Sequential Learning - Online Appendices

Yair Antler, Daniel Bird and Santiago Oliveros

Online Appendix A: 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 $q_0 \le \overline{q}^{\mathbf{F}}$ (Lemma 1 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.[24]

The main results in this appendix correspond to the characterization, existence, and efficiency results of Section IV. The next proposition establishes that, if there is a mixed strategy equilibrium, the equilibrium characterized in Proposition 3 is the Pareto-dominant equilibrium in the class of equilibria in mixed strategies.

Proposition 9. *If $\boldsymbol{E}$ is a Pareto-undominated equilibrium in which $\boldsymbol{F}$ uses a mixed strategy, then it is the equilibrium characterized in Proposition 3.*

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 IV.C is obtained by comparing the mixed strategy equilibrium characterized in Proposition 3 to the Pareto-dominant equilibrium in the class of pure strategy equilibria, 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 A.1. *If $\underline{q}^{\boldsymbol{S}} \ge \frac{1}{1+v^{\boldsymbol{F}}}$, 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-undominated pure strategy equilibria is $\{\Sigma(\tau^{\boldsymbol{F}}) : \tau^{\boldsymbol{F}} \le \hat{\tau}\}$, where $\hat{\tau} > 0$ is defined implicitly by*

$$1 - l_{\underline{q}^S} v^{\boldsymbol{F}} \equiv \frac{c^{\boldsymbol{F}}}{\lambda^{\boldsymbol{F}}} l_{q_0} \left(1 - e^{-\lambda^{\boldsymbol{F}} \hat{\tau}}\right) + c^{\boldsymbol{F}} \hat{\tau}. \tag{A1}$$

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

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

deviation for $\mathbf{F}$ 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, $\mathbf{F}$ cannot manipulate $\mathbf{S}$ 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-undominated as $\mathbf{F}$'s ($\mathbf{S}$'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 Pareto-dominant equilibrium in pure strategies is for $\mathbf{S}$ to take over the project at $t = 0$. As the Pareto-dominant equilibrium in pure strategies in a small conflict is identical to the Pareto-dominant equilibrium in pure strategies in a large conflict, the same argument used to derive Proposition 5 yields the following result.

Proposition 10. *Assume that $\underline{q}^{\boldsymbol{S}} \geq \frac{1}{1+v^{\boldsymbol{F}}}$. If the equilibrium characterized in Proposition 3 exists, then it is the Pareto-dominant 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-undominated equilibria in pure strategies. Note that $\mathbf{F}$'s preferred equilibrium in pure strategies is the one in which $\mathbf{S}$ takes over the project at time zero. Hence, by the same argument used in Section IV.C, if the equilibrium characterized in Proposition 3 exists, then it is Pareto-undominated. By contrast, $\mathbf{S}$'s preferred equilibrium in pure strategies is $\Sigma(\hat{\tau})$. Hence, comparing $\mathbf{S}$'s payoff in the mixed strategy equilibrium to his payoff as a DM—the comparison used to derive Proposition 5—is not sufficient to establish that $\mathbf{S}$ is better off in the equilibrium with strategic uncertainty.

Nevertheless, the method of the proof used to establish Proposition 5 can still be used to derive an efficiency result. As explained in Section IV.C, $\mathbf{S}$'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 $\mathbf{F}$'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 $\Sigma(\hat{\tau})$, $\mathbf{S}$'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 $\mathbf{S}$ prefers the mixed strategy equilibrium to the pure strategy equilibrium depends on the value of these signals to $\mathbf{S}$.

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 $\mathbf{S}$'s belief after receiving the signal remains near the prior. As $\mathbf{S}$'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 $\mathbf{S}$'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 Pareto-dominant equilibrium if $\hat{\tau}$ is small enough. However, if $\hat{\tau}$ is large, it may be the case that $\mathbf{S}$ 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 11. *If the equilibrium characterized in Proposition 3 exists, then: (i) it is Pareto-undominated, and (ii) there exists some $\epsilon > 0$ such that if $\hat{\tau} < \epsilon$, it is the Pareto-dominant equilibrium.*

We now establish that Proposition 3 characterizes the Pareto-dominant 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 $\mathbf{F}$ to take over the project at τ^*, which determines the efficient continuation value and hence the equilibrium behavior.

PROOF OF PROPOSITION 9:

To establish this result, we first show that mixed strategies cannot be part of an efficient equilibrium when $q_0 > \overline{q}^{\mathbf{S}}$ (Lemma A.2). We then show that, at some point in time, moral hazard no longer hinders the interaction (this result requires using the technical Lemma A.3). Lemma A.4 shows that the interaction evolves in the same manner that we described in Proposition 3. Finally, we establish that in an efficient equilibrium where $\mathbf{F}$ uses a mixed strategy she must take over the project when she can be trusted to do so (Lemmas A.5 and A.6).

In this proof we often use the strategy profile where 1) $\mathbf{F}$ does not fabricate breakthroughs if $q_t > q_{\omega(\mathbf{E})}$ and submits the project otherwise, and 2) $\mathbf{S}$ best responds to $\mathbf{F}$'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 $\mathbf{F}$ 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 $\mathbf{F}$'s equilibrium payoff in **E'** at time zero is greater than her payoff from launching the project immediately.

Lemma A.2. *If $q_0 > \overline{q}^{\mathbf{S}}$, then $\mathbf{F}$ uses a pure strategy in any efficient equilibrium.*

PROOF:

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

If $t = 0$ belongs to the support of $G^{\mathbf{F}}$, then there is an atom at $t = 0$. In this case, $\mathbf{F}$'s equilibrium payoff at $t = 0$ is equal to the payoff from launching the project at $t = 0$. Note that the sum of $\mathbf{F}$'s and $\mathbf{S}$'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 1 we defined the function

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

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 $\mathbf{E}$. Therefore, $t = 0$ cannot be in the support of $G^{\mathbf{F}}$.

Since $t = 0$ is not in the support of $\mathbf{F}$'s strategy, the support of $\mathbf{F}$'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 1 it follows that **E'** Pareto dominates $\mathbf{E}$. To see that **E'** is an equilibrium, note that, in $\mathbf{E}$, 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 A.2, we focus on the case where $q_0 \leq \overline{q}^{\mathbf{S}}$ in the rest of this proof.

Lemma A.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 A.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 1, **E'** Pareto dominates **E**. Hence, for the rest of the proof we focus on the case where $\psi > 0$. Note that since we are focusing on the case where $q_0 \leq \overline{q}^{\mathbf{S}}$, it follows that $q_\psi < \overline{q}^{\mathbf{S}}$.

Lemma A.4. *$q_t^{\mathbf{S}}$ 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 A.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 1, **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 A.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 A.5 it follows that $\psi < \omega$. Together with Lemma 1, 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 3 show that if **E** is a Pareto-efficient equilibrium, then the continuation equilibrium at τ^* must be efficient. Thus, it suffices to show that **F** taking over the project at τ^* is the Pareto-dominant equilibrium among all equilibria where the project is terminated at ω.

Lemma A.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 Pareto-dominant equilibrium.*

PROOF:

Proposition 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 1 shows that **F**'s utility from any equilibrium in which she terminates the project at ω is weakly less than her utility from learning 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 until ω and then terminating the project and his utility from launching the project immediately.

Online Appendix B: Monitoring

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

B1. 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

$$EU^{\mathbf{S}}(\phi, z_t, q_t) = (q_t v^{\mathbf{S}} - (1-q_t))(1-\phi)z_t + v^{\mathbf{S}}(1-z_t) - \kappa(\phi). \tag{B1}$$

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

$$\kappa'(\phi) = z_\tau \left(1 - q_\tau - q_\tau v^{\mathbf{S}}\right). \tag{B2}$$

Note that the RHS of (B2) 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.,

$$EU^{\mathbf{S}}(z_t, q_t) \equiv (1-z)v + (1-\phi^*(z,q))z(q(v+1)-1) - \kappa(\phi^*(z,q)) \geq 0. \tag{B3}$$

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 (B2) 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 (B3) is decreasing in z and is strictly negative if $z = 1$. Hence, there exists a unique z for which (B3) 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 B.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

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

B2. 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 Pareto-dominant 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 maximizes the utility of both players. Since for such priors **F** is willing to learn as a DM, this is the Pareto-dominant 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 B.1 implies that this probability increases over time.

For convenience, we restate the main result of Section V.B, namely Proposition 7.

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 Pareto-dominant 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 7 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-mistrust 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

$$\lambda^{\mathbf{F}} q_\tau \left[W_\tau^{NB} - \theta_\tau v^{\mathbf{F}} \right] + c^{\mathbf{F}} = \frac{dW_\tau^{NB}}{d\tau}. \tag{B4}$$

To construct the Pareto-dominant equilibrium, we analyze separately this indifference condition under the assumption that $\theta_t < 1$ on some interval (a partial-mistrust 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 B.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 (B4) becomes

$$(B4b) \qquad c^{\mathbf{F}} + \lambda^{\mathbf{F}} q_t \theta_t v^{\mathbf{F}}(\underline{P}_t - 1) = (q_t(v^{\mathbf{F}}+1)-1)\frac{d(\theta_t \underline{P}_t)}{dt}.$$

From this representation we can derive the following result.

Lemma B.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 (B4) becomes

$$(B4c) \qquad c^{\mathbf{F}} + \lambda^{\mathbf{F}} q_t v^{\mathbf{F}}(P_t - 1) = \dot{P}_t(q_t(v^{\mathbf{F}}+1)-1).$$

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 (B4c) 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 (B4c) 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 (B4c) 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 (B4b) 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 (B4c), 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 B.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 (B4). 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, $\hat{W}^{NB}_t < W^{NB}_t$, 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 a Pareto-efficient equilibrium. 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.