{\tau^*} - \int_t^{\tau^*} \frac{c^{\boldsymbol{F}}}{q_s v^{\boldsymbol{F}} - (1-q_s)\frac{l(\underline{q}^{bd})}{l(\overline{q}^{bd})}} ds,$$

with the boundary condition

$$\frac{\eta_{\tau^*}}{1-\eta_{\tau^*}} = l(\overline{q}^{\boldsymbol{F}}) \times \frac{v^{\boldsymbol{F}} - \frac{1}{l(\overline{q}^{\boldsymbol{F}})}}{1 - \frac{l(\underline{q}^{bd})}{l(\overline{q}^{bd})}}.$$

Since $\eta_{\tau^*} < 1$ and η_t is increasing, this mixed strategy equilibrium exists if $\eta_0 \geq 0$. Moreover, if this equilibrium exists, then it is the unique Pareto-efficient equilibrium since, for $q_0 \in (\overline{q}^{\mathbf{F}}, \underline{q}^{bd})$, the project is terminated immediately in any equilibrium in pure strategies. Furthermore, as in the baseline model, **F** fabricates breakthroughs with a probability that decreases over time.

[20] The proof of this result is analogous to that of Proposition 5.1 and is omitted.

6.2. **Monitoring.** In various sequential interactions and, in particular, in standardization processes, the player who moves second in the process does not have the ability to learn directly about the project. Instead, that player can gather information about the first mover's findings. For instance, pharmaceutical companies run clinical trials and the FDA scrutinizes the data generated in the clinical trials.[21]

In this section, we capture this idea by endowing **S** with a monitoring technology. Specifically, upon receiving the project, **S** chooses the quality of a signal that is informative about whether the breakthrough was fabricated or not.[22] Unlike in the previous sections where **S** learns about the project's quality, in this section **S**'s learning does not provide direct value to **F**. Nevertheless, our main findings apply in this context as well.

Formally, we assume that, after receiving the project, **S** can choose the quality ϕ of a signal with realizations in the set $\{s_Y, s_N\}$. If the breakthrough was fabricated and the chosen quality is ϕ, then the signal realization is s_Y with probability ϕ and s_N with the complementary probability. On the other hand, if the breakthrough was authentic, then the signal realization is s_N with probability 1 (regardless of the choice of quality). Therefore, the realization s_Y is conclusive evidence that the breakthrough was fabricated, while the realization s_N is inconclusive evidence about the nature of the breakthrough. We assume that **S**'s cost of acquiring a signal of quality ϕ is $\kappa(\phi)$, where $\kappa : [0, 1) \to \mathbb{R}$ is strictly increasing, convex, and twice continuously differentiable. Furthermore, we assume that $\kappa(0) = \kappa'(0) = 0$ and $\lim_{\phi \to 1} \kappa(\phi) \to \infty$.

Denote by $\hat{q}^{\mathbf{S}} = \frac{1}{1+v^{\mathbf{S}}}$ the belief for which **S** is indifferent between launching the project and terminating it. If $q_\tau > \hat{q}^{\mathbf{S}}$, then information has no value for **S** as he would rather launch the project regardless of the signal's realization. Thus, for the rest of the section we assume that $q_0 < \hat{q}^{\mathbf{S}}$. If $q_0 \leq \overline{q}^{\mathbf{F}}$, then, in the unique Pareto-efficient equilibrium, **F** takes over the project at $t = 0$ as in previous sections. In the remainder of this section, we

[21]Surprisingly, the role of the FDA as a monitor has not received much attention in the learning literature. Typically, papers studying related issues in this strand of the literature take a contracting approach that assumes away fraud and data manipulation on the equilibrium path (despite both these phenomena being well documented; see, e.g., George and Buyse, 2015 and Seife, 2015), phenomena that our model predicts.

[22]This application contributes to the literature on costly verification initiated by Townsend (1979). We introduce a general monitoring technology instead of the standard assumption that at a fixed cost the monitor learns the truth.

assume that $\overline{q}^{\mathbf{F}} < \hat{q}^{\mathbf{S}}$, which corresponds to the large conflict assumption, and focus on priors in $(\overline{q}^{\mathbf{F}}, \hat{q}^{\mathbf{S}})$. In this case, if **S** acquires information he terminates the project after s_Y and launches it after s_N.

S's choice regarding the quality of information depends on two factors. The first is the probability that the breakthrough he observed was fabricated, which we denote by z_t. The second is the probability that the project is good conditional on the breakthrough being fabricated, which, as before, we denote by q_t. Let $\phi^*(z_t, q_t)$ denote **S**'s optimal choice. Intuitively, $\phi^*(z_t, q_t)$ is increasing in z_t and decreasing in[23] q_t.

Furthermore, as the marginal cost of acquiring better information is increasing, there is an upper bound on the value of z_t for which **S** collects information. We denote by $\overline{z}(q_t)$ the upper bound for a given q_t. Moreover, we denote by $\overline{\phi}(q_t) = \phi^*(\overline{z}(q_t), q_t)$ the quality of information that **S** acquires given the pair $(\overline{z}(q_t), q_t)$. It can be shown that $\overline{z}(q_t)$ is increasing in q_t so that $\overline{\phi}(q_t)$ is decreasing in q_t

While $q_t > \overline{q}^{\mathbf{F}}$, **F** must randomize between continuing to learn and fabricating a breakthrough, and so her behavior is characterized by the indifference condition (3). In this case, the continuation values from submitting the project to **S** with and without an authentic breakthrough are, respectively,

$$\text{(2c)} \qquad \begin{aligned} W_t^B &= \theta_t v^{\mathbf{F}}, \\ W_t^{NB} &= \theta_t(1 - \phi^*(z_t, q_t))(q_t(v^{\mathbf{F}} + 1) - 1), \end{aligned}$$

where θ_t is the probability that **S** acquires information of quality $\phi^*(z_t, q_t)$ (with the complementary probability that he terminates the project).

When $\theta_t < 1$, **S** randomizes between terminating the project and monitoring, which resembles the *partial-mistrust* phase in the quality assurance application. However, there are two differences worth mentioning. First, under the breakdown technology analyzed in Section 6.1, **S**'s indifference occurs for a single belief ($\underline{q}^{bd}$). By contrast, when **S** monitors **F** this indifference occurs for all pairs (q_τ, z_τ) for which $z_\tau = \overline{z}(q_\tau)$. Moreover, the amount of monitoring that **F** is subject to is given by $\overline{\phi}(q_\tau)$, which is not constant over time. Second, the fact that submitting a good project after

[23]All formal proofs of the statements in this section appear in Online Appendix C.

an authentic breakthrough leads to a higher approval probability than submitting a good project after a fabricated breakthrough creates incentives to continue learning that do not exist in the breakdown technology case.

While this new incentive is strong when q_t is high (as this is when authentic breakthroughs are highly likely), it does not have a qualitative impact on the efficient equilibrium for intermediate priors. The reason for this is that the evolution of **F**'s behavior toward honesty requires a verification phase when as t approaches τ^*, according to the boundary condition

$$\lim_{t \to \tau^*} W_t^{NB} = \overline{q}^{\mathbf{F}} v^{\mathbf{F}} - (1 - \overline{q}^{\mathbf{F}}) \quad \text{or} \quad \lim_{t \to \tau^*} \theta_t (1 - \phi^*(0, q_{\tau^*})) = 1.$$

Hence, we can establish the following proposition (the proof is relegated to Online Appendix C).

Proposition 6.2. *There exists $q^{**} \in (\overline{q}^{\boldsymbol{F}}, \hat{q}^{\boldsymbol{S}})$ such that, for any $q_0 \in (\overline{q}^{\boldsymbol{F}}, q^{**})$, the unique Pareto-efficient equilibrium is characterized by $\tau^{**} < \tau^*$ such that*

(1) *for all $t \leq \tau^{**}$, if **S** receives the project at t, he terminates the project with probability $1 - \theta_t$ and otherwise monitors **F** according to $\phi_t = \overline{\phi}(q_t)$ (partial-mistrust phase),*
(2) *for all $t \in (\tau^{**}, \tau^*)$, if **S** receives the project at t, he monitors **F** according to $\phi_t = \phi_t^*(z_t, q_t) < \overline{\phi}(q_t)$ (verification phase),*

*and **F** takes over the project at τ^*.*

Note that in the partial-mistrust phase $z_t = \overline{z}(q_t)$, which implies that z_t decreases over time. In the verification phase, on the other hand, the evolution of $\phi^*(z_t, q_t)$ follows from (2c) and (3). As in this phase $\theta_t = 1$, it follows that $\phi^*(z_t, q_t)$ is decreasing, which, in turn, implies that z_t also decreases over time. Hence, **F**'s behavior resembles her behavior in the previous sections: although she fabricates evidence, she does so at a lower rate as time progresses.

6.3. **Expert Advice.** The objective of this paper is to understand interactions in which there is sequential learning about the quality of a project and both players have veto power. However, our insights also apply in situations in which the second mover has the authority to approve the project but he himself cannot learn. This is a limit case of our baseline model (as well as that of the other variants we consider) where $c^{\mathbf{S}}$, $\lambda^{\mathbf{S}}$, and $v^{\mathbf{S}}$ are such that $\frac{c^{\mathbf{S}}}{\lambda^{\mathbf{S}}} \geq \frac{v^{\mathbf{S}}}{1+v^{\mathbf{S}}}$. Thus, the results in this section are a special case of the baseline model and do not require independent proof.

S's behavior as a decision maker in this case is straightforward. If he receives the project and $q_t^{\mathbf{S}} > \hat{q}^{\mathbf{S}}$ he launches it, whereas if $q_t^{\mathbf{S}} < \hat{q}^{\mathbf{S}}$ he terminates it. Since **S** cannot learn, if $q_t \leq \overline{q}^{\mathbf{F}}$ then the unique Pareto-efficient equilibrium has **F** taking over the project. Moreover, if $q_0 \geq \hat{q}^{\mathbf{S}}$ and $q_0 > \overline{q}^{\mathbf{F}}$, then **F** prefers launching the project to examining it, and **S** is willing to launch the project at $t = 0$. Hence, we focus on priors $q_0 \in \left(\overline{q}^{\mathbf{F}}, \hat{q}^{\mathbf{S}}\right)$ and study the mixed strategy equilibria of the game. In this case, in equilibrium, **S** must be indifferent between launching the project and terminating it at all $t < \tau^*$ and so $q_t^{\mathbf{S}} = \hat{q}^{\mathbf{S}}$, and **F** must mix throughout when $q_t \in \left(\overline{q}^{\mathbf{F}}, \hat{q}^{\mathbf{S}}\right)$.

F's values from submitting the project while mixing reduce to

$$(2d) \qquad \begin{aligned} W_t^B &= \mu_t v^{\mathbf{F}} \\ W_t^{NB} &= q_t W_t^B - \mu_t(1-q_t), \end{aligned}$$

where μ_t is the probability that the project is launched and $1 - \mu_t$ is the probability that the project is terminated. The indifference condition (3) determines the evolution of μ_t, which is given by

$$c^{\mathbf{F}} = \frac{d\mu_t}{dt}(q_t v^{\mathbf{F}} - (1-q_t)).$$

Since the transition to the region where **F** learns honestly must be smooth, $\mu_{\tau^*} = 1$ and the next proposition follows directly.

Proposition 6.3. *Assume that **S** cannot learn and that $q_0 \in \left(\overline{q}^{\boldsymbol{F}}, \hat{q}^{\boldsymbol{S}}\right)$. If in an efficient equilibrium $\omega > 0$, then $q_t^{\boldsymbol{S}} = \hat{q}^{\boldsymbol{S}}$ for all $t < \tau^*$ and*

$$(8) \qquad \mu_t = 1 - \int_t^{\tau^*} \frac{c^{\boldsymbol{F}}}{q_s v^{\boldsymbol{F}} - (1-q_s)} ds.$$

Since μ_t is increasing, this equilibrium exists if $\mu_0 \geq 0$. It is the unique Pareto-efficient equilibrium since for such priors the project is terminated immediately in any equilibrium in pure strategies.

This application is related to Escobar and Zhang (2020) who characterize the optimal contract offered by a principal that can commit to a launching policy.[24] In contrast to our paper, Escobar and Zhang assume that the

[24] Henry and Ottaviani (2019) and McClellan (2020) study similar problems in which information is modeled as a Brownian motion.

agent's learning is private and costless, and that players discount the future. They find that the principal delays approval of the project after its submission, and that this delay is shorter the longer it takes for the agent to report a breakthrough. This decreasing delay serves as an incentive for the agent to continue learning about the project.

Our paper and Escobar and Zhang (2020) complement each other by taking a different approach to a similar moral hazard problem. In their model, the principal's commitment power allows him to screen agents by using delay as a money-burning device. In our model, the players cannot commit, and so they must incentivize each other to solve the moral hazard problem in a way that is self-enforcing. As a result, while in their paper the agent reports honestly throughout the interaction, as is standard in the mechanism design literature, in our paper **F** may initially fabricate breakthroughs, and there is a gradual transition from dishonesty to honesty.

7. Concluding Remarks

We study the moral hazard problem that arises in sequential learning when players have the ability to fabricate evidence but do not have commitment power. We assume that **F** publicly collects information using a technology that generates a signal only if the project is good, and endow the second mover with various learning technologies in line with the various applications considered in the paper: learning in organizations, quality assurance, monitoring, and expert advice.

We find that efficient sequential collaboration may require occasional deception and a gradual transition toward honesty. In particular, in such a collaboration, **F** fabricates breakthroughs with a probability that decreases over time and **S**'s relevant beliefs diverge from **F**'s beliefs. This incentivizes **S** to treat **F** more trustingly as time goes by. Eventually, **F** stops fabricating breakthroughs and submits the project only after an authentic breakthrough, and, as a result, **S** launches the project once he receives it.

While **F** mixes, the probability that **S** will launch the project conditional on receiving it increases over time. This raises **F**'s return on learning and obtaining an authentic breakthrough, which compensates her for the decrease in the likelihood of a breakthrough occurring. Moreover, the increase in the probability of launching both good and bad projects follows a similar pattern regardless of **S**'s particular learning technology. This leads to a common pattern in the evolution of mistakes across all applications: the probability that a good project is terminated decreases over time while the probability that a bad project is launched increases over time.

The fact that our model does not rely on explicit contractual arrangements and can accommodate a second mover with a qualitatively different learning technology than that of the first mover makes it amenable to a wide range of other contexts. For instance, consider a criminal justice system where the police and district attorney interact sequentially. Our model predicts that efficient sequential collaboration between the police and the district attorney will occasionally lead to false accusations being made, on the one hand, and suspects against whom the police have obtained evidence not being prosecuted, on the other hand.

References

Antler, Yair, Daniel Bird, and Santiago Oliveros. 2019. "Sequential learning." CEPR Discussion Paper 13934.

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.

Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier. 2020. "Testing, disclosure and approval." *Journal of Economic Theory*, 187: 105002.

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

Bonatti, Alessandro, and Johannes Hörner. 2017. "Learning to disagree in a game of experimentation." *Journal of Economic Theory*, 169: 234–269.

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. 2020. "Uncertainty-driven cooperation." *Theoretical Economics*, 15(3): 1023–1058.

Che, Yeon-Koo, and Johannes Hörner. 2018. "Recommender systems as mechanisms for social learning." *Quarterly Journal of Economics*, 133(2): 871–925.

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.*

Escobar, Juan F., and Qiaoxi Zhang. 2020. "Delegating learning." *Theoretical Economics*, Forthcoming.

George, Stephen L., and Marc Buyse. 2015. "Data fraud in clinical trials." *Clinical Investigation*, 5(2): 161–173.

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." *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." *RAND Journal of Economics*, 44(4): 632–663.

Keller, Godfrey, and Sven Rady. 2015. "Breakdowns." *Theoretical Economics*, 10(1): 175–202.

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. 2020. “Experimentation and approval mechanisms.” *University of Chicago, 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.” *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.

Seife, Charles. 2015. “Research misconduct identified by the US Food and Drug Administration: Out of sight, out of mind, out of the peer-reviewed literature.” *JAMA Internal Medicine*, 175(4): 567–577.

Townsend, Robert M. 1979. “Optimal contracts and competitive markets with costly state verification.” *Journal of Economic Theory*, 21(2): 265–293.

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

A. Proofs

Proof of Proposition 4.1

A DM’s value from learning until time t when her prior is q_0 is

$$EU(t, q_0) = q_0 \int_0^t \lambda e^{-\lambda s} (v - cs)\, ds - \left(q_0 e^{-\lambda t} + (1 - q_0)\right) ct.$$

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:

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

Since $l(q_t) = l(q_0)e^{-\lambda t}$ we obtain the cutoff belief $q_t = \underline{q}(v, \frac{c}{\lambda}) = \frac{c}{\lambda v}$. 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 $\frac{EU^*(q)}{q} - \left(v - \frac{1}{l(q)}\right) \geq 0$, where $EU^*(q)$ is the value of optimal learning with prior q, i.e.,

$$\mathbf{M}(q) \equiv \frac{\lambda}{c} - 1 + \log\left(l(\underline{q}(v, \frac{c}{\lambda}))\right) - (l(q) + \log(l(q))) \geq 0.$$

Note first that $\mathbf{M}(q)$ is continuous and decreasing in q. Second, it is positive at $\underline{q}(v, \frac{c}{\lambda})$ if $v < \frac{1}{l(\underline{q}(v,\frac{c}{\lambda}))}$, which holds since $c < \lambda \frac{v}{v+1}$. Third, $\lim_{q \to 1} \mathbf{M}(q) = -\infty$. Thus, by the mean value theorem, there is a unique cutoff $\overline{q}(v, \frac{c}{\lambda}) \in \left(\underline{q}(v, \frac{c}{\lambda}), 1\right)$ for which $\mathbf{M}(\overline{q}(v, \frac{c}{\lambda})) = 0$. Since launching the project is better than terminating the project at $\overline{q}(v, \frac{c}{\lambda})$, it follows that $\underline{q}(v, \frac{c}{\lambda}) < \overline{q}(v, \frac{c}{\lambda})$. □

Proof of Proposition 4.2

In a pure strategy equilibrium, if **S** receives the project at $t < \tau^{\mathbf{F}}$, then he infers that an authentic breakthrough has occurred and updates his belief to $q_t^{\mathbf{S}} = 1$. However, if **S** receives the project at $t = \tau^{\mathbf{F}}$, then he infers that the breakthrough at t was fabricated, and so $q_t^{\mathbf{S}} = q_t$.

If $q_0 \leq \overline{q}^{\mathbf{F}}$, **F** taking over the project is an equilibrium as she would rather examine the project than launch it. 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. 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 she would rather fabricate a breakthrough at $t = 0$ and have the project approved. Consider an equilibrium in which both players learn. There are two possible events: either a breakthrough occurs before the player's (common) belief reaches $\underline{q}^{\mathbf{S}}$ or not. In the first case, **F** would rather deviate and fabricate a breakthrough at $t = 0$ (and have **S** respond by launching it). In the second case, the project is terminated when the belief is $\underline{q}^{\mathbf{S}} > \overline{q}^{\mathbf{F}}$, and it follows that **F** would also rather deviate by fabricating a breakthrough 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 ω, then $q_\omega = q_\omega^{\boldsymbol{S}}$; otherwise, there exists a sequence $t_n \to \omega$ such that* $\lim_{n \to \infty}(q_{t_n}^{\boldsymbol{S}} - q_{t_n}) = 0$.

Proof of Lemma A.1. First, we show that $\omega < \infty$. If $\omega = \infty$, then for every stopping time τ 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 τ from the stopping time τ' is bounded from above by $q_0 v^{\mathbf{F}} - (1-q_0)c^{\mathbf{F}}\tau$, which is negative for sufficiently large τ.

By Bayes' law, if t is an atom of $G^{\mathbf{F}}$ then $q_t^{\mathbf{S}} = q_t$. Hence, if either ω 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

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

denote the hazard ratio of $G(\cdot)$ on (τ, ω), and note that by Bayes' law $q_t^{\mathbf{S}} = q_t \frac{\lambda^{\mathbf{F}} + h(t)}{\lambda^{\mathbf{F}} q_t + h(t)}$ in this interval. If there exists a sequence $t_n \to \omega$ such that $\lim_{n\to\infty} g^{\mathbf{F}}(t_n) > 0$, then for that sequence $\lim_{n\to\infty} h(t_n) = \infty$ and $\lim_{n\to\infty}(q_{t_n}^{\mathbf{S}} - q_{t_n}) = 0$. Otherwise, then as $g^{\mathbf{F}}(\omega) = 0$ and $g^{\mathbf{F}}$ is left-continuous at ω, there exists another sequence $t_n \to \omega$ such that for every t_n, we have that $g^{\mathbf{F}}(s) < g^{\mathbf{F}}(t_n)$ for all $s > t_n$. For t_n in this sequence,

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

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

First, consider the case where **F** uses the stopping time ω. If **F** submits the project at ω, by Bayes' law we have that $q_\omega = q_\omega^{\mathbf{S}}$ and hence **S**'s response at ω in **E** is also a best response to the profile in which **F** learns honestly until ω (we denote this profile by **E'**). If **F** terminates the project at ω she gets a continuation utility of 0, which is less than her utility from having **S** learn with belief q_ω. Hence, **F**'s continuation utility at ω in **E** is weakly lower than it is in **E'**. In **E'** or when **F** uses the stopping time ω in **E**, **S** receives the project before ω only after an authentic breakthrough. However, in the former case he 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 an authentic breakthrough occurs while **F** is learning or she submits the project at ω, then **S** weakly prefers his payoff in **E'** to that in **E**. If

F fabricates a breakthrough at $t < \omega$ (and submits the project), then since $q_0 < \overline{q}^{\mathbf{S}}$ **S** strictly prefers that **F** continue learning until ω instead of submitting the project. Note that this is exactly what occurs in **E'**.

If **F** does not use the stopping time ω, 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 along t_n, and so by the same arguments used above it follows that **E'** Pareto dominates[25] **E**. □

Proof of Corollary 4.4

For the first part of the corollary, consider a mixed strategy equilibrium and note that, since $q_0 \leq \underline{q}^{\mathbf{S}}$, Lemma A.1 implies that the project is terminated at ω. It follows that $q_\omega \geq \underline{q}^{\mathbf{F}}$. By Lemma 4.3, that mixed strategy equilibrium is Pareto dominated by the strategy profile where $\tau^{\mathbf{F}} = \omega$ and **S** best responds. Since $q_0 \leq \overline{q}^{\mathbf{F}}$, (i) the latter profile is dominated by **F** taking the project, and (ii) **F** taking over the project is an equilibrium.

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 her 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 ω 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 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 a positive probability upon receiving it, it must be the case that $\Delta(q_t) > 0$. Otherwise, if $\Delta(q_t) < 0$, then submitting the project gives **F** a higher utility than her utility from launching it.

[25]Since ω is finite the strategy profile **E'** is well defined in this case.

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 show that $q_t^{\mathbf{S}}$ is left-continuous at τ; the proof that it is right-continuous is analogous. Since $q_t^{\mathbf{S}}$ is bounded there exists a sequence $t_n \to \tau$ such that the subsequence $q_{t_n}^{\mathbf{S}}$ 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 $1 - \sigma_{t_{n_k}} \geq 0$,

$$(9)\qquad \lim_{k\to\infty} (1 - \sigma_{t_{n_k}}) \left(q_{t_{n_k}} v^{\mathbf{F}} - (1 - q_{t_{n_k}})\right) + \sigma_{t_{n_k}} q_{t_{n_k}} P^{\mathbf{S}}(\overline{q}^{\mathbf{S}}) > q_\tau P^{\mathbf{S}}(q_\tau^{\mathbf{S}}) \\ \Rightarrow \lim_{t_{n_k}\uparrow\tau} W_{t_{n_k}}^{NB} > W_\tau^{NB},$$

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

Now consider the case of a convergent sequence for which $\lim_{n\to\infty} q_{t_n}^{\mathbf{S}} < q_\tau^{\mathbf{S}}$ (upward jump). Assume without loss of generality that $q_{t_n}^{\mathbf{S}} < \overline{q}^{\mathbf{S}}$ for all n so 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 τ, 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 τ.

We now show that $q_t^{\mathbf{S}}$ is continuous at $t = 0$. Since $G^{\mathbf{F}}$ 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 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. As **F** mixes until τ^* (Corollary 4.4), we can use **F**'s continuation utility at τ^* 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 τ^* (by Lemma 4.3). It remains to show that increasing continuation utilities at τ^* 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* τ^* *characterizes an equilibrium if and only if* $q_0 < q_0^{\mathbf{S}}$ *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 τ^*, denote by $q_t^{\mathbf{S}}(U)$ the belief associated with this solution, and by $\sigma_t(U)$ the probability that **S** learns in this equilibrium.

Because equation (3) assumes that **S** best responds to his belief, its solution 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, by Bayes' law,

$$\frac{g^{\mathbf{F}}(t)}{1 - G^{\mathbf{F}}(t)} = \lambda^{\mathbf{F}} \frac{l(q_t)}{l(q_t^{\mathbf{S}}) - l(q_t)} \tag{10}$$

holds. If $q_t^{\mathbf{S}} \leq q_t$, then (10) 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 (10). To see this, note that by integrating (10) we have that

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

and substituting into (10) we have that **F**'s strategy is

$$g^{\mathbf{F}}(s) = \lambda^{\mathbf{F}} \frac{l(q_s)}{l(q_s^{\mathbf{S}}) - l(q_s)} \left(1 - G^{\mathbf{F}}(\tau^*)\right) e^{-\lambda^{\mathbf{F}} \int_s^{\tau^*} \left(\frac{l(q_u)}{l(q_u^{\mathbf{S}}) - l(q_u)}\right) du}. \tag{11}$$

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 τ^* 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 $\sigma_0 \leq 1$. $\square$

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 τ^*; 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 τ^* 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, τ^* 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, 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 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 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. **S**'s expected utility is

$$V_0^{\mathbf{S}} = \int_0^{\tau^*} \left(q_0\left(1 - G^{\mathbf{F}}(s)\right)\lambda^{\mathbf{F}} e^{-\lambda^{\mathbf{F}} s} + \left(q_0 e^{-\lambda^{\mathbf{F}} s} + (1 - q_0)\right) g^{\mathbf{F}}(s)\right)$$

$$\left[q_s^{\mathbf{S}}\left(v^{\mathbf{S}} - \frac{c^{\mathbf{S}}}{\lambda^{\mathbf{S}}}\right) P^{\mathbf{S}}\left(q_s^{\mathbf{S}}\right) - (1 - q_s^{\mathbf{S}}) c^{\mathbf{S}}\left(\tau^* - s\right)\right] ds \tag{12}$$

$$+ \left(1 - G^{\mathbf{F}}(\tau^*)\right)\left(q_0 e^{-\lambda^{\mathbf{F}} \tau^*} + (1 - q_0)\right) V_{\tau^*}^{\mathbf{S}},$$

where $V^{\mathbf{S}}_{\tau^*}$ is $\mathbf{S}$'s continuation utility at τ^*. Clearly, increasing $V^{\mathbf{S}}_{\tau^*}$ increases $V^{\mathbf{S}}_0$. Using (11), we can write the value function (12) as

$$\text{(12}b\text{)}\quad \frac{V^{\mathbf{S}}_0}{(1-q_0)\left(1-G^{\mathbf{F}}(\tau^*)\right)} = \frac{V^{\mathbf{S}}_{\tau^*}}{1-q_{\tau^*}} + \int_0^{\tau^*} \frac{d\left(e^{-\lambda^{\mathbf{F}}\int_s^{\tau^*}\left(\frac{l(q_u)}{l(q^{\mathbf{S}}_u)-l(q_u)}\right)du}\right)}{ds} + \left[l(q^{\mathbf{S}}_s)\left(v^{\mathbf{S}}-\frac{c^{\mathbf{S}}}{\lambda^{\mathbf{S}}}\right)P^{\mathbf{S}}\left(q^{\mathbf{S}}_s\right) - c^{\mathbf{S}}\left(\tau^*-s\right)\right] ds.$$

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 5.3

Recall that we abstract away from the non-generic cases where $q_0 = \overline{q}^{\mathbf{S}}$ or $\Delta(\overline{q}^{\mathbf{F}}) = 0$. Moreover, as discussed above, when $q_0 > \overline{q}^{\mathbf{S}}$ there is no strategy for $\mathbf{F}$ that induces belief $q^{\mathbf{S}}_t \leq \overline{q}^{\mathbf{S}}$ by Bayes' law. Hence, in this proof we assume that $q_0 < \overline{q}^{\mathbf{S}}$. First, we characterize τ^{**} 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 $\Delta(\overline{q}^{\mathbf{F}}) > 0$. It is easier to solve equation (3c) in terms of strategies instead of value functions:

$$\frac{c^{\mathbf{F}}}{q_t} = -\dot{\sigma}_t \Delta(q_t).$$

The solution of this differential equation is

$$\begin{aligned}\sigma_t &= \int_t^{\tau^*} \frac{c^{\mathbf{F}}}{q_s\,\Delta(q_s)} \times ds \\ &= \frac{c^{\mathbf{F}} + v^{\mathbf{F}}\left(1-P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right)\right)}{v^{\mathbf{F}}\left(1-P^{\mathbf{S}}\left(\overline{q}^{\mathbf{S}}\right)\right)} \times \log\left(\frac{l(q_{\tau^*})}{l(q_t)} \times \frac{\Delta(q_{\tau^*})}{\Delta(q_t)}\right) - c^{\mathbf{F}}\left(\tau^*-t\right),\end{aligned}$$

where we use the boundary condition $\sigma_{\tau^*} = 0$. Integrating we get that in the partial-trust phase, σ_t is given by

$$\text{(13)}\qquad \sigma_t = \int_t^{\tau^*} \frac{c^{\mathbf{F}}}{q_s\,\Delta(q_s)} \times ds.$$

If $1 > \int_0^{\tau^*} \frac{c^{\mathbf{F}}}{q_s\,\Delta(q_s)} \times ds$, then $\sigma_t < 1$ for all $t > 0$, there is no verification phase, and $\tau^{**} = 0$. Otherwise, τ^{**} is given by setting $\sigma_t = 1$ in (13).

Recall that $q^{\mathbf{S}}_0 > q_0$ and $\sigma_0 \leq 1$ are necessary and sufficient conditions for the existence of the mixed strategy equilibrium (Lemma A.4). First, consider the case where $\Delta(\overline{q}^{\mathbf{F}}) > 0$. If $\tau^{**} = 0$, existence is trivial; if $\tau^{**} > 0$, then in the verification phase the differential equation (3b) becomes

$$W_t^B = W_{\tau^{**}}^B - c^{\mathbf{F}} \int_t^{\tau^{**}} \frac{1}{q_u} du = W_{\tau^{**}}^B - \frac{c^{\mathbf{F}}}{\lambda^{\mathbf{F}}} \left(\frac{1}{l(q_{\tau^{**}})} + \lambda^{\mathbf{F}}(\tau^{**} - t) \right).$$

Given that in this case $W_t^B = v^{\mathbf{F}} P^{\mathbf{S}}(q_t)$, we have that $q_0^{\mathbf{S}} > q_0$ if

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

Now consider the case where $\Delta(\overline{q}^{\mathbf{F}}) < 0$ and there is only a verification phase. Since $W_{\tau^*}^B = v^{\mathbf{F}} - \frac{1}{l(\overline{q}^{\mathbf{F}})}$, the differential equation (3b) becomes

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

and the condition for existence reduces to

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

□

B. Online Appendix: Small conflict

In this appendix we show that strategic uncertainty can be mutually beneficial in situations where the conflict between the players is not large. Specifically, we say that there is a *small conflict* between the players if $\underline{q}^{\mathbf{F}} < \underline{q}^{\mathbf{S}} < \overline{q}^{\mathbf{F}} < \overline{q}^{\mathbf{S}}$. Since the pure strategy equilibrium in which **F** takes over the project Pareto dominates any mixed strategy equilibrium for $q_0 \leq \overline{q}^{\mathbf{F}}$ (Lemma 4.3 applies regardless of the size of the conflict), we restrict attention to $q_0 \in (\overline{q}^{\mathbf{F}}, \overline{q}^{\mathbf{S}})$ and we also assume that there is a small conflict between the players throughout this appendix.[26]

The main results in this appendix correspond to the characterization, existence, and efficiency results of Section 5. The next proposition establishes that, if there is a mixed strategy equilibrium, the equilibrium characterized in Proposition 5.1 is the Pareto-efficient equilibrium in the class of equilibria in mixed strategies.

Proposition B.1. *If **E** is an efficient equilibrium in which **F** uses a mixed strategy, then it is the equilibrium characterized in Proposition 5.1.*

We relegate the proof of this key technical result to the end of the appendix. This result implies that the characterization and existence results in the main text remain unaltered under the small conflict assumption. However, the efficiency result in Section 5.3 is obtained by comparing the mixed strategy equilibrium characterized in Proposition 5.1 to the unique efficient pure strategy equilibrium, namely, the equilibrium in which **S** takes over the project at $t = 0$. Under the small conflict assumption, there may exist additional equilibria in pure strategies in which **F** learns, and so to obtain an efficiency result we must first characterize these additional equilibria.

For the rest of this appendix we use $\Sigma(\tau^{\mathbf{F}})$ to denote the pure strategy profile in which **F** uses the strategy $\tau^{\mathbf{F}}$ and **S** best responds to it. Note that a strategy profile with a higher index $\tau^{\mathbf{F}}$ is a strategy profile in which **F** learns more.

Lemma B.1. *If $\underline{q}^{\boldsymbol{S}} \geq \frac{1}{1+v^{\boldsymbol{F}}}$, then $\Sigma(0)$ in the unique pure strategy equilibrium. If $\underline{q}^{\boldsymbol{S}} < \frac{1}{1+v^{\boldsymbol{F}}}$, then the set of Pareto-efficient pure strategy equilibria*

[26] As established in the text, the mixed strategy equilibrium we characterized does not exist if $q_0 \geq \overline{q}^{\mathbf{s}}$.

is $\{\Sigma(\tau^{\boldsymbol{F}}) : \tau^{\boldsymbol{F}} \leq \hat{\tau}\}$*, where* $\hat{\tau} > 0$ *is defined implicitly by*

$$1 - l_{\underline{q}^{S}} v^{\boldsymbol{F}} \equiv \frac{c^{\boldsymbol{F}}}{\lambda^{\boldsymbol{F}}} l_{q_0} \left(1 - e^{-\lambda^{\boldsymbol{F}} \hat{\tau}}\right) + c^{\boldsymbol{F}} \hat{\tau}. \tag{14}$$

Proof. Since $q_0 < \overline{q}^{\mathbf{S}}$, in a pure strategy equilibrium **S** never launches the project without inspecting it if he receives it at $\tau^{\mathbf{F}}$. Hence, when considering the optimal pure strategy equilibrium, there is no loss of generality in assuming that **F** submits the project at $\tau^{\mathbf{F}}$ (rather than terminating it) and **S** best responds.

Consider the strategy profile $\Sigma(\tau^{\mathbf{F}})$. **F**'s expected payoff from this profile is

$$V^{\mathbf{F}}(\tau) = (1 - q_0) \left[\left(l_{q_0} - l_{\underline{q}^{\mathbf{S}}} \right) v^{\mathbf{F}} - \frac{c^{\mathbf{F}}}{\lambda^{\mathbf{F}}} \left(1 - e^{-\lambda^{\mathbf{F}} \tau}\right) - c^{\mathbf{F}} \tau \right].$$

If $\tau^{\mathbf{F}} > 0$, then **S** launches the project if he receives it at $t < \tau^{\mathbf{F}}$. Hence, any deviation to $\tau < \tau^{\mathbf{F}}$ by **F** will result in the project being launched with certainty. It follows that the most profitable deviation for **F** is to fabricate a breakthrough and submit the project at $t = 0$, which, in turn, implies that $\Sigma(\tau^{\mathbf{F}})$ is an equilibrium if and only if

$$V^{\mathbf{F}}(\tau^{\mathbf{F}}) \geq q_0 v^{\mathbf{F}} - (1 - q_0).$$

The previous inequality is equivalent to

$$1 - l_{\underline{q}^{\mathbf{S}}} v^{\mathbf{F}} \geq \frac{c^{\mathbf{F}}}{\lambda^{\mathbf{F}}} l_{q_0} \left(1 - e^{-\lambda \tau^{\mathbf{F}}}\right) + c^{\mathbf{F}} \tau^{\mathbf{F}}.$$

The right-hand side of this inequality is increasing in $\tau^{\mathbf{F}}$. As $\hat{\tau}$ is the value for which this condition holds with equality, it follows that if $\tau^{\mathbf{F}} < \hat{\tau}$, then $\Sigma(\tau^{\mathbf{F}})$ is an equilibrium. Moreover, note that $\Sigma(0)$ is an equilibrium in pure strategies as, for this strategy profile, **F** cannot manipulate **S** into launching the project.

To summarize, the strategy profiles $\Sigma(\tau^{\mathbf{F}})$ for $\tau^{\mathbf{F}} \in \{0\} \cup (0, \hat{\tau}]$ are equilibria of this game. Moreover, all these equilibria are Pareto-efficient as **F**'s (**S**'s) payoff is increasing (decreasing) in $\tau^{\mathbf{F}}$. To complete the lemma, note that $\hat{\tau} = 0$ if $\underline{q}^{\mathbf{S}} = \frac{1}{1+v^{\mathbf{F}}}$. □

The lemma shows that, if $\underline{q}^{\mathbf{S}} \geq \frac{1}{1+v^{\mathbf{F}}}$, then the unique Pareto-efficient equilibrium in pure strategies is for **S** to take over the project at $t = 0$. As the efficient equilibrium in pure strategies in a small conflict is identical to the efficient equilibrium in pure strategies in a large conflict, the same argument used to derive Proposition 5.5 yields the following result.

Proposition B.2. *Assume that* $\underline{q}^{\boldsymbol{S}} \geq \frac{1}{1+v^{\boldsymbol{F}}}$*. If the equilibrium characterized in Proposition 5.1 exists, then it is the unique Pareto-efficient equilibrium.*

If, on the other hand, $\underline{q}^{\mathbf{S}} < \frac{1}{1+v^{\mathbf{F}}}$, then under the small conflict assumption there are multiple Pareto-efficient equilibria in pure strategies. Note that **F**'s preferred equilibrium in pure strategies is the one in which **S** takes over the project at time zero. Hence, by the same argument used in Section 5.3, if the equilibrium characterized in Proposition 5.1 exists, then it is Pareto-efficient. By contrast, **S**'s preferred equilibrium in pure strategies is $\Sigma(\hat{\tau})$. Hence, comparing **S**'s payoff in the mixed strategy equilibrium to his payoff as a DM—the comparison used to derive Proposition 5.5—is not sufficient to establish that **S** is better off in the equilibrium with strategic uncertainty.

Nevertheless, the method of the proof used to establish Proposition 5.5 can still be used to derive an efficiency result. As explained in Section 5.3, **S**'s payoff in the equilibrium with strategic uncertainty is the payoff he obtains as a DM with prior q_0 who receives a signal about the state that is induced by **F**'s learning. In particular, this signal has the support $\{1\} \cup \{\underline{q}^{\mathbf{F}}\} \cup \{q_t^{\mathbf{S}}\}_{t\in[0,\tau^*]}$. On the other hand, in $\Sigma(\hat{\tau})$, **S**'s payoff is that of a DM with prior q_0 who receives a binary signal that induces belief 1 with probability $q_0(1-e^{-\lambda^{\mathbf{F}}\hat{\tau}})$ and belief $q_{\hat{\tau}}$ with the complementary probability. Whether **S** prefers the mixed strategy equilibrium to the pure strategy equilibrium depends on the value of these signals to **S**.

If $\hat{\tau}$ is small, then the value of the signal in the pure strategy equilibrium is also small. To see this, observe that with high probability **S**'s belief after receiving the signal remains near the prior. As **S**'s value function is bounded from above and is continuous in his belief, it follows that the value of this signal vanishes as $\hat{\tau}$ converges to zero. As **S**'s value for the signal he receives in the equilibrium with strategic uncertainty is strictly positive, it follows that if the latter equilibrium exists, then it is the unique Pareto-efficient equilibrium if $\hat{\tau}$ is small enough. However, if $\hat{\tau}$ is large, it may be the case that **S** prefers the signal he obtains in the pure strategy equilibrium to the signal he obtains in the mixed strategy equilibrium. The above discussion is summarized in the following proposition.

Proposition B.3. *If the equilibrium characterized in Proposition 5.1 exists, then: (i) it is Pareto-efficient, and (ii) there exists some $\epsilon > 0$ such that if $\hat{\tau} < \epsilon$, it is the unique Pareto-efficient equilibrium.*

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

Proof of Proposition B.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 B.2). We then show that, at some point in time, moral hazard no longer hinders the interaction (this result requires using the technical Lemma B.3). Lemma B.4 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 B.5 and B.6).

In this proof we often use the strategy profile where 1) **F** does not fabricate breakthroughs 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 B.2. *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 to a DM with a value of $v^{\mathbf{F}} + v^{\mathbf{S}}$ who has the choice of using the learning technology $(\lambda^{\mathbf{F}}, c^{\mathbf{F}})$ or $(\lambda^{\mathbf{S}}, c^{\mathbf{S}})$.

The value such a DM obtains from learning at any $q > \overline{q}^{\mathbf{S}}$ is strictly less than the DM's value from launching the project immediately. To see this recall that in Proposition 4.1 we defined the function

$$\mathbf{M}(q) \equiv \frac{\lambda}{c} - 1 + \log\big(l(\underline{q}(v,\lambda,c))\big) - (l(q) + \log(l(q))) \geq 0,$$

which is continuous and decreasing in q. As $\overline{q}(v,\lambda,c)$ is implicitly defined by $\mathbf{M}(\overline{q}(v,\lambda,c)) = 0$, implicit differentiation gives that $\frac{\partial \overline{q}(v,\lambda,c)}{\partial v} = -l(\overline{q}(v,\lambda,c))\frac{\overline{q}(v,\lambda,c)}{\left(v-\frac{c}{\lambda}\right)} < 0$, and so $\overline{q}(v,\lambda,c)$ is decreasing in v. 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 for which $q_t > \overline{q}^{\mathbf{S}}$. 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 time zero induces **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 her payoff from launching the project immediately. Thus, **E'** is an equilibrium that Pareto dominates **E**. □

By Lemma B.2, we focus on the case where $q_0 \leq \overline{q}^{\mathbf{S}}$ in the rest of this proof.

Lemma B.3. *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 ω 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 ω, in contradiction to **F**'s indifference condition. Next, consider the case where ω is an atom. If there exists a decreasing sequence $q_{t_n}^{\mathbf{S}}$ for $t_n \to \omega$, the same argument used above establishes the claim. Otherwise, since $\Delta(q_t) \geq 0$ whenever $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$, it follows that the value from submitting the project at ω where **S**'s belief is correct is strictly less than the value from submitting the project just prior to ω. □

Denote by ψ 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 B.3 this set is not empty). If $q_t^{\mathbf{S}} > \overline{q}^{\mathbf{S}}$ 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 from launching the project. By the definition of ψ, **F**'s continuation utility is less than that, and so it must be that $q_t^{\mathbf{S}} \leq \overline{q}^{\mathbf{S}}$ for all $t < \psi$. This implies that **F** fabricates breakthroughs with a positive intensity at any $t < \psi$. Moreover, if $q_t^{\mathbf{S}} = \overline{q}^{\mathbf{S}}$ 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 B.4. *$q_t^{\boldsymbol{S}}$ 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 τ^* replaced by ψ. □

Lemma B.5. *The project is terminated at ω in **E**.*

Proof. Assume by way of contradiction that **S** learns after receiving the project at ω. By Lemma 4.3, **F**'s continuation utility at ψ in **E** is weakly less than her continuation utility in **E'**. Moreover, **F**'s continuation payoff at ψ in **E'** is weakly less than her payoff from submitting the project to **S** and having **S** best respond to belief q_ψ. Hence, **F**'s continuation utility at ψ 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 B.4 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 ψ than at ψ. □

Note that if **F** would rather launch the project than terminate it at ω (and receive a payoff of zero), then submitting the project just prior to ω is a profitable deviation regardless of **S**'s response. Hence, from Lemma B.5 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 τ^* must be efficient. Thus, it suffices to show that **F** taking over the project at τ^* is the unique efficient equilibrium among all equilibria where the project is terminated at ω.

Lemma B.6. *Assume that $q_0 = \overline{q}^{\boldsymbol{F}}$ and that **F** must use a strategy in which she terminates the project at ω. **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 ω is weakly less than her utility from learning honestly until ω and then terminating the project. Similarly, it shows that **S**'s utility from any equilibrium in which

F terminates the project at ω is weakly less than the maximum between his utility from **F** learning honestly until ω and then terminating the project and his utility from launching the project immediately. $\square$

$\square$

C. Online Appendix: Monitoring

In this appendix, we conduct the formal analysis for Section 6.2 (Monitoring). The description of the model is provided in the main text.

C.1. **The Monitor's Problem.** We now discuss **S**'s optimal behavior when he receives the project. Consider the case where **S** acquires information of quality $\phi > 0$. In the event that the breakthrough was indeed fabricated, **S** launches the project if the signal's realization is s_N and terminates it if the signal's realization is s_Y. If the project is launched, which occurs with probability $1 - \phi$, **S**'s expected payoff is $q_t v^{\mathbf{S}} - (1 - q_t)$. If, on the other hand, the breakthrough was not fabricated, **S** always obtains the realization s_N and his payoff from launching it is $v^{\mathbf{S}}$ since the project is good. Thus, **S**'s ex-ante expected payoff from acquiring information of quality ϕ and following the realization of the signal is given by

$$EU^{\mathbf{S}}(\phi, z_t, q_t) = (q_t v^{\mathbf{S}} - (1 - q_t))(1 - \phi)z_t + v^{\mathbf{S}}(1 - z_t) - \kappa(\phi). \tag{15}$$

Since this is a concave problem, its solution is given by the first-order condition

$$\kappa'(\phi) = z_\tau \left(1 - q_\tau - q_\tau v^{\mathbf{S}}\right). \tag{16}$$

Note that the RHS of (16) is positive since, by assumption, $q < \hat{q} = \frac{1}{1+v}$. The assumptions that κ is convex, $\kappa'(0) = 0$, and $lim_{\phi \to 1} k(\phi) = \infty$ imply that at the optimum the FOC holds with equality. Denote the optimal choice of quality by $\phi^*(z_t, q_t)$.

Since **S** may decide to terminate the project without scrutiny, strictly positive monitoring is optimal only if his expected value from optimal monitoring is nonnegative, i.e.,

$$EU^{\mathbf{S}}(z_t, q_t) \equiv (1-z)v + (1-\phi^*(z,q))z(q(v+1)-1) - \kappa(\phi^*(z,q)) \geq 0. \tag{17}$$

Since $\kappa \geq 0$ this condition also implies that, conditional on receiving the realization s_N, **S** would rather launch the project than terminate it.

From equation (16) and the assumptions that κ is increasing and convex, it follows that, for a fixed q, $\phi^*(z, q)$ is decreasing in z and converges to zero as z does. Hence, for any $q < \hat{q}$, this condition holds for a positive z that is sufficiently small. Moreover, for a given q, the LHS of (17) is decreasing in z and is strictly negative if $z = 1$. Hence, there exists a unique z

for which (17) holds with equality. Denote this critical probability of the breakthrough being fabricated by $\overline{z}(q)$ and denote $\overline{\phi}(q) \equiv \phi^*(\overline{z}(q), q)$.

Lemma C.1. *$\overline{\phi}(q)$ is decreasing in q and $\overline{z}(q)$ is increasing in q.*

Proof. This result follows from the implicit function theorem and the definitions

$$0 = EU^{\mathbf{S}}(\overline{z}(q_t), q_t)$$
$$0 = \kappa'(\overline{\phi}(q_t)) - \overline{z}(q_t)(1 - q_t(1 + v^{\mathbf{S}})).$$

□

C.2. **The Strategic Interaction.** Having analyzed the monitor's problem, we now revert to analyzing the strategic interaction between **F** and **S**, when **S** has access to a monitoring technology. For this analysis, we impose the assumption of large conflict, $\hat{q}^{\mathbf{S}} > \overline{q}^{\mathbf{F}}$.

First, as in all of the variants considered in this paper, if $q_0 \leq \overline{q}^{\mathbf{F}}$, then the unique efficient equilibrium is for **F** to take over the project. To see this, note that since **S** cannot learn directly about the quality of the project, **F** learning honestly maximizes the utility of both players. Since for such priors **F** is willing to learn honestly as a DM, this is the unique Pareto-efficient equilibrium.

In the rest of this section, we focus on priors such that $q_0 \in (\overline{q}^{\mathbf{F}}, \hat{q}^{\mathbf{S}})$. If $\overline{q}^{\mathbf{F}} < q_0 < \hat{q}^{\mathbf{S}}$, then in any pure strategy equilibrium the project is terminated. Hence, any equilibrium in which the players obtain strictly positive expected utility is in mixed strategies. Importantly, in a mixed strategy equilibrium, **F** must randomize between fabricating a breakthrough and continuing to lean at every t such that $q_t > \overline{q}^{\mathbf{F}}$. To see this, note that were **S** to approve the project with probability one for such t (which is his unique best response if **F** does not fabricate breakthroughs), then **F** would rather have **S** launch the project immediately at that point than continue learning, as she does not derive any direct benefit from **S**'s monitoring.

Recall that $\overline{\phi}(q_\tau)$ is the maximal probability with which **S** can detect that a breakthrough was fabricated if he decides to use his monitoring technology optimally. Moreover, since q_t decreases over time, Lemma C.1 implies that this probability increases over time.

For convenience, we restate the main result of Section 6.2, namely Proposition 6.2.

Proposition. *There exists some* $q^{**} \in (\overline{q}^{\boldsymbol{F}}, \hat{q}^{\boldsymbol{S}})$ *such that, for any* $q_0 \in (\overline{q}^{\boldsymbol{F}}, q^{**})$*, the unique Pareto-efficient equilibrium is characterized by* $\tau^{**} < \tau^*$ *such that*

(1) *for all* $t \leq \tau^{**}$*, if* ***S*** *receives the project at* t*, he terminates the project with probability* $1 - \theta_t$ *and otherwise monitors* ***F*** *according to* $\phi_t = \overline{\phi}(q_t)$ *(partial-mistrust phase),*

(2) *for all* $t \in (\tau^{**}, \tau^*)$*, if* ***S*** *receives the project at* t*, he monitors* ***F*** *according to* $\phi_t = \phi_t^*(z_t, q_t) < \overline{\phi}(q_t)$ *(verification phase),*

and ***F*** *takes over the project at* τ^**.*

The path by which we establish Proposition 6.2 is as follows. First, we define a specific equilibrium in which **S**'s behavior can be decomposed into at most two phases: an initial partial-misstrust phase in which he randomizes between terminating the project and maximal scrutiny given belief q_t, and a later verification phase in which he scrutinizes the project (non-maximally) with certainty. In the former phase, **F** is incentivized to learn by increasing the probability of scrutinizing the project as time goes by, and in the later phase, **F** is incentivized to learn by decreasing the amount of scrutiny the project receives as time goes by. Finally, we show that if this equilibrium exists, it Pareto dominates any other equilibrium.

Recall that **F** must be indifferent between all stopping times $\tau < \tau^*$. In this application this indifference condition (equation (3)) can be written as

$$\lambda^{\mathbf{F}} q_\tau \left[W_\tau^{NB} - \theta_\tau v^{\mathbf{F}}\right] + c^{\mathbf{F}} = \frac{dW_\tau^{NB}}{d\tau}. \tag{18}$$

To construct the efficient equilibrium, we analyze separately this indifference condition under the assumption that $\theta_t < 1$ on some interval (a partial-misstrust phase), and under the assumption that $\theta_t = 1$ on some interval (verification phase).

Case 1: $\theta_t < 1$ *on some interval.* In this case, **S** must be indifferent between terminating the project and scrutinizing **F**'s submission. Thus, his beliefs about the probability that the breakthrough was fabricated must equal $\overline{z}(q_t)$. However, as q_t changes over time, the probability of **S** approving a project submitted with a fabricated breakthrough, conditional on examining it, is not constant over time. In particular, this probability is given by $\overline{\phi}(q_t)$, which is increasing in t (Lemma C.1). Let $\underline{P}_t = 1 - \overline{\phi}(q_t)$ denote the probability that the fabrication of the breakthrough is not detected by **S** (and hence the project is launched).

In an interval where $\theta_t < 1$ it holds that $W^{NB} = \theta_t \underline{P}_t(q_t(v^{\mathbf{F}} + 1) - 1)$ and so the differential equation (18) becomes

$$c^{\mathbf{F}} + \lambda^{\mathbf{F}} q_t \theta_t v^{\mathbf{F}}(\underline{P}_t - 1) = (q_t(v^{\mathbf{F}} + 1) - 1)\frac{d(\theta_t \underline{P}_t)}{dt}. \tag{18b}$$

From this representation we can derive the following result.

Lemma C.2. *If $c^{\mathbf{F}} + \lambda^{\mathbf{F}} q_t \theta_t v^{\mathbf{F}}(\underline{P}_t - 1) > 0$, then $\dot{\theta}_t > 0$.*

Proof. In the mixing region **F** would rather launch the project than terminate it and so $(q_t(v^{\mathbf{F}} + 1) - 1) > 0$. Since $\underline{P}_t$ is decreasing in t, it follows that if $c^{\mathbf{F}} + \lambda^{\mathbf{F}} q_t \theta_t v^{\mathbf{F}}(\underline{P}_t - 1) > 0$, it must be the case that $\dot{\theta}_t > 0$. □

Case 2: $\theta_t = 1$ on some interval. In this case $W^{NB} = P_t(q_t(v^{\mathbf{F}} + 1) - 1)$, where P_t is the probability that a project submitted with a fabricated breakthrough at t is launched, and so the differential equation (18) becomes

$$c^{\mathbf{F}} + \lambda^{\mathbf{F}} q_t v^{\mathbf{F}}(P_t - 1) = \dot{P}_t(q_t(v^{\mathbf{F}} + 1) - 1). \tag{18c}$$

We construct an equilibrium in which **F** takes over the project at τ^*. Thus, the boundary condition for W^{NB} is $W^{NB}_{\tau^*} = q_{\tau^*}(v^{\mathbf{F}} + 1) - 1$, which implies that $P_{\tau^*} = 1$. Note that this implies that $\dot{P}_{\tau^*} > 0$. Moreover, recall that by assumption $c^{\mathbf{F}} < v^{\mathbf{F}} \lambda^{\mathbf{F}}$ and so from equation (18c) it follows that once P_t becomes increasing it remains so. Let t'' be the infimum of times $\{s : \dot{P}_{t''} \geq 0, \forall t \geq s\}$ (note that $t'' < \tau^*$), and define $q^{**} = q_{t''}$.

In equilibrium it must be the case that $P_t \geq \underline{P}_t$. Hence, the evolution of P_t derived in case 2 above need not be feasible in equilibrium. To construct an equilibrium, let τ^{**} denote the unique t for which $P_t = \underline{P}_t$, where P_t is the solution to (18c) with the boundary condition $P_{\tau^*} = 1$. Such a t exists and is unique as P_t is increasing and satisfies $P_{\tau^*} = 1$, while $\underline{P}_s$ is decreasing and equals 1 if $q_s = \hat{q}^{\mathbf{S}}$.

We construct our equilibrium as follows. On $[\tau^{**}, \tau^*)$ the probability of the project being approved after a fabricated breakthrough is given by the solution to equation (18c) and $\theta_{\tau^*} = 1$, note that by the definition of τ^{**}, this probability is a best response to some belief for **S**. For $t < \tau^{**}$ the probability of the project being approved following a fabricated breakthrough is $\underline{P}_t$, and θ_t is given by the solution to equation (18b) with the boundary condition $\theta_{\tau^{**}} = 1$, which is required for the continuity of W^{NB} at τ^{**}. Since in this range $\underline{P}_t$ is greater than the value of P_t that would have been obtained from the solution to equation (18c), it follows that $(q_t(v^{\mathbf{F}}+1)-1)\frac{d(\theta_t \underline{P}_t)}{dt} > (q_t(v^{\mathbf{F}}+1)-1)\dot{P}_t) > 0$, and so by Lemma C.2, $\dot{\theta}_t > 0$. Hence, our equilibrium exists if $\tau^{**} \leq 0$, or $\tau^{**} > 0$ and $\theta_0 \geq 0$.

Next, we show that if there exists an alternative equilibrium $\hat{E}$, then the equilibrium described above also exists. To see this, recall that since **S** cannot learn about the project's quality, **F** must mix with a full support on $[0, \tau^*]$ and continuation values in $[0, \tau^*]$ are given by the differential equation (18). The continuation play in $\hat{E}$ generates a continuation value for **F**, which, in turn, pins down $\hat{\theta}_{\tau^*}$ and $\hat{P}_{\tau^*}$ in the alternative equilibrium. If $\hat{\theta}_{\tau^*} = \hat{P}_{\tau^*} = 1$, then **F**'s strategy in $[0, \tau^*]$ is identical in the alternative equilibrium and the one we are characterizing. Hence, our equilibrium exists. If $\hat{\theta}_{\tau^*} \leq 1, \hat{P}_{\tau^*} \leq 1$ (with one strict inequality), then the boundary condition for $\hat{E}$ is below the boundary condition for our equilibrium. Hence, $\hat{W}_t^{NB} < W_t^{NB}$, for every $t \leq \tau^*$. In particular, this implies that $\hat{\theta}_0 \leq \theta_0$ and $\hat{P}_0 \leq P_0$. Since $\hat{\theta}_0, \hat{P}_0$ are feasible strategies, our equilibrium also exists.

Finally, we show that our equilibrium is efficient. For **F** this is immediate. **F**'s equilibrium value is W_0^{NB}, and since this value is given by an ODE it is increasing in $W_{\tau^*}^{NB}$. Since $W_{\tau^*}^{NB}$ is increasing in P_{τ^*} and θ_{τ^*}, it follows that the $W_{\tau^*}^{NB}$ is maximized in our equilibrium, where $\theta_{\tau^*} = P_{\tau^*} = 1$.

For **S** the argument is different. First, note that in our equilibrium the probability that **F** fabricates a breakthrough at $t < \tau^*$ is weakly lower than in any other equilibrium. This is the case, as the amount of monitoring is increasing in the probability of the breakthrough being fabricated, and in our equilibrium the amount of monitoring is lower than in any other equilibrium (P is higher). Prior to τ^{**} there must be a partial-mistrust phase in both equilibria; thus, in this range **F**'s strategy is the same in both equilibria. Thus, the difference between **S**'s payoff in the alternative equilibrium and in our equilibrium can be analyzed by determining the impact of reducing the probability that **F** submits the project with a fabricated breakthrough at $t \in (\tau^{**}, \tau^*)$ and increasing the probability that the project is submitted after an authentic breakthrough at some $s \in (t, \omega)$ or terminated at ω. Note that **S**'s value from receiving a project following a fabricated breakthrough is strictly negative; his utility from receiving the project after an authentic breakthrough, and monitoring it for any plausible amount of time, is positive; and his utility from the project being terminated at ω is zero. Hence, **S** also prefers our equilibrium to any alternative equilibrium.