Online Appendix – NOT FOR PUBLICATION

"Experimentation and Persuasion in Political Organizations"

Alexander V. Hirsch. July 16, 2015.

This Online Appendix is divided into six parts. Appendix A provides a general statement of strategies and equilibria, and describes how we do equilibrium selection when there are multiple equilibria. Appendix B proves Proposition 3. Appendix C proves verbal statements in the analysis of experimentation about the principal's commitment problem and institutional solutions. Appendix D contains accessory lemmas used in the other proofs. Appendix E treats the variant of the model with differing preferences rather than beliefs discussed in the analysis of experimentation. Appendix F contains supplemental Mathematica code employed in the proofs.

A Equilibrium Characterization and Selection

We begin by introducing additional notation and providing a general equilibrium characterization for the baseline model. This requires allowing mixed strategies for the principal, which are omitted from the main text for simplicity. As in the main text, a strategy for the agent consists of two functions $e^1(x^1), e^2(x^1, e^1, y^1, x^2)$ to $[0,1]$ mapping histories to effort. To allow for mixed strategies for the principal, we now denote the principal's strategy as a probability $p^1 \in [0,1]$ of initially choosing policy a, and a set of probabilities $p^{x^1}_{y^1} \in [0,1]$ of sticking with the initial policy x^1 after outcome y^1 for every (x^1, y^1). Throughout we will also use $P_i(\cdot)$ to denote probabilities evaluated with respect to the prior of player i – so $P_2(\omega = b) = 1 - \theta_2$ denotes the agent's prior belief that b is the correct policy.

The Agent's Problem In the second period the agent exerts effort $e^2(x^1, e^1, y^1, x^2) = \lambda P_2(\omega = x^2 \mid x^1, e^1, y^1)$ and his expected utility is $\frac{\lambda}{2}\left[P_2(\omega = x^2 \mid x^1, e^1, y^1)\right]^2$. For simplicity, denote his prior that $\omega = a$ as θ, his initial effort e^1 as e, and p^a_s and p^a_f as p_s and p_f. Then it is straightforward to show that his expected two-period utility when $x^1 = a$ as a function of first period effort is

$$\begin{aligned} U_2(e, \theta, p_s, p_f) &= \left(\theta e - \frac{e^2}{2\lambda}\right) + \frac{\lambda}{2}\left(p_f \cdot \theta\left(e + (1-e)\,h(e,\theta)\right) + (1-p_f)\cdot(1-\theta)\left(1 - h(e,\theta)\right)\right) \\ &\quad + (p_s - p_f)\,\frac{\lambda}{2}\theta e \end{aligned} \tag{A.1}$$

By symmetry the agent's expected utility from exerting effort e^1 on an arbitrary policy $x \in \{a, b\}$ is $U_2\left(e^1, P_2(\omega = x), p^x_s, p^x_f\right)$. Also note that $U_2(\cdot)$ is distinct from the expression $U(\cdot)$ for the agent's objective function as defined in the main Appendix because the latter amalgamated utilities from the NLB and the GWL, and did not account for mixed strategies.

The Principal's Problem The principal's second period policy choices must be interim-optimal given $e^2(x^1, e^1, y^1, x^2)$ and her posteriors computed with the agent's equilibrium strategy. Thus, she must stay with the initial policy x if it succeeds ($p_s^x = 1$), and can only stay with the initial policy if it fails $\left(p_f^x > 0\right)$ i.f.f. $h\left(e^1(x), P_1(\omega = x)\right) \geq 1 - h\left(e^1(x), P_2(\omega = x)\right)$. If the inequality is strict then $p_f^x = 1$.

For simplicity denote e^1 as e and p_s^a and p_f^a as p_s and p_f. In period 1 the principal's expected utility from selecting policy a when she expects first period effort e and future equilibrium behavior is,

$$U_1\left(e^1, \theta_1, \theta_2, p_s, p_f\right) = \theta_1\left(e + \lambda\left(ep_s + (1-e)\, p_f \cdot h\left(e, \theta_2\right)\right)\right) + (1-\theta_1)\left(1 - p_f\right)\lambda\left(1 - h\left(e, \theta_2\right)\right) \tag{A.2}$$

By symmetry her expected utility from some x is $U_1\left(e^1, P_1(\omega = x), P_2(\omega = x), p_s^x, p_f^x\right)$.

Equilibrium Conditions Strategies $\left(x^1, p_s^a, p_f^a, p_s^b, p_f^b\right)$ and $\left(e^1(x^1), e^2(x^1, e^1, y^1, x^2)\right)$ are an equilibrium if and only if they satisfy the following conditions.

(Agent Optimality)

1. $e^2(x^1, e^1, y^1, x^2) = \lambda P_2\left(\omega = x^2 \mid x^1, e^1, y^1\right)$ (the agent optimizes in the second period)

2. $e^1(x) \in \underset{e^1 \in [0,1]}{\arg\max}\left\{U_2\left(e^1, P_2(\omega = x), p_s^x, p_f^x\right)\right\}$ $\forall x \in \{a, b\}$ (the agent optimizes in the first period given the principal's strategy and expectations about his own future effort)

(Principal Optimality)

1. $p_s^x = 1$ $\forall x \in \{a, b\}$ (the principal always stays after success)

2. $\forall x \in \{a, b\}$ $p_f^x \geq 0 \iff h\left(e^1(x), P_1(\omega = x)\right) \geq 1 - h\left(e^1(x), P_2(\omega = x)\right)$ and $= 1$ if satisfied with strict inequality (the principal only stays after failure if it is interim-optimal given on-path posteriors)

3. $x^1 \in \underset{x \in \{a,b\}}{\arg\max}\left\{U_1\left(e^1(x), P_1(\omega = x), P_2(\omega = x), p_s^x, p_f^x\right)\right\}$ (the principal's initial policy choice maximizes her expected continuation value).

Equilibrium Selection and Notation Lemma 2 in the Appendix D proves the following two statements: (i) whenever experimenting is an equilibrium in a subgame x^1, it is the optimal strategy for the principal even if she could precommit to her future decisions, and (ii) if experimenting with $x^1 = a$ is *not* an equilibrium of the subgame following $x^1 = a$, then the unique equilibrium is rigid implementation. Together, these statements imply that we can select the equilibria that are best for the principal by considering only pure strategy equilibria, and choosing experimentation for sure in the subgame commencing with policy $x^1 \in \{a, b\}$ (i.e., $p_s^{x^1} = 1$ and $p_f^{x^1} = 0$) whenever it is an equilibrium.

With this selection and restriction to pure strategies, we now introduce simplified notation for the agent's best responses and principal's utility. First let $e^s(\theta_2)$ denote the agent's first-period best response when the principal's pure strategy is $(x^1 = a, s)$, where $s \in \{R, E\}$ denotes whether the principal (**R**)igidly implements or (**E**)xperiments with the initial policy. It is also helpful to state implicit characterizations of $e^R(\theta_2)$ and $e^E(\theta_2)$ so we can approximate their values in several proofs. The FOC from the proof of Lemma 1 yields that

$$e^R(\theta_2) = \lambda\theta_2 \left(1 + \frac{\lambda}{2} k(e, \theta_2)\right) \quad \text{and} \quad e^E(\theta_2) = \lambda\theta_2 \left(1 + \frac{\lambda}{2}(k(e, \theta_2) + 1)\right),$$

where $k(e, \theta) = (1 - h(e, \theta))^2$.

Second, let $V_1^s(\theta_1, \theta_2)$ denote the principal's two-period expected utility when her pure strategy is $(x^1 = a, s)$ and the agent best-responds, so $V_1^R(\theta_1, \theta_2) = U_1\left(e^R(\theta_2), \theta_1, \theta_2, 1, 1\right)$ and $V_1^E(\theta_1, \theta_2) = U_1\left(e^E(\theta_2), \theta_1, \theta_2, 1, 0\right)$. Third, let $\bar{\theta}(\theta_2)$ be the unique solution to

$$h\left(e^R(\theta_2), \bar{\theta}(\theta_2)\right) = 1 - h\left(e^R(\theta_2), \theta_2\right), \tag{A.3}$$

By the equilibrium characterization, experimenting with $x^1 = a$ is an equilibrium of that subgame i.f.f. $\theta_1 \leq \bar{\theta}(\theta_2)$, and it is easy verified that $\bar{\theta}(\theta_2) > 1 - \theta_2$. Finally, by symmetry the agent's effort on $x^1 = b$ is $e^s(1 - \theta_2)$, the principal's utility is $V_1^s(1 - \theta_1, 1 - \theta_2)$, and the threshold for experimentation with b is $1 - \theta_1 < \bar{\theta}(1 - \theta_2) \iff \theta_1 > 1 - \bar{\theta}(1 - \theta_2)$.

B Proof of Proposition 3

In Lemma 3 in the Supplemental Proofs, we prove the following handy property: fixing the principal's experimentation decisions down each path of play, if she prefers $x^1 = a$ given beliefs $\hat{\theta}_1$ (in the sense of ex-ante expected utility) then she also prefers it for all higher beliefs $\theta_1 > \hat{\theta}_1$ (of course, by symmetry if she prefers $x^1 = b$ at beliefs $\hat{\theta}_1$ then she also prefers it for all beliefs $\theta_1 < \hat{\theta}_1$). We call this property "preference monotonicity" and employ it in this and the subsequent proofs.

Part 1 We prove that deference expands in GWL as compared to the NLB both in the baseline model when $\lambda \geq \underline{\lambda} \approx .23505$, and when the principal can commit $\forall \lambda$. When λ goes below $\underline{\lambda}$ and the principal can't commit, there opens up a very tiny interval of beliefs where the principal *would* have deferred in the No Learning Benchmark, but does not in the Game with Learning. Using Mathematica we find that the size of this interval is maximized when $(\lambda = .123, \theta_2 = .545)$, and at these values is $\theta_1 \in (.5, .508)$. In this interval the principal is actually best off deferring by experimenting with b, but can't credibly commit do so because the agent is working so (unrealistically) little that almost nothing is learned from failure. Rather than rigidly implement b, she experiments with a to better motivate the agent.

To prove that deference strictly expands when the principal can commit, we argue that the following property proved in analytically in Lemma 4 in the Supplemental Proofs is sufficient: a principal with beliefs $\theta_1 = 1 - \theta_2$ strictly prefers experimenting with b to experimenting with a. To see this, observe that for $\theta_1 \leq 1 - \theta_2$ experimenting with a is strictly better than any other strategy with a, by $1 - \theta_2 < \bar{\theta}(\theta_2)$ and Lemma 2. If experimenting with b is also strictly better than experimenting with a, then it is strictly better than any other strategy with a, and $x^1 = b$ must be chosen (either experimenting or rigidly implementing). Thus within the entire deference region from the NLB the principal defers in the GWL. Because the preference for deference is there strict, the deference region must strictly expand.

To prove that the deference region strictly expands when the principal can't commit, observe that the preceding argument proves the principal defers whenever her beliefs θ_1 are $\in \left[1-\bar{\theta}\left(1-\theta_2\right), 1-\theta_2\right]$, because this condition implies that the principal experiments in the subgame following $x^1 = b$ without commitment. However, to show that the deference region in the GWL contains the entire deference region in the NLB, we must also show that the principal will still choose $x^1 = b$ even when $\theta_1 \in \left[\frac{1}{2}, 1-\bar{\theta}\left(1-\theta_2\right)\right]$ and she rigidly implements in the subgame following $x^1 = b$. If θ_2 is such that $\frac{1}{2} \geq 1-\bar{\theta}\left(1-\theta_2\right)$ then this region is empty. If $1-\bar{\theta}\left(1-\theta_2\right) > \frac{1}{2}$, then we require that a principal with beliefs $\theta_1 = 1-\bar{\theta}\left(1-\theta_2\right)$ weakly prefers rigidly implementing b to experimenting with a. If this holds then by preference monotonicity a principal with beliefs $\theta_1 \in \left[\frac{1}{2}, 1-\bar{\theta}\left(1-\theta_2\right)\right]$ who would rigidly implement b if chosen also prefers that to experimenting with a, and so selects b initially. If it fails, then for some principal beliefs a little bit below $1-\bar{\theta}\left(1-\theta_2\right) < 1-\theta_2$, the principal will experiment with a in the GWL when she would have deferred in the NLB.

Finally, we prove that a principal with beliefs $\theta_1 = 1-\bar{\theta}\left(1-\theta_2\right)$ weakly prefers rigidly implementing b to experimenting with a in Lemma 5 in the Supplemental Proofs with the aid of Mathematica, if and only if $\lambda > \underline{\lambda} \approx .23505$.

Part 2 We first argue that following three conditions on θ_2 are jointly *sufficient* for the principal to always defer in the first period regardless of her own beliefs; 1) $V_1^E\left(0, 1-\theta_2\right) \geq V_1^R\left(1, \theta_2\right)$, 2) $V_1^R\left(1, \theta_2\right) > V_1^E\left(1, \theta_2\right)$, 3) $V_1^R\left(\bar{\theta}\left(1-\theta_2\right), 1-\theta_2\right) > V_1^E\left(1-\bar{\theta}\left(1-\theta_2\right), \theta_2\right)$. Conditions (1) and (2) jointly imply that experimenting with b is better than both experimenting with or rigidly implementing a when $\theta_1 = 1$; by preference monotonicity (proved in Lemma 3 in the Supplemental Proofs) this also implies that experimenting with b is better $\forall \theta_1 \in [0,1]$. Thus, whenever experimenting with b is an equilibrium strategy $\left(\theta_1 \geq 1-\bar{\theta}\left(1-\theta_2\right)\right)$ it is chosen. Now whenever experimenting with b is not an equilibrium strategy $\left(\theta_1 < 1-\bar{\theta}\left(1-\theta_2\right)\right)$, the principal compares rigidly implementing b to experimenting with a; again applying preference monotonicity, condition (3) implies that she prefers the former $\forall \theta_1 \leq 1-\bar{\theta}\left(1-\theta_2\right)$. All possible principal beliefs are covered, which completes the argument. We next argue that condition (1) is *necessary* for the principal to always defer regardless of her own beliefs; if it fails then $V_1^R(1,\theta_2) > V_1^E(0,1-\theta_2)$. For $\theta_1 > \bar{\theta}(\theta_2)$ the principal would rigidly implement a and experiment with b, and by continuity she also prefers rigidly implementing a to experimenting with b for θ_1 sufficiently close to 1. Thus for such θ_1 she selects $x^1 = a$ in equilibrium and does not defer.

Finally, Lemma 6 in the Supplemental Proofs proves analytically that when $\lambda > \hat{\lambda}$ (where $\hat{\lambda}$ is the unique solution to $\lambda(1+\lambda)\left(1+\frac{\lambda}{2}\right) = 1$ and $\approx .5214$), each of the three conditions $k \in \{1,2,3\}$ holds for θ_2 in a nonempty interval $(0,\varepsilon_k)$.[27] Since they then all hold for $\theta_2 \in (0,\varepsilon_k)$, $\lambda > \hat{\lambda}$ is therefore sufficient for existence of a range of θ_2 where the principal always defers. In addition, the supplemental mathematica code verifies that condition 1 fails $\forall\theta_2 > 0$ when $\lambda \leq \hat{\lambda}$, which is equivalent to

$$\frac{V_1^E(0,1-\theta_2) - V_1^R(1,\theta_2)}{\lambda\theta_2} < 0 \ \forall\theta_2 \in \left[0,\frac{1}{2}\right] \text{ when } \lambda < \hat{\lambda}.$$

$\lambda > \hat{\lambda}$ is thus also *necessary* for existence of a range of θ_2 where the principal always defers.

C Underexperimentation and Commitment

In this section we formally state, and then prove, several verbal statements in the analysis of experimentation about the principal sometimes "underexperimenting," and institutional arrangements that can help the principal solve this commitment problem.

The first result pertains the possibility of "underexperimentation" in equilibrium. Formally, we say that the principal *underexperiments* if she rigidly implements a policy x^1 on the equilibrium path of play, but a strategy of experimenting with some policy x^1 or $\neg x^1$ would yield higher ex-ante higher expected utility. In other words, she underexperiments if she implements a policy in equilibrium, but would experiment with some policy if she could

[27]Note that we have already shown that property (3) holds $\forall\theta_2$ when $\lambda > \underline{\lambda}$ in Lemma 5; however, the proof is computational. In Lemma 6 we prove analytically that property (3) holds for sufficiently low θ_2 for any $\lambda \in [0,1]$.

commit to her entire two period strategy ex-ante.

Underexperimentation happens in the model because it is possible that a principal with relatively strong beliefs in favor of a policy would be better off ex-ante committing to experiment with that policy in order to better motivate the agent, but after actually observing failure she will want to renege on the experiment and persist with the initial policy. The agent will anticipate this rigidity, work accordingly, experimentation will collapse in equilibrium. Formally, we have the following result.

Proposition C.1. *For some beliefs (θ_1, θ_2) the principal underexperiments. Conversely, the principal never experiments with a policy in equilibrium when rigidly implementing some policy would yield higher ex-ante expected utility.*

The next result considers an institutional arrangement that can eliminate underexperimentation, and in particular that will result in the principal's "optimal policy experiment" becoming the equilibrium outcome. The phrase "optimal policy experiment" refers to the policy experiment that would yield the highest ex-ante expected utility if the principal could commit to her strategy ex-ante. The result states that creating exogenous "costs to rigidity" can induce the optimal policy experiment when she underexperiments.

Proposition C.2. *For any beliefs (θ_1, θ_2) s.t. the principal underexperiments, there is a cost c of maintaining policy after failure that makes her optimal policy experiment an equilibrium.*

The final result pertains to a model variant in which the principal can first "appoint" a player with different beliefs $\hat{\theta}_1 \in [0, 1]$ to make decisions in her place, and *that* player and the agent will then play the equilibrium that is best for the "original" principal with beliefs θ_1. In particular, we look for conditions under which the principal is strictly better off appointing somebody with beliefs that differ from her own. This yields the following result.

Proposition C.3. *Suppose that a principal with beliefs θ_1 could appoint a player with beliefs $\hat{\theta}_1$ to make policy decisions in her place. If appointing herself is not optimal, then any optimal appointee $\hat{\theta}_1^*$ believes less strongly in the resulting policy x^{*1} than the principal does.*

The intuition here is that the principal would appoint somebody with different beliefs when she would like to commit ex-ante to experiment with some policy x^{*1}, but her beliefs are such that she must rigidly implement x^{*1} whenever she selects it, and so in equilibrium she either rigidly implements x^{*1} or experiments with $\neg x^{*1}$. The optimal appointee will be anybody whose beliefs allow her to experiment with x^{*1}, and this must necessarily be somebody who believes less strongly in it.

Proofs of underexperimentation and commitment

Proof of Proposition C.1 First, there is never *over*experimentation in equilibrium, i.e., the principal never experiments with x^1 when rigidly implementing either x^1 or $\neg x^1$ would be better. The former is immediately ruled out by Lemma 2. The latter also ruled out – if rigidly implementing $\neg x^1$ were optimal with commitment then it must be better than experimenting with $\neg x^1$, and by implication the unique equilibrium of the subgame following $\neg x^1$ without commitment; thus, the principal failing to choose it would be a contradiction.

Next, we there $\exists (\theta_1, \theta_2)$ s.t. underexperimentation occurs, i.e. the principal rigidly implements x^1 when experimenting would be better. In part 3 of Lemma 6 in the Supplemental Proofs we show there exists a nonempty interval of θ_2 s.t. the principal prefers rigidly implementing b to experimenting with a when $\theta_1 = 1 - \bar{\theta}(1-\theta_2)$. By continuity, rigidly implementing b is thus the equilibrium outcome for $\theta_1 = 1 - \bar{\theta}(1-\theta_2) - \varepsilon$ when $\varepsilon > 0$ is sufficiently close to 0, and it is also worse than experimenting with b since the agent's effort drops discretely. Formally,

$$
\begin{aligned}
& V_1^E\left(\bar{\theta}(1-\theta_2), 1-\theta_2\right) = U_1\left(e^E(1-\theta_2), \bar{\theta}(1-\theta_2), 1-\theta_2, 1, 0\right) \\
= \;& U_1\left(e^E(1-\theta_2), \bar{\theta}(1-\theta_2), 1-\theta_2, 1, 1\right) \; > U_1\left(e^R(1-\theta_2), \bar{\theta}(1-\theta_2), 1-\theta_2, 1, 1\right) \\
= \;& V_1^R\left(\bar{\theta}(1-\theta_2), 1-\theta_2\right)
\end{aligned}
$$

The first and last equalities follow from the definitions, the second follows from the definition of $\bar{\theta}(\cdot)$, and the inequality from $U_1(\cdot)$ increasing in e^1 and $e^E(1-\theta_2) > e^R(1-\theta_2)$. ■

Proof of Proposition C.2 If there were an exogenous cost $c > 0$ of maintaining policy after failure, then for experimentation to fail to be an equilibrium of the subgame following policy x requires that the principal prefer to reselect x given initial effort $e^E(P_2(\omega = x))$, the players' resulting posterior beliefs, *and* the cost c. This condition is,

$$h\left(e^E(P_2(\omega = x)), P_1(\omega = x)\right) - \frac{c}{\lambda} > 1 - h\left(e^E(P_2(\omega = x)), P_2(\omega = x)\right)$$

Now suppose the principal underexperiments by rigidly implementing policy $x^{1*} = a$ when $c = 0$. Experimenting must then be the unique equilibrium of the subgame following b, it will remain so with any higher cost $c > 0$, and also rigidly implementing a is better than experimenting with b. So experimenting with a must be the optimal policy experiment, but not an equilibrium of the subgame following a without commitment. However, it will become one when

$$c \geq \lambda\left(h\left(e^E(\theta_2), \theta_1\right) - \left(1 - h\left(e^E(\theta_2), \theta_2\right)\right)\right),$$

and so the principal will select it in equilibrium. A symmetric argument holds when the principal underexperiments by rigidly implementing b. ■

Proof of Proposition C.3 Suppose appointing herself is not optimal; then for a principal with beliefs θ_1, the resulting equilibrium (x^*, s^*) is strictly worse than the equilibrium that would result if an optimal appointee with beliefs $\hat{\theta}_1$ were making policy decisions. Denote this equilibrium $(\hat{x}, \hat{s})$. First note $\hat{s}$ must not an equilibrium experimentation decision the subgame following $\hat{x}$ for θ_1 (since otherwise the principal would choose it). We next argue that $\hat{s} = E$. If $\hat{s} = R$ then by implication experimenting must be an equilibrium of the subgame following $x^1 = \hat{x}$ for θ_1; but then by Lemma 2 it is also strictly better than rigidly implementing $\hat{s}$ and we have a contradiction. Finally, since $(\hat{x}, \hat{s} = E)$ is an equilibrium of the subgame following $\hat{x}$ for $\hat{\theta}_1$ but not θ_1, it follows that

$$\begin{aligned} h\left(e^E\left(P_2\left(\omega=\hat{x}\right)\right), \hat{P}_1\left(\omega=\hat{x}\right)\right) &< 1-h\left(e^E\left(P_2\left(\omega=\hat{x}\right)\right), P_2\left(\omega=\hat{x}\right)\right) \\ &< h\left(e^E\left(P_2\left(\omega=\hat{x}\right)\right), P_1\left(\omega=\hat{x}\right)\right), \end{aligned}$$

implying $\hat{P}_1\left(\omega=\hat{x}\right)<P_1\left(\omega=\hat{x}\right)$. ■

D Supplemental Proofs

In this section we prove a sequence of lemmas that are employed in the previous proofs.

Lemma 1. *Say that the agent's problem is "well-behaved" when* $e^*\left(\theta,\eta_1,\eta_2\right)=\underset{e\in[0,1]}{\arg\max}\left(U\left(e,\theta,\eta_1,\eta_2\right)\right)$ *is unique and* $\in(0,1)$. *The set of* λ *s.t. the agent's problem is well-behaved* $\forall\left(\theta\in(0,1),\eta_1,\eta_2\right)$ *is an interval* $\lambda\in[0,\bar{\lambda})$, *where* $\bar{\lambda}\in\left(\frac{\sqrt{5}-1}{2},\sqrt{3}-1\right)$ *and is* $\approx .68466$.

Proof: First, it is simple to verify that the derivative of the agent's objective function $U\left(e,\theta,\eta_1,\eta_2\right)$ from the proof of Propositions 1 and 2 in the main appendix is:

$$\frac{\partial U\left(e,\theta,\eta_1,\eta_2\right)}{\partial e}=\frac{1}{\lambda}\left(-e+\lambda\theta\left(1+\frac{\lambda}{2}\eta_1\left[k\left(e,\theta\right)+\eta_2\right]\right)\right),$$

where $k\left(e,\theta\right)=\left(1-h\left(e,\theta\right)\right)^2$.

Second, observe $\frac{\partial U}{\partial e}>0$ at $e=0$, and that $\frac{\partial U}{\partial e}$ is convex in e, which follows from the convexity of $k\left(e,\theta\right)$. The set of maximizers is thus either a singleton in the interior (the first and possibly only point where the FOC is satisfied), a singleton on the boundary $e=1$, or a pair where one of the two elements is $e=1$. This further implies that whenever the problem is not "well behaved," $e=1$ is a maximizer.

Third, observe that $\frac{\partial^2 U}{\partial e\partial\lambda}=\frac{e}{\lambda^2}+\frac{\theta}{2}\eta_1\left(k\left(e,\theta\right)+\eta_2\right)>0$. Thus by Milgrom and Shannon (1994) the set of maximizers of $U\left(\cdot\right)$ is weakly increasing in λ. This implies that the set of λ s.t. the problem is well behaved for a given $\left(\eta_1,\eta_2,\theta\right)$ is an interval; if it were well behaved for λ' but not $\lambda'' < \lambda'$, then $e = 1$ would be a maximizer for the former but not the latter, contradicting weak set increasingness. Lastly, this implies that the set of λ that are well behaved for *all* feasible parameters $\forall \left(\theta \in (0,1), \eta_1, \eta_2\right)$ is also an interval $[0, \bar{\lambda})$; if it were not then it would also not an interval for some specific profile of parameters (η_1, η_2, θ), a contradiction.

We can bound $\bar{\lambda}$ below and above analytically, and computationally compute an estimated value. We must have $\bar{\lambda} > \frac{\sqrt{5}-1}{2}$ since the problem is strictly concave for all feasible parameters (and by implication well behaved) when $\frac{\partial U}{\partial e}\Big|_{e=1} = -\frac{1}{\lambda} + \theta\left(1 + \frac{\lambda}{2}\eta_1\left(1+\eta_2\right)\right) \leq -\frac{1}{\lambda} + (1+\lambda) < 0$, which holds i.f.f. $\lambda < \frac{\sqrt{5}-1}{2}$. We also must have $\bar{\lambda} < \sqrt{3}-1$ since best-response effort at $\theta = 1$ when the principal experiments is $e^*(1,1,1) = \lambda\left(1+\frac{\lambda}{2}\right) < 1$ i.f.f. $\lambda < \sqrt{3}-1$. In the supplemental mathematica code to this document, we verify that $\bar{\lambda} \approx .68466$. Since the set of maximizers is weakly increasing in η_2 (by $\frac{\partial^2 U}{\partial e \partial \eta_2} = \lambda\theta > 0$), to find $\bar{\lambda}$ it suffices to check that the problem is well-behaved $\forall\theta$ when $\eta_2 = 1$. We thus identify $\bar{\lambda}$ by compute the highest λ s.t. $\forall\theta \in [0,1]$, the agent's utility at the lowest solution to the first-order condition is greater than his utility from $e = 1$. ■

Lemma 2. *The following two statements hold when the principal can play mixed strategies.*

(i) Whenever experimenting is an equilibrium of the subgame following initial policy $x^1 \in \{a, b\}$, then it is also the optimal (pure or mixed) strategy for the principal if she could precommit to her responses to success and failure.

(ii) Whenever experimenting with $x^1 = a$ is not an equilibrium of the subgame following $x^1 \in \{a, b\}$, then the unique equilibrium is rigid implementation.

(Part 1) Because of symmetry we can restrict attention to the subgame following $x^1 = a$. We first must characterize the agent's best response effort $\arg\max_{e\in[0,1]} \{U_2(e, \theta_2, p_s, p_f)\}$ to a general mixed strategy by the principal in the Game with Learning as characterized in

Appendix A. Taking the derivative of equation (A.1) w.r.t. e yields:

$$\frac{\partial U_2\left(e,\theta,p_s,p_f\right)}{\partial e}=\frac{1}{\lambda}\left(-e+\lambda\theta\left(1+\frac{\lambda}{2}\left[k\left(e,\theta\right)+\left(p_s-p_f\right)\right]\right)\right).$$

Now, it is easily verified that $\frac{\partial U_2\left(e,\theta,p_s,p_f\right)}{\partial e}=\frac{\partial U\left(e,\theta,1,p_s-p_f\right)}{\partial e}$, where $U\left(\cdot\right)$ is the form of the agent's objective function that we used in the Main Appendix, which amalgamated payoffs from the NLB and the GWL, but did not account for principal mixed strategies. This means we can "piggyback" off of the analysis of that problem in Lemma 1. First, it remains true that the agent's problem is well behaved $\forall\left(\theta,p_s,p_f\right)$ when $\lambda<\bar{\lambda}$. Second, when $\lambda<\bar{\lambda}$ we have that (i) $\arg\max_{e\in[0,1]}\left\{U_2\left(e,\theta_2,p_s,p_f\right)\right\}=e^*\left(\theta,1,p_s-p_f\right)$ where $e^*\left(\theta,\eta_1,\eta_2\right)$ is the maximizer of $U\left(\theta,\eta_1,\eta_2\right)$, and (ii) $e^*\left(\theta,1,p_s-p_f\right)$ is strictly increasing in p_s and strictly decreasing in p_f by $\frac{\partial^2 U}{\partial e\partial\eta_2}=\lambda\theta>0$ and Theorem 1 of Edlin and Shannon (1998). For notational simplicity, for the remainder of the proof we will write $e^*\left(\theta,1,p_s-p_f\right)$ as $\hat{e}\left(\theta,p_s-p_f\right)$ so as not to carry around unecessary terms.

(Part 2) We now prove (i). Employing the characterization in Appendix A, the principal's utility from choosing $x^1=a$ if she could precommit to her responses to success and failure $\left(p_s,p_f\right)$ would be $U_1\left(\hat{e}\left(\theta_2,p_s-p_f\right),\theta_1,\theta_2,p_s,p_f\right)$. It is easily verified that $U_1\left(\cdot\right)$ is strictly increasing in p_s holding first period effort e fixed, and increasing in e when $p_s=1$. Also recall that $\hat{e}\left(\theta_2,p_s-p_f\right)$ is increasing in p_s. Hence,

$$U_1\left(\hat{e}\left(\theta_2,p_s-p_f\right),\theta_1,\theta_2,p_s,p_f\right)<U_1\left(\hat{e}\left(\theta_2,p_s-p_f\right),\theta_1,\theta_2,1,p_f\right)<U_1\left(\hat{e}\left(\theta_2,1-p_f\right),\theta_1,\theta_2,1,p_f\right)$$

and the optimal choice after success is therefore to always stay, i.e. $p_s^*=1$. This feature is shared with the baseline model without commitment. Intuitively, the reason is that a higher probability of staying with the initial policy after success is both interim-better for the principal, and also better motivates the agent ex-ante.

Given the above analysis, the principal's optimal choice after policy failure satisfies $p_f^*\in$

$\arg\max_{p_f \in [0,1]} \{U_1(\hat{e}(\theta_2, 1-p_f), \theta_1, \theta_2, 1, p_f)\}$. Now her utility $U_1(\cdot)$ can be rewritten as,

$$\begin{aligned}&\theta_1 e \cdot (1 + \lambda p_s)\\&+ (1-\theta_1 e) \cdot \lambda (p_f \cdot h(e,\theta_1) h(e,\theta_2) + (1-p_f) \cdot (1-h(e,\theta_1))(1-h(e,\theta_2))),\end{aligned}$$

and it is easily verified that this is decreasing in p_f whenever $h(e,\theta_1) \leq 1 - h(e,\theta_2)$, i.e., if posteriors after failure are s.t. it is better to switch. If experimenting is an equilibrium of the subgame $x^1 = a$ then by definition this property holds for $e = \hat{e}(\theta_2, 1)$ (see equilibrium conditions in Appendix A). In addition, recall that $U_1(\cdot)$ is increasing in e when $p_s = 1$ and that $\hat{e}(\theta_2, 1-p_f)$ is decreasing in p_f. Combining these observations yields,

$$U_1(\hat{e}(\theta_2,1), \theta_1, \theta_2, 1, 0) > U_1(\hat{e}(\theta_2,1), \theta_1, \theta_2, 1, p_f) > U_1(\hat{e}(\theta_2, 1-p_f), \theta_1, \theta_2, 1, p_f)$$

whenever experimenting is an equilibrium. Consequently $\left(p_s^* = 1, p_f^* = 0\right)$, i.e. experimenting, is the optimal strategy with commitment.

(Part 3) We now prove (ii). If experimentation is not an equilibrium of the subgame following $x^1 = a$, then by the equilibrium characterization in Appendix A, $h\left(\hat{e}\left(\theta_2, 1-p_f^a\right), \theta_1\right) > 1 - h\left(\hat{e}\left(\theta_2, 1-p_f^a\right), \theta_2\right)$ when $p_f^a = 0$. Another equilibrium with $p_f^a > 0$ would require that the l.h.s. $\leq$ r.h.s. – but this cannot be since $\hat{e}\left(\theta_2, 1-p_f^a\right)$ is decreasing in p_f^a, so the l.h.s. is increasing and the r.h.s. is decreasing. ■

Lemma 3 (Preference Monotonicity)**.** *For all* $(s, s') \in \{R, E\}^2$,

$$V_1^s\left(\hat{\theta}_1, \theta_2\right) > V_1^{s'}\left(1-\hat{\theta}_1, 1-\theta_2\right) \to V_1^s(\theta_1, \theta_2) > V_1^{s'}(1-\theta_1, 1-\theta_2) \text{ for all } \theta_1 > \hat{\theta}_1.$$

In words, fixing the principal's experimentation decisions down each path of play, if she prefers $x^1 = a$ given beliefs $\hat{\theta}_1$ then she also prefers it for any higher belief.

Proof: Because the principal's expected utility for each first period policy is linear in her prior beliefs (holding her future experimentation decisions fixed), a nonmonotonicity would imply that $x^1 = b$ is better when $\omega = a$ and $x^1 = a$ is better when $\omega = b$. The former could not be true if she is rigidly implementing b since it would always fail, and the latter could not be true if she is rigidly implementing a. Thus, a nonmonotonicity requires that she be experimenting down both paths of play. To rule it out, it therefore suffices to show that experimenting with b is better than experimenting with a when $\omega = b$ (given the agent is predisposed to b, i.e. $\theta_2 \le \frac{1}{2}$). This is,

$$(1+\lambda)\, e^E(1-\theta_2) > \lambda\left(1 - h\left(e^E(\theta_2), \theta_2\right)\right)$$

Applying the definition from Appendix A, we know $e^E(1-\theta) > \lambda(1-\theta) \rightarrow$ the l.h.s. is $> (1+\lambda)\lambda(1-\theta_2)$. Also, $e^E(\theta_2) < \lambda\theta_2(1+\lambda) \rightarrow$ the r.h.s. $< \frac{\lambda(1-\theta_2)}{1-\theta_2^2\lambda(1+\lambda)}$. The above inequality will thus hold when

$$(1+\lambda)\lambda(1-\theta_2) > \frac{\lambda(1-\theta_2)}{1-\theta_2^2\lambda(1+\lambda)} \iff (1+\lambda)^2 < \frac{1}{\theta_2^2}.$$

The l.h.s. is $<$ the r.h.s. $\forall \theta_2 \le \frac{1}{2}$ when $\lambda < 1$, which always holds by assumption since $\lambda < \bar{\lambda} < 1$.■

Lemma 4. *The principal strictly prefers experimenting with b to experimenting with a when $\theta_1 = 1-\theta_2$ and $\theta_2 < \frac{1}{2}$.*

Proof: Let $\phi(\theta_2) = \lambda\theta_2\left(1 + \frac{\lambda}{2}\left(k\left(e^E(\theta_2), \theta_2\right) + 1\right)\right)(1+\lambda) - \lambda\left(1 - h\left(e^E(1-\theta_2), 1-\theta_2\right)\right)$; this is the principal's utility difference between experimenting with a and experimenting with b when $\omega = a$. Applying symmetry, her expected utility difference between experimenting with a and experimenting with b with prior θ_1 is $\theta_1\phi(\theta_2) - (1-\theta_1)\phi(1-\theta_2)$. We now wish to show that this is < 0 when $\theta_1 = 1-\theta_2$ given that $\theta_2 < \frac{1}{2}$, i.e.

$$(1-\theta_2)\phi(\theta_2) - \theta_2\phi(1-\theta_2) < 0 \iff \frac{\phi(\theta_2)}{\lambda\theta_2} - \frac{\phi(1-\theta_2)}{\lambda(1-\theta_2)} < 0 \text{ when } \theta_2 < \frac{1}{2}.$$

It is simple to verify that

$$\frac{\phi\left(\theta_2\right)}{\lambda\theta_2}=\left(1+\frac{\lambda}{2}\left(k\left(e^E\left(\theta_2\right),\theta_2\right)+1\right)\right)(1+\lambda)-\frac{1}{1-e^E\left(1-\theta_2\right)\cdot\left(1-\theta_2\right)},$$

and by symmetry it suffices to show that $\frac{\phi(\theta_2)}{\lambda\theta_2}-\frac{\phi(1-\theta_2)}{\lambda(1-\theta_2)}>0$ when $\theta_2>\frac{1}{2}$. Using substitution and rearranging, we then have that $\frac{\phi(\theta_2)}{\lambda\theta_2}-\frac{\phi(1-\theta_2)}{\lambda(1-\theta_2)}>0 \iff$

$$\frac{\theta_2 e^E\left(\theta_2\right)-\left(1-\theta_2\right)e^E\left(1-\theta_2\right)}{\left(1-\theta_2 e^E\left(\theta_2\right)\right)\cdot\left(1-\left(1-\theta_2\right)e^E\left(1-\theta_2\right)\right)}>\frac{\lambda}{2}(1+\lambda)\left(k\left(e^E\left(1-\theta_2\right),1-\theta_2\right)-k\left(e^E\left(\theta_2\right),\theta_2\right)\right) \tag{D.1}$$

Since the denominator of the l.h.s. is <1, and $1+\lambda<2$, the above inequality holds if the following yet stronger inequality holds,

$$\theta_2 e^E\left(\theta_2\right)-\left(1-\theta_2\right)e^E\left(1-\theta_2\right)>\lambda\left(k\left(e^E\left(1-\theta_2\right),1-\theta_2\right)-k\left(e^E\left(\theta_2\right),\theta_2\right)\right). \tag{D.2}$$

Now again substituting in the definition of $e^E(\cdot)$, the l.h.s. can be rewritten

$$\lambda\left(\left(2\theta_2-1\right)\left(1+\frac{\lambda}{2}\right)+\frac{\lambda}{2}\left(\theta_2^2 k\left(e^E\left(\theta_2\right),\theta_2\right)-\left(1-\theta_2\right)^2 k\left(e^E\left(1-\theta_2\right),1-\theta_2\right)\right)\right),$$

implying that the desired inequality holds i.f.f.

$$\left(2\theta_2-1\right)\left(1+\frac{\lambda}{2}\right)>\left(1+\frac{\lambda}{2}\left(1-\theta_2\right)^2\right)k\left(e^E\left(1-\theta_2\right),1-\theta_2\right)-\left(1+\frac{\lambda}{2}\theta_2^2\right)k\left(e^E\left(\theta_2\right),\theta_2\right)$$

It is easily verified that the r.h.s. $<\left(1+\frac{\lambda}{2}\right)\left(k\left(e^E\left(1-\theta_2\right),1-\theta_2\right)-k\left(e^E\left(\theta_2\right),\theta_2\right)\right)$ when $\theta_2>\frac{1}{2}$, implying that the above inequality holds if the following yet stronger inequality holds:

$$2\theta_2-1>k\left(e^E\left(1-\theta_2\right),1-\theta_2\right)-k\left(e^E\left(\theta_2\right),\theta_2\right) \tag{D.3}$$

We now show that eqn. $(D.3)$ holds. Substituting in the definition of $k(\cdot)$, observing that

$2\theta_2 - 1 = \theta_2^2 - (1-\theta_2^2)$, and rearranging, we have that the inequality is equivalent to

$$\theta_2^2 \left(1 - \frac{1}{\left(1-(1-\theta_2)\, e^E\,(1-\theta_2)\right)^2}\right) > (1-\theta_2)^2 \left(1 - \frac{1}{\left(1-\theta_2 e^E\,(\theta_2)\right)^2}\right)$$

Since $e^E(\theta_2)$ is increasing in θ_2, this clearly holds because $(1-\theta_2)\, e^E(1-\theta_2) < \theta_2 e^E(\theta_2)$ when $\theta_2 > \frac{1}{2}$. The desired property is hence shown. ■

Lemma 5. *For all θ_2 s.t. $1-\bar{\theta}(1-\theta_2) > \frac{1}{2}$, the principal weakly prefers rigidly implementing b to experimenting with a when $\theta_1 = 1-\bar{\theta}(1-\theta_2)$.*

Proof: Observe that $1-\bar{\theta}(1-\theta_2)$ is $=1$ at $\theta_2 = 0$, less than $\frac{1}{2}$ at $\theta_2 = \frac{1}{2}$, and strictly decreasing in θ_2. Thus, there exists some unique $\hat{\theta}_2$ satisfying $1-\bar{\theta}\left(1-\hat{\theta}_2\right) = \frac{1}{2}$ s.t. the set of beliefs $\left[\frac{1}{2}, 1-\bar{\theta}(1-\theta_2)\right]$ where the principal would rigidly implement b in that subgame is nonempty if and only if $\theta_2 < \hat{\theta}_2$. Note that $\hat{\theta}_2$ is a function of λ so we henceforth write $\hat{\theta}_2(\lambda)$ for clarity. Now applying the definitions and rearranging, we wish to show that

$$\forall \lambda \in \left[\underline{\lambda}, \bar{\lambda}\right],$$
$$\bar{\theta}(1-\theta_2) \cdot \left(V_1^R(1, 1-\theta_2) - V_1^E(0,\theta_2)\right) > \left(1-\bar{\theta}(1-\theta_2)\right) \cdot V_1^E(1,\theta_2) \;\; \forall \theta_2 \in \left[0, \hat{\theta}_2(\lambda)\right],$$

where $\underline{\lambda} \approx .23505$. We verify this step in the supplemental mathematical code. ■

Lemma 6. *The following three properties hold.*

1. *When $\lambda > \hat{\lambda}$, $V_1^E(0, 1-\theta_2) - V_1^R(1,\theta_2) > 0$ for θ_2 in a nonempty interval $[0,\varepsilon_1]$.*
2. *When $\lambda \in [0,1]$, $V_1^R(1,\theta_2) - V_1^E(1,\theta_2) > 0$ for θ_2 in a nonempty interval $[0,\varepsilon_2]$.*
3. *When $\lambda \in [0,1]$, $V_1^R\left(\bar{\theta}(1-\theta_2), 1-\theta_2\right) - V_1^E\left(1-\bar{\theta}(1-\theta_2), \theta_2\right) > 0$ for θ_2 in a nonempty interval $[0,\varepsilon_3]$.*

To show that some function $f(\theta_2)$ satisfying $f(0)=0$ is >0 for $\theta_2 \in (0,\varepsilon)$ where $\varepsilon>0$, it suffices to show by continuity that $\left.\frac{f(\theta_2)}{\theta_2}\right|_{\theta_2=0} > 0$ provided that this quantity is finite. We now show this for each of the desired expressions.

Property 1: When $\omega = a$, the principal's utility from experimenting with b is,

$$\lambda\left(1-h\left(e^E(1-\theta_2),1-\theta_2\right)\right) = \frac{\lambda\theta_2}{1-(1-\theta_2)\,e^E(1-\theta_2)},$$

and from rigidly implementing a is,

$$e^R(\theta_2)(1+\lambda)+\left(1-e^R(\theta_2)\right)h\left(e^R(\theta_2),\theta_2\right) = \lambda\theta_2\left((1+\lambda)\left(1+\frac{\lambda}{2}k\left(e^R(\theta_2),\theta_2\right)\right)+\frac{\left(1-e^R(\theta_2)\right)^2}{1-\theta_2 e^R(\theta_2)}\right)$$

Since $e^R(0)=0$, $e^E(1)=\lambda\left(1+\frac{\lambda}{2}\right)$, and $k\left(e^R(0),0\right)=1$,

$$\left.\frac{1}{\lambda\theta_2}\left(V_1^E(0,1-\theta_2)-V_1^R(1,\theta_2)\right)\right|_{\theta_2=0} = \left(\frac{1}{1-\lambda\left(1+\frac{\lambda}{2}\right)}\right)-\left((1+\lambda)\left(1+\frac{\lambda}{2}\right)+1\right).$$

Manipulating the above expression demonstrates that it is ≥ 0 i.f.f.

$$\lambda(1+\lambda)\left(1+\frac{\lambda}{2}\right) > 1. \tag{D.4}$$

which holds i.f.f. $\lambda > \hat{\lambda}$ by definition.

Property 2: When $\omega = a$, the principal's utility from experimenting with a is,

$$e^E(\theta_2)(1+\lambda) = \lambda\theta_2\left(1+\frac{\lambda}{2}\left(k\left(e^E(\theta_2),\theta_2\right)+1\right)\right)(1+\lambda),$$

Since $k\left(e^E(0),0\right)=1$ and applying Part 1, the desired condition is equivalent to

$$(1+\lambda)\left(1+\frac{\lambda}{2}\right)+1 > (1+\lambda)^2$$

which holds if $\lambda(1+\lambda) < 2 \iff \lambda < 1$.

Property 3: Using the definition in equation (A.3), the threshold function $\bar{\theta}(\theta_2)$ in closed form is $\bar{\theta}(\theta_2) = \frac{1-\theta_2}{(1-\theta_2)+\theta_2(1-e^E(\theta_2))^2}$. Thus, when $\theta_1 = 1 - \bar{\theta}(1-\theta_2)$ the difference in the principal's utility between rigidly implementing b and experimenting with a is,

$$\bar{\theta}(1-\theta_2) \cdot \left(V_1^R(1, 1-\theta_2) - V_1^E(0, \theta_2)\right) - \left(1 - \bar{\theta}(1-\theta_2)\right) \cdot V_1^E(1, \theta_2) \tag{D.5}$$

The first term is the subjective probability that $\omega = b$ times the utility difference between rigidly implementing b and experimenting with a when $\omega = b$. The second term is the subjective probability that $\omega = a$ times the utility of experimenting with a when $\omega = a$ (the payoff from rigidly implementing b when the state is a is 0).

Now to get the desired expression we divide through by θ_2 and then evaluate at $\theta_2 = 0$ – we do so by dividing $\bar{\theta}(1-\theta_2)$ by θ_2 in the first term, then $V_1^E(1, \theta_2)$ by θ_2 in the second term, and then evaluating all expressions at $\theta_2 = 0$. First,

$$\left(\frac{1}{\theta_2}\right) \cdot \bar{\theta}(1-\theta_2) = \frac{1}{\theta_2 + (1-\theta_2)\left(1 - e^E(1-\theta_2)\right)^2}$$

which is equal to $\left(1 - \lambda\left(1 + \frac{\lambda}{2}\right)\right)^{-2}$ at $\theta_2 = 0$. Second, $V_1^R(1, 1-\theta_2) = 2\lambda$, $V_1^R(0, \theta_2) = \lambda$, and $1 - \bar{\theta}(1-\theta_2) = 1$ at $\theta_2 = 0$. Third, from the proof of property 2 we know $\left(\frac{1}{\theta_2}\right) \cdot V_1^E(1, \theta_2) = \lambda\left(1 + \frac{\lambda}{2}\left(k\left(e^E(\theta_2), \theta_2\right) + 1\right)\right)(1+\lambda)$ which is $= \lambda(1+\lambda)^2$ evaluated at $\theta_2 = 0$. Assembling these observations, the desired expression is

$$\left(1 - \lambda\left(1 + \frac{\lambda}{2}\right)\right)^{-2} \cdot (2\lambda - \lambda) - \lambda(1+\lambda)^2 > 0 \iff \left(1 - \lambda\left(1 + \frac{\lambda}{2}\right)\right)^{2} \cdot (1+\lambda)^2 < 1$$

Now the l.h.s. is $< (1-\lambda)^2 \cdot (1+\lambda)^2 = (1-\lambda^2)^2 \leq 1 \, \forall \lambda \in [0,1]$, so the result is shown. ■

E Preferences vs. Beliefs

Suppose the state $\omega \in \{a, b\}$ is payoff-irrelevant for the agent, and he shares the principal's prior θ_1 over $P(\omega = a)$. Instead, he simply has his own return to effort of $\pi_x e$ when working on each policy $x \in \{a, b\}$, and $\pi_b > \pi_a$.

In the second period, his objective function from exerting second period effort e on some policy $x \in \{a, b\}$ is simply $\pi_x e - \frac{e^2}{2\lambda}$ regardless of the history. It is straightforward to show that his optimal level of effort is $\lambda \pi_x$, and his expected payoff is $\frac{\lambda}{2}\pi_x^2$.

Now consider the first period. Deriving the agents expected payoff from exerting first period effort e when the principal rigidly implement policy a is immediate. His utility when the principal experiments with a is:

$$-\frac{e^2}{2\lambda} + \pi_a e + \theta_1 \left(e \left(\frac{\lambda}{2}\pi_a^2 \right) + \left(1 - e^1\right) \left(\frac{\lambda}{2}\pi_b^2 \right) \right) + (1 - \theta_2) \left(\frac{\lambda}{2}\pi_b^2 \right),$$

which easily reduces to the expression in the main text.

Both objective functions are strictly concave, so the unique solution is given by the first-order condition. The agent's effort on a given rigid implementation is simply $\lambda \pi_a$. If the principal experiments then the derivative of the objective function is

$$-\frac{e}{\lambda} + \pi_a - \frac{\lambda}{2}\theta_1 \left(\pi_b^2 - \pi_a^2\right),$$

so optimal effort is $\max\left\{\lambda\pi_a - \frac{\lambda^2}{2}\theta_1\left(\pi_b^2 - \pi_a^2\right), 0\right\}$ and is strictly lower as stated in the text.

Finally, by symmetry the agent's effort if the principal rigidly implements b is $\lambda\pi_b$, and if she experiments with b is $\min\left\{\lambda\pi_b + \frac{\lambda^2}{2}(1 - \theta_1)\left(\pi_b^2 - \pi_a^2\right), 1\right\}$. The agent therefore works harder on a policy experiment with the policy b that he prefers; the "threat" of switching to his less preferred policy a after failure motivates him, as intuition would suggest.

F Supplemental Mathematica Code

In[25]:=
```
(*SUPPLEMENTAL MATHEMATICA CODE TO
 "EXPERIMENTATION AND PERSUASION IN POLITICAL ORGANIZATIONS", 02-21-2015*)



(*-------------------------------------BEGIN
 CODE----------------------------------*)
```

In[26]:=
```
(*----------Threshold Parameters------------------------*)
(*-----------------------------------------------------*)
```

In[27]:=
```
lLbar := .23505
lbar := .684660
lfig := .653
lhat = l /. Solve[l * (1 + l) * (1 + (l / 2)) == 1, l][[1]]
```

Out[30]=
$$-1+\frac{1}{3}\left(27-3\sqrt{78}\right)^{1/3}+\frac{\left(9+\sqrt{78}\right)^{1/3}}{3^{2/3}}$$

In[31]:=
```
(*----------Posterior Beliefs After Failure-------------*)
(*-----------------------------------------------------*)
h[e_, p_] := (1 - e) * p / (1 - p * e)
k[e_, p_] := (1 - h[e, p]) ^ 2
```

In[33]:=

In[34]:=
```
(*---------------The Agent's Problem--------------------*)
(*-----------------------------------------------------*)

(*objective function*)
U2[e_, t2_, pS_, pF_, l_] := (t2 * e - e ^ 2 / (2 * l)) +
   (l / 2) * (pF * t2 * (e + (1 - e) * h[e, t2]) + (1 - pF) * (1 - t2) * (1 - h[e, t2])) +
   (pS - pF) * (l / 2) * t2 * e

(*optima*)
eStar[t2_, d_, l_] := Root[(#1 - l * t2 * (1 + (l / 2) * (k[#1, t2] + d))) &, 1]
eR[t2_, l_] := eStar[t2, 0, l]
eE[t2_, l_] := eStar[t2, 1, l]
```

In[38]:=

```
(*---------------The Principal's Problem--------------------*)
(*-------------------------------------------------------*)

(*Equation 10*)
U1[e_, t1_, t2_, pS_, pF_, l_] :=
 t1 * (e + l * (e * pS + (1 - e) * pF * h[e, t2])) + (1 - t1) * (1 - pF) * l * (1 - h[e, t2])

(*Pure Strategy Payoffs with agent best-responses - from Proposition 3*)
V1R[t1_, t2_, l_] := U1[eR[t2, l], t1, t2, 1, 1, l]
V1E[t1_, t2_, l_] := U1[eE[t2, l], t1, t2, 1, 0, l]

(*Threshold for experimentation for a given e1*)
thetaBar[t2_, e_] := (1 - t2) / (1 - 2 e * t2 + e^2 * t2)

(*Indifference Belief Thresholds for Principal*)
bEaR[t2_, l_] := t1 /. Solve[V1E[1 - t1, 1 - t2, l] == V1R[t1, t2, l], t1][[1]]
bEaE[t2_, l_] := t1 /. Solve[V1E[1 - t1, 1 - t2, l] == V1E[t1, t2, l], t1][[1]]
bRaE[t2_, l_] := t1 /. Solve[V1R[1 - t1, 1 - t2, l] == V1E[t1, t2, l], t1][[1]]
aEaR[t2_, l_] := If[V1E[1, t2, l] ≥ V1R[1, t2, l],
  1, t1 /. Solve[V1E[t1, t2, l] == V1R[t1, t2, l], t1][[1]]]



(*-------------------------------------END
 CODE----------------------------------*)


(*-------------------------------------BEGIN
 FIGURES----------------------------------*)


(*-------------PLOTS FOR FIGURE 2------------------*)
Manipulate[Plot[{l * t2, l * (1 - t2), eR[t2, l], eR[1 - t2, l], eE[t2, l], eE[1 - t2, l]},
  {t2, 0, .5}, PlotRange → {0, 1}], {l, lfig, lfig}]
```

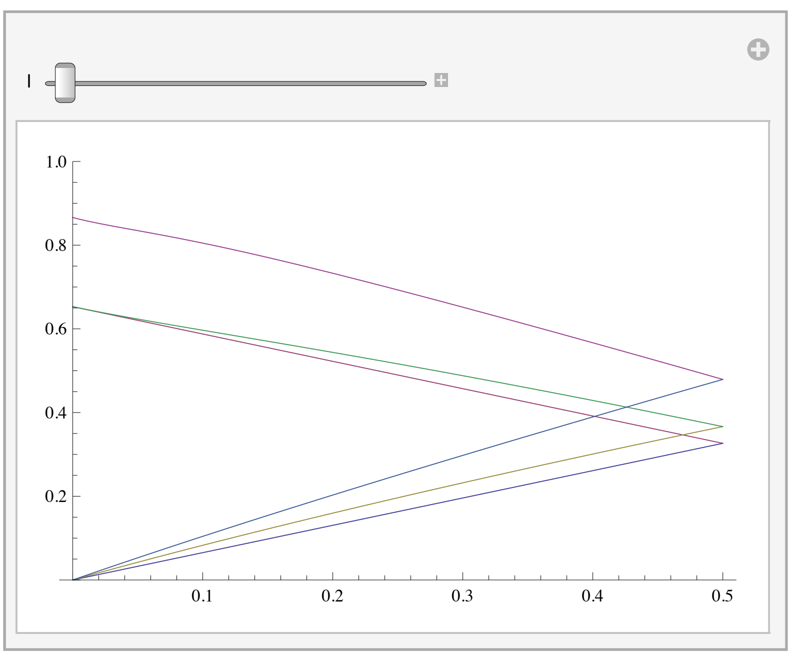

In[46]:=

```
(*-------------PLOT FOR FIGURE 3------------------*)

Manipulate[
 Show[

  Plot[aEaR[t2, l], {t2, 0, 1},
   PlotRange → {0, 1}, PlotStyle → Directive[Blue, Thick]],

  Plot[bEaR[t2, l], {t2, 0, 1},
   PlotRange → {0, 1}, PlotStyle → Directive[Red, Thick]],
  Plot[thetaBar[t2, eE[t2, l]], {t2, 0, 1}, PlotRange → {0, 1},
   PlotStyle → Directive[Black, Dashed]],
  Plot[{1 - t2, .5, 1, 1000 * (t2 - .5)}, {t2, 0, 1}, PlotRange → {0, 1},
   PlotStyle → Directive[Gray, Dashed]],
  Plot[bEaE[t2, l], {t2, 0, 1}, PlotRange → {0, 1},
   PlotStyle → Directive[Orange, Thick]],

  RegionPlot[t1 < 1 - t2, {t2, 0, 1}, {t1, 0, 1}, PlotStyle → Gray]
 ], {l, lfig, lfig}]
```

Out[46]=

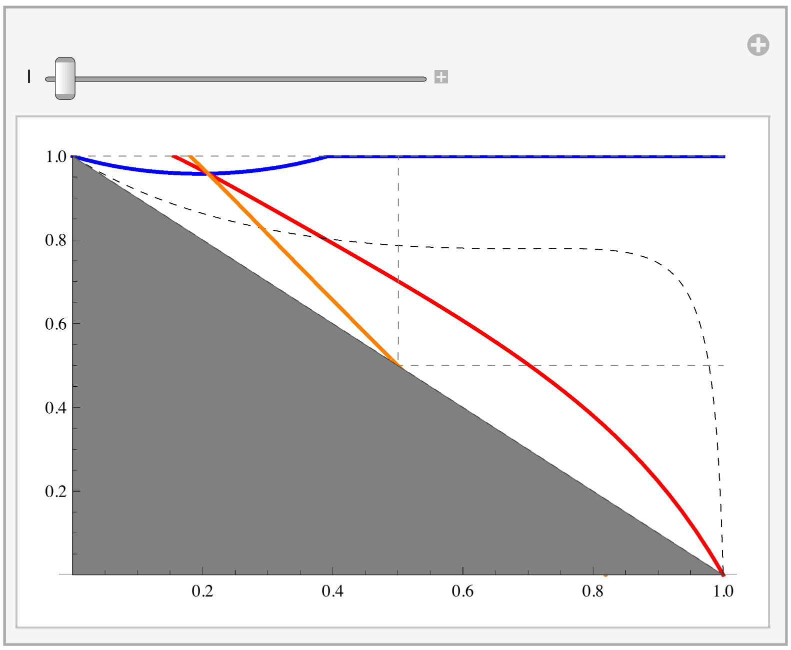


```
(*-----------------------------------END
 FIGURES--------------------------------*)



(*-----------------------------------BEGIN
 PROOF SUPPLEMENTS--------------------------------*)
```

```
(*--------For Proof of Lemma 1 (Good Behavior)-----------------*)
(*The following plots verify that the agent's utility for the interior
 solution exceeds his utility at boundary solution for all t2 \in [0,1]
 when \lambda \approx .68466. The first plot illustrates the entire
 parameter space. The second plot focuses on the range t2 \in [.92,.96]*)

Manipulate[Plot[{U2[eE[t2, l], t2, 1, 0, l] - U2[1, t2, 1, 0, l]},
  {t2, 0, 1}, PlotRange → {-.1, .8}], {l, .684660, 0}]
Manipulate[Plot[{1000 * (U2[eE[t2, l], t2, 1, 0, l] - U2[1, t2, 1, 0, l])},
  {t2, .92, .96}, PlotRange → {0, .7}], {l, .684660, 0}]
```

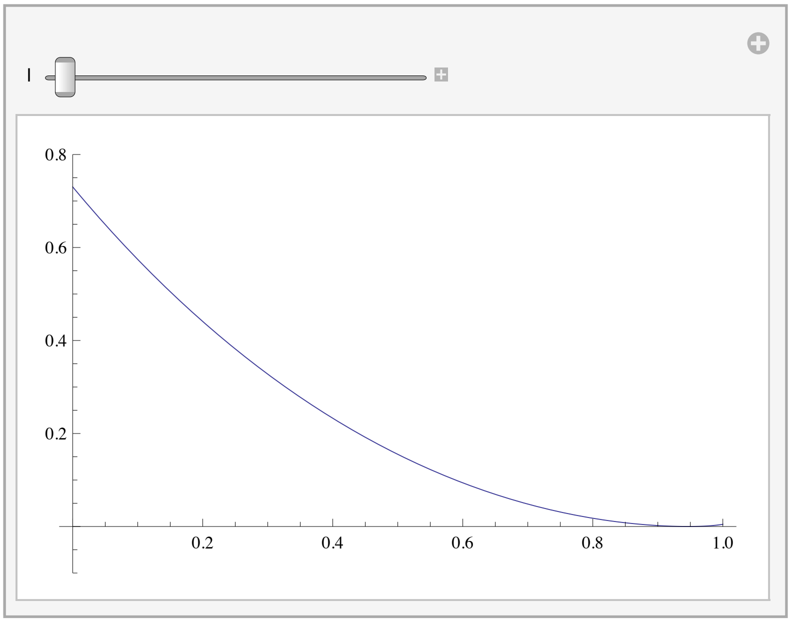

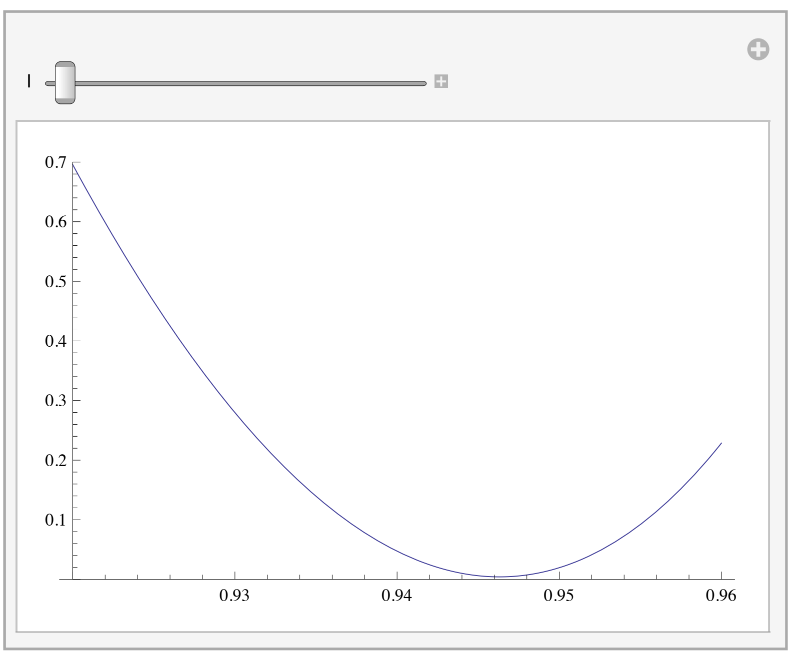

In[47]:=

```
(*-------------For Proof of Lemma 4 -----------------*)

t2hat[l_] := t2 /. NSolve[thetaBar[1 - t2, eE[1 - t2, l]] == 1 / 2, t2][[1]]

(*To verify the condition, observe that the blue line is higher
 than the purple line for all \theta_2 to the left of \hat{\theta}_2
 (\lambda) for any value of \lambda in [lUbar,lbar] in the slider *)
Manipulate[Show[
  Plot[{
    thetaBar[1 - t2, eE[1 - t2, l]] * (V1R[1, 1 - t2, l] - V1E[0, t2, l]),
    (1 - thetaBar[1 - t2, eE[1 - t2, l]]) * V1E[1, t2, l]},
   {t2, 0, .5}, PlotStyle → {Blue, Purple}],
  Plot[Evaluate[10 000 * (t2 - t2hat[l])], {t2, 0, .5}, PlotStyle → {Green}]
 ]
 , {l, lLbar, lbar}]
```

Out[48]=

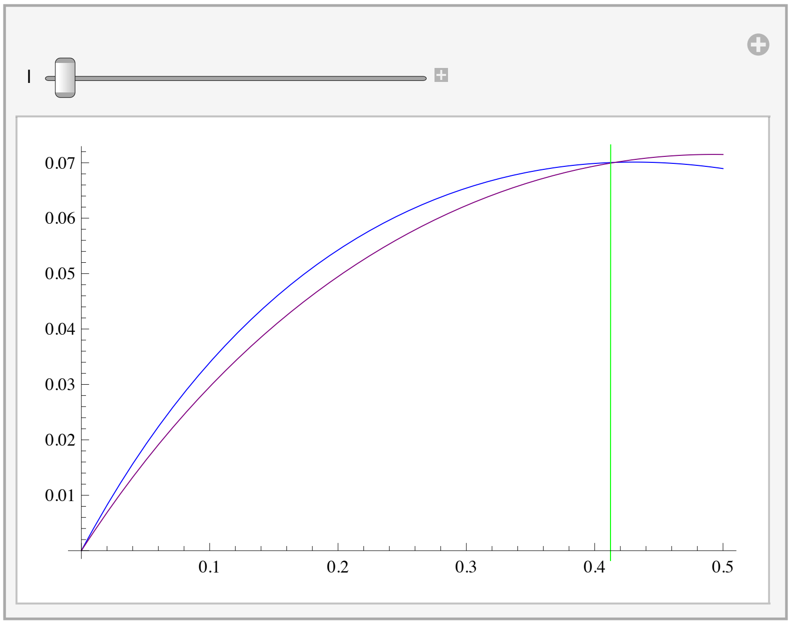

```
(*-------------For Proof of Proposition 3, Part 2 -----------------*)
(*To verify the condition,
observe that the function is <0 for for all \theta_2 \in [0,1]
  and any value of \lambda \in [0,lbar] in the slider *)

Manipulate[Plot[{V1E[0, 1 - t2, l] / (t2 * l) - V1R[1, t2, l] / (t2 * l)},
  {t2, 0, 1}, PlotRange → {-1.5, 0}], {l, .001, lhat}]
```

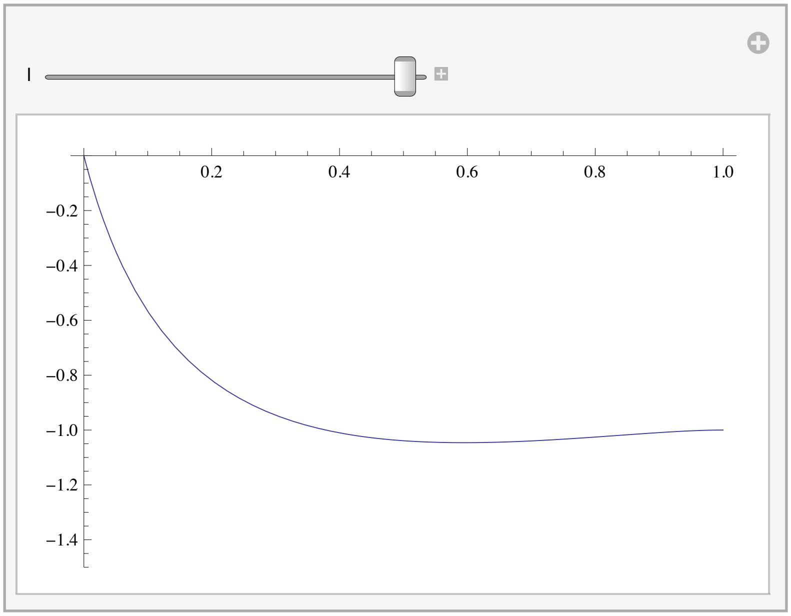


```
(*------------------------------------END
 PROOF SUPPLEMENTS----------------------------------*)
```