Prolonged Learning and Hasty Stopping: the Wald Problem with Ambiguity

Sarah Auster, Yeon-Koo Che and Konrad Mierendorff *

—Online Appendix—

B Additional Results and Proofs

B.1 Bayesian optimal behavior

Proposition B.1. *The Bayesian DM takes action ℓ if $p_t \leq p_\ell^B$, takes action r if $p_t \geq p_r^B$, and experiments if $p_t \in (p_\ell^B, p_r^B)$, where the threshold $p_\ell^B > 0$ maximizes the value (3) and $p_r^B \in [p_\ell^B, 1)$ equalizes the two terms in (4). The experimentation region is non-empty, i.e. $p_\ell^B < p_r^B$, if $c < \bar{c}$, where $\bar{c}$ is defined in Section I.*

Proof. In order to describe the Bayesian optimal stopping rule, consider the DM's value from experimentation under the assumption that she takes action r if a breakthrough occurs:

$$\phi(p) = \lambda p u_r^R dt + (1 - \lambda p dt)\phi(p + dp) - cdt + o(dt).$$

Taking the limit $dt \to 0$, we obtain the following ODE

$$c = \lambda p(u_r^R - \phi(p)) + \phi'(p)\eta(p). \tag{1}$$

Let $\phi(p; q, \Phi_0)$ denote the solution with boundary value (p', Φ_0), given by

$$\phi(p\,; p', \Phi_0) = \frac{p - p'}{1 - p'} u_r^R + \frac{1 - p}{1 - p'}\Phi_0 - \left[\frac{p - p'}{1 - p'} + (1 - p)\log\left(\frac{p}{1 - p}\frac{1 - p'}{p'}\right)\right]\frac{c}{\lambda}.$$

*Auster: University of Bonn, auster.sarah@gmail.com. Che: Columbia University, yeonkooche@gmail.com. Mierendorff: University College London.

Conditional on taking action ℓ after stopping, the value of this stopping problem for a Bayesian DM is

$$\Phi(p) := \max_{p' \in [0,p]} \phi(p; p', U_\ell(p')).$$

We know that at $p = p_\ell^B$, we have $\phi'(p; p, U_\ell(p)) = U_\ell'(p) = \delta_\ell$. Solving this condition for p delivers

$$p_\ell^B = \frac{1}{u_r^R - u_\ell^R} \frac{c}{\lambda}.$$

The DM optimally follows the derived stopping rule as long as it delivers a higher payoff than taking action r. The right boundary p_r^B is thus determined by the indifference condition $\Phi(p_r^B) = U_r(p_r^B)$. The Bayesian value function is

$$\Phi^*(p) = \max\{\Phi(p), U_r(p)\}.$$

Experimentation is optimal for some $p \in (0,1)$ if the left boundary p_ℓ^B is smaller than the belief at which the stopping payoffs $U_\ell(\cdot)$ and $U_r(\cdot)$ cross, given by

$$\hat{p} = \frac{u_\ell^L - u_r^L}{\delta_\ell + \delta_r}.$$

Substituting for p_ℓ^B and $\hat{p}$ and solving for c, the inequality $p_\ell^B < \hat{p}$ can be written as

$$c < \bar{c} := \lambda \frac{(u_r^R - u_\ell^R)(u_\ell^L - u_r^L)}{(u_r^R - u_\ell^R) + (u_\ell^L - u_r^L)}.$$

□

B.2 Maxmin Commitment Solution

Proposition B.2. *The maxmin value for the DM with a commitment ability is*

$$\min_{p \in \mathcal{P}_0} \Phi^*(p).$$

- *If $p_* < p_r^B$ or $p_r^B \notin \mathcal{P}_0$, the DM's maxmin strategy coincides with her Bayesian optimal strategy, as described in Proposition B.1, for prior belief $p_{min} \in \arg\min_{p \in \mathcal{P}_0} \Phi^*(p)$.*

- *If $p_* = p_r^B \in \mathcal{P}_0$, then the DM randomizes as follows:*
  - *$c < \bar{c}$: the DM randomizes between action r and acquiring information using the Bayesian optimal stopping rule for p_r^B with probabilities ξ and $1 - \xi$, respectively, where ξ is such that*

$$\xi U_r' + (1 - \xi)\Phi'(p_r^B) = 0; \tag{2}$$

  - *$c \geq \bar{c}$: the DM randomizes between actions r and ℓ with probabilities $\hat{\rho}$ and $1 - \hat{\rho}$, respectively, where $\hat{\rho}$ is specified by (6).*

Proof. We consider a zero-sum game between the DM and adversarial nature seeking to minimize the DM's expected payoff. In this game, the DM chooses a distribution over stopping times and a distribution over $\{\ell, r\}$ conditional on stopping, while nature chooses $p \in \mathcal{P}_0$. We want to show that the DM's strategy described in Proposition B.2, together with nature choosing $p_{min} \in \min_{p \in \mathcal{P}_0} \Phi^*(p)$, forms a saddle point. As is well-known, if a saddle point exists, then the DM's strategy in the saddle point is maxmin-optimal.[1]

Since the DM's strategy is Bayesian optimal against nature's choice p_{min},[2] it suffices to show is that p_{min} minimizes the DM's expected payoff from the strategy described in Proposition B.2.

To this end, we first examine the value $V_0(p)$ of the DM's strategy specified in the Proposition

[1] To see this, let $u(s_1, s_2)$ be player 1's payoff when she chooses $s_1 \in S_1$ in a zero-sum game against player 2 who chooses s_2, where S_1 and S_2 are arbitrary (nonempty) sets. Suppose there exists a saddle point (s_1^*, s_2^*) such that

$$u(s_1^*, s_2) \geq u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall (s_1, s_2) \in S_1 \times S_2. \tag{*}$$

Then,

$$\inf_{s_2 \in S_2} u(s_1^*, s_2) \geq u(s_1^*, s_2^*) = \sup_{s_1 \in S_1} u(s_1, s_2^*) \geq \inf_{s_2 \in S_2} \sup_{s_1 \in S_1} u(s_1, s_2) \geq \sup_{s_1 \in S_1} \inf_{s_2 \in S_2} u(s_1, s_2),$$

where the first inequality follows from the first inequality of (*) and the last inequality follows from the min-max inequality. The above inequalities prove that

$$s_1^* \in \arg \max_{s_1 \in S_1} \inf_{s_2 \in S_2} u(s_1, s_2).$$

Symmetrically, one can show that $s_2^* \in \arg\min_{s_2 \in S_2} \sup_{s_1 \in S_1} u(s_1, s_2)$. Conversely, it can be proven that if s_1^* is maxmin optimal and s_2^* is minmax optimal, then (s_1^*, s_2^*) is a saddle point.

[2] Note that the Bayesian DM with belief p_r^B is indifferent between taking action r and taking action ℓ if $c \geq \bar{c}$ or acquiring information using the Bayesian optimal stopping rule for p_r^B if $c < \bar{c}$.

as a function of nature's choice $p \in [0,1]$:

$$V_0(p) = \begin{cases} \Phi^*(p_{min}) + (p - p_{min})\Phi^{*\prime}(p_{min}) & \text{if } p_{min} \neq p_r^B \\ U_r(p) & \text{if } p_{min} = p_r^B > p_* \\ \xi U_r(p) + (1-\xi)(\Phi(p_*) + (p - p^*)\Phi'(p_*)) & \text{if } p_{min} = p_r^B = p_* \text{ and } c < \overline{c} \\ \hat{\rho} U_r(p) + (1-\hat{\rho})U_\ell(p) & \text{if } p_{min} = p_r^B = p_* \text{ and } c \geq \overline{c}. \end{cases}$$

This function is obtained by means of several observations. First, note that $V_0(p)$ is a convex combination of the expected payoff from the DM's strategy in state R and in state L with weight p. Hence, $V_0(p)$ is linear in p. Second, the value of any strategy, including the one in question, is no higher than that of the Bayesian optimal strategy. Hence, we have $V_0(p) \leq \Phi^*(p)$ for all $p \in [0,1]$. Third, recall that the DM's strategy is Bayesian optimal for the minimizing belief p_{min}, so $V_0(p_{min}) = \Phi^*(p_{min})$. These three observations imply that $V_0(\cdot)$ is tangent to $\Phi^*(\cdot)$ at $(p_{min}, \Phi^*(p_{min}))$. Recalling that Φ^* is differentiable everywhere except possibly at p_r^B pins down the above value function.

We are now in a position to prove that $p_{\min} \in \arg\min_{p \in \mathcal{P}_0} V_0(p)$. Note first that since $\Phi^*(\cdot)$ is a convex function, we have $p_{min} = \max \mathcal{P}_0$ if $\max \mathcal{P}_0 < p_*$, $p_{min} = \min \mathcal{P}_0$ if $p_* < \min \mathcal{P}_0$, and $p_{min} = p_*$ otherwise. Hence, it suffices to show that $V_0'(\cdot) \leq 0$ if $p_{min} = \max \mathcal{P}_0 < p_*$, $V_0'(\cdot) \geq 0$ if $p_{min} = \min \mathcal{P}_0 > p_*$ and $V_0'(\cdot) = 0$ if $p_{min} = p_*$.

Suppose first $p_{min} \neq p_r^B$. In this case, $V_0'(\cdot) = \Phi^{*\prime}(p_{min})$, which is negative if $p_{min} < p_*$, positive if $p_{min} > p_*$, and zero if $p_{min} = p_*$, as desired. Suppose next $p_{min} = p_r^B > p_*$. Then, $V_0'(\cdot) = U_r'(\cdot) > 0$, so nature's optimality condition is again satisfied. Finally, if $p_{min} = p_r^B = p_*$, then, by definition of ξ and $\hat{\rho}$, $V_0(p)$ is constant in p. Hence, nature is indifferent between all values of $p \in \mathcal{P}_0$. □

B.3 Stopping-time formulation of strategies

The strategy the DM chooses can be formulated as a stopping time adapted to the filtration generated by the underlying state and the Poisson signal. In the current setting, the DM faces a nontrivial decision only when no information is revealed; once breakthrough news is received, the DM stops immediately and chooses r. Taking this as given, we can describe the DM's strategy simply by the *contingent stopping times*—a family of CDFs $\{F_t\}_{t \in [0,\infty)}$, where

$F_t : [0, \infty] \to [0, 1]$, for each $t \geq 0$,[3] and $F_t(\tau)$ denotes the probability of stopping by time $t + \tau$ starting at time t, conditional on the breakthrough news not being received by then. The reason we consider a family of stopping times instead of a single stopping time is that the strategy must specify actions off (as well as on) the path. For example, the original strategy may prescribe stopping by time t with probability one, but if the DM deviates from the original plan and has not stopped by t, the strategy must prescribe her stropping plan from then on.

We require a couple of technical assumptions:

(1) (Admissibility) F_t is nondecreasing and right-differentiable.

(2) (Consistency) For any $s < s' \leq t$,

$$F_{s'}(t - s') = \frac{F_s(t - s) - F_s(s' - s)}{(1 - F_s(s' - s))},$$

whenever $F_s(s' - s) < 1$.

Admissibility requires the distribution functions to be sufficiently well-behaved. Note also that a monotonic function has countably many points of discontinuities, which will be invoked later. Consistency means that the stopping time distribution starting at s' must form a consistent conditional stopping distribution starting at $s < s'$ according to Bayes formula, as long as the latter distribution does not prescribe stopping before s with probability one. Let $\mathcal{F}$ be the set of all families of distributions satisfying these two requirements. Let $\tilde{\rho} : [0, \infty) \to [0, 1]$ be the choice strategy specifying the probability of choosing r when stopping at time t. Then, the strategy for the DM consists of a pair $(\{F_t\}, \tilde{\rho})$ such that $\{F_t\} \in \mathcal{F}$.

As an intermediate step, we next define the following notion of strategy. A *time-indexed strategy* is a function $\tilde{\sigma} : [0, \infty) \to [0, \infty) \times [0, 1] \times [0, 1]$, $\tilde{\sigma}_t = (\tilde{\nu}_t, \tilde{m}_t, \tilde{\rho}_t)$, where $\tilde{\nu}_t$ is the stopping rate and $\tilde{m}_t$ is the instantaneous stopping probability used at time t if no R-signal has been received up until time t. ρ_t is the probability of action r conditional on stopping at time t without receiving an R-signal. We require the following *admissibility* conditions on $\tilde{\sigma}$:

(a) any connected set of t's in which $m_t = 1$ contains its infimum,

(b) the set $\{t \geq 0 | m_t \in (0, 1)\}$ is countable.

[3] We let the stopping time be ∞ if the DM never stops.

Condition (a) follows from the right continuity of F_t's, which in turn follows from its right differentiability (assumed in (1)). Condition (b) follows from the fact that each F_t has countably many points of discontinuities. Let $\widetilde{\Sigma}$ denote the set of all admissible time-indexed strategies. We now prove the equivalence between the two notions of strategies.

Lemma B.1. *Each $(\{F_t\}, \tilde{\rho}), \{F_t\} \in \mathcal{F}$ induces an outcome-equivalent admissible strategy $\tilde{\sigma} \in \widetilde{\Sigma}$. Conversely, any admissible strategy $\tilde{\sigma} \in \widetilde{\Sigma}$ induces an outcome-equivalent $(\{F_t\}, \rho), \{F_t\} \in \mathcal{F}$.*

Proof. For both statements, $\tilde{\rho}$ is fixed to be the same, so we can focus on the specification of stopping times.

Fix any $\{F_t\} \in \mathcal{F}$. For each $t \geq 0$, we define:

$$\tilde{\nu}_t := F_t'(0),$$

where $F_t'(\cdot)$ is the right-derivative, which is well-defined by Admissibility (requirement (1) above), and

$$\tilde{m}_t := F_t(0).$$

Then, $(\tilde{\nu}_t, \tilde{m}_t)_{t\geq 0}$ is outcome-equivalent to $\{F_t\}_{t\in[0,\infty)}$. To see this, fix any $s < t$. If $1 - F_s(t-s) < 1$, then by Consistency (requirement (2) above),

$$\frac{F_s'(t-s)}{1 - F_s(t-s)} = F_t'(0) = \tilde{\nu}_t$$

and

$$\frac{F_s(t-s) - F_s^-(t-s)}{1 - F_s(t-s)} = F_t(0) = \tilde{m}_t,$$

where F_s^- is the left limit, which is well defined. Hence, for any $s < t$,

$$\begin{aligned} & \Pr\{\text{stopping by } t \text{ under } \tilde{\sigma} | \text{ starting at } s\} \\ =& 1 - e^{-\int_s^t \tilde{\nu}_{s'} ds'} \prod_{s' \in [s,t]} (1 - \tilde{m}_{s'}) \\ =& F_s(t-s). \end{aligned}$$

Finally, the right-continuity and monotonicity of each $F_s(\cdot)$ imply that the admissibility of $\tilde{\sigma}$ is satisfied. In particular, $\tilde{m}_t > 0$ for countably many t. We thus conclude $(\tilde{\nu}, \tilde{m}, \tilde{\rho}) \in \widetilde{\Sigma}$.

Conversely, fix any $\tilde{\sigma} \in \widetilde{\Sigma}$. Then, for any $s \leq t$, define

$$F_s(t-s) := 1 - e^{-\int_s^t \tilde{\nu}_{s'} ds'} \prod_{s' \in [s,t]} (1 - \tilde{m}_{s'}).$$

Again, the admissibility implies that $\{F_t\}$ satisfies (1). Requirement (2) follows from the construction. Hence, $\{F_t\} \in \mathcal{F}$. $\square$

B.4 Proof of Lemma 1

Since we are considering the left limit of the state $\overline{p}$, a possible mass point at state $\overline{p}$ does not appear in $W_\epsilon^\sigma(p, \overline{p}_-)$. Specifically, we can write

$$\begin{aligned}
W_\epsilon(\pi_\epsilon, \overline{p}_-) =& W_\epsilon^{\sigma_\epsilon}(\pi_\epsilon, \overline{p}_-) \\
=& -\pi_\epsilon \int_{\overline{p}^\epsilon}^{\overline{p}} e^{-\int_0^{\tau(\overline{p}, \overline{p}')} (\lambda + \nu_\epsilon(p^\tau)) d\tau} \left(-c + \lambda u_r^R + \nu_\epsilon(\overline{p}') u_\ell^R\right) \frac{1}{\eta(\overline{p}')} d\overline{p}' \\
& - (1 - \pi_\epsilon) \int_{\overline{p}^\epsilon}^{\overline{p}} e^{-\int_0^{\tau(\overline{p}, \overline{p}')} \nu_\epsilon(p^\tau) d\tau} \left(-c + \nu_\epsilon(\overline{p}') u_\ell^L\right) \frac{1}{\eta(\overline{p}')} d\overline{p}' \\
& + \left(\pi_\epsilon e^{-\int_0^{\tau(\overline{p}, \overline{p}^\epsilon)} (\lambda + \nu_\epsilon(p^\tau)) d\tau} + (1 - \pi_\epsilon) e^{-\int_0^{\tau(\overline{p}, \overline{p}^\epsilon)} \nu_\epsilon(p^\tau) d\tau}\right) V^{\sigma^*}(\pi_\epsilon^\epsilon, \overline{p}^\epsilon).
\end{aligned}$$

Note as $\epsilon \to 0$, $\overline{p}^\epsilon \to \overline{p}$ and $\tau(\overline{p}, \overline{p}^\epsilon) \to 0$. Hence, as $\epsilon \to 0$, $W_\epsilon(\pi_\epsilon, \overline{p}_-) \to V^{\sigma^*}(\pi^*, \overline{p}_-)$. Next, consider

$$\begin{aligned}
\frac{\partial W_\epsilon(\pi_\epsilon, \overline{p}_-)}{\partial p} =& -\int_{\overline{p}^\epsilon}^{\overline{p}} e^{-\int_0^{\tau(\overline{p}, \overline{p}')} (\lambda + \nu_\epsilon(p^\tau)) d\tau} \left(-c + \lambda u_r^R + \nu_\epsilon(\overline{p}') u_\ell^R\right) \frac{1}{\eta(\overline{p}')} d\overline{p}' \\
& + \int_{\overline{p}^\epsilon}^{\overline{p}} e^{-\int_0^{\tau(\overline{p}, \overline{p}')} \nu_\epsilon(p^\tau) d\tau} \left(-c + \nu_\epsilon(\overline{p}') u_\ell^L\right) \frac{1}{\eta(\overline{p}')} d\overline{p}' \\
& + \left(e^{-\int_0^{\tau(\overline{p}, \overline{p}^\epsilon)} (\lambda + \nu_\epsilon(p^\tau)) d\tau} - e^{-\int_0^{\tau(\overline{p}, \overline{p}^\epsilon)} \nu_\epsilon(p^\tau) d\tau}\right) V^{\sigma^*}(\pi^\epsilon, \overline{p}^\epsilon) \\
& + \left(\pi_\epsilon e^{-\int_0^{\tau(\overline{p}, \overline{p}^\epsilon)} (\lambda + \nu_\epsilon(p^\tau)) d\tau} + (1 - \pi_\epsilon) e^{-\int_0^{\tau(\overline{p}, \overline{p}^\epsilon)} \nu_\epsilon(p^\tau) d\tau}\right) V_p^{\sigma^*}(\pi^\epsilon, \overline{p}^\epsilon) \frac{\eta(\pi_\epsilon^\epsilon)}{\eta(\pi^\epsilon)} \\
& \to V_p^{\sigma^*}(\pi^*, \overline{p}_-) \text{ as } \epsilon \to 0,
\end{aligned}$$

since $\pi_\epsilon \to \pi^*, \overline{p}^\epsilon \to \overline{p}$, and $\tau(\overline{p}, \overline{p}^\epsilon) \to 0$, as $\epsilon \to 0$, provided that $V_p^{\sigma^*}(\cdot, \cdot)$ is continuous at $(\pi^*, \overline{p}_-)$.

Finally, consider

$$\begin{aligned}
&\frac{\partial W_\epsilon(\pi_\epsilon,\overline{p}_-)}{\partial \overline{p}}\\
=&-c+\nu_\epsilon(U_\rho(\pi_\epsilon)-W_\epsilon^{\sigma_\epsilon}(\pi_\epsilon,\overline{p}))+\pi_\epsilon\lambda(u_r^R-W_\epsilon^{\sigma_\epsilon}(\pi_\epsilon,\overline{p}))+\frac{\partial W_\epsilon^{\sigma_\epsilon}(\pi_\epsilon,\overline{p})}{\partial p}\eta(p)+\frac{\partial W_\epsilon^{\sigma_\epsilon}(\pi_\epsilon,\overline{p})}{\partial \overline{p}}\eta(\overline{p})\\
&+\left(\pi_\epsilon e^{-\int_0^{\tau(\overline{p},\overline{p}^\epsilon)}(\lambda+\nu_\epsilon(p^\tau))d\tau}+(1-\pi_\epsilon)e^{-\int_0^{\tau(\overline{p},\overline{p}^\epsilon)}\nu_\epsilon(p^\tau)d\tau}\right)V_{\overline{p}}^{\sigma^*}(\pi^\epsilon,\overline{p}^\epsilon)\frac{\eta(\overline{p}^\epsilon)}{\eta(\overline{p})}\\
=&\left(\pi_\epsilon e^{-\int_0^{\tau(\overline{p},\overline{p}^\epsilon)}(\lambda+\nu_\epsilon(p^\tau))d\tau}+(1-\pi_\epsilon)e^{-\int_0^{\tau(\overline{p},\overline{p}^\epsilon)}\nu_\epsilon(p^\tau)d\tau}\right)V_{\overline{p}}^{\sigma^*}(\pi^\epsilon,\overline{p}^\epsilon)\frac{\eta(\overline{p}^\epsilon)}{\eta(\overline{p})}\\
\to& V_{\overline{p}}^{\sigma^*}(\pi^*,\overline{p}_-)\text{ as }\epsilon\to 0,
\end{aligned}$$

where the second line holds because of the HJB for ϵ-commitment solution, i.e., by (A7).

B.5 Proof of Theorem 1

For our purpose, we first establish a preliminary result about the Bayesian value function. Recall the solution ϕ to the ODE (2). Then, the following lemma holds.

Lemma B.2. *For any $p > p_\ell^B$, $\phi'(p;p,U_\ell(p)) - U_\ell'(p) > 0$ whenever $\phi(p;p,U_\ell(p)) = U_\ell(p)$.*

Proof. Define $\delta^\omega := |\delta_r^\omega - \delta_\ell^\omega|, \omega = L, R$. It follows from (2) that

$$\begin{aligned}
\lambda p(1-p)(\phi'(p;p,U_\ell(p))-U_\ell'(p)) &= \lambda p(u_r^R-\phi(p;p,U_\ell(p)))-c+\lambda(1-p)(u_\ell^L-U_\ell(p))\\
&=\lambda\left(pu_r^R+(1-p)u_\ell^L-U_\ell(p)\right)-c\\
&=\lambda p\delta^R-c,
\end{aligned}$$

where the second equality uses the fact that $\phi(p;p,U_\ell(p))) = U_\ell(p)$. Since the last line is strictly increasing in p and equals 0 when $p = p_\ell^B$, the above claim is proven. □

For the proof, for each region of $\overline{p}$'s, we will specify the strategy profile (σ,π) and the associated value $V^\sigma(p,\overline{p})$ that are not fully specified in Theorem 1. We will then verify the value function together with the strategy profile (σ,π) satisfies the equilibrium conditions.

For ease of presentation, we shall mainly focus on Case 1. Further, we will assume that either $c < \underline{c}$ or $c \in (\underline{c},\overline{c}]$ but $\Delta < \Delta_c$, where Δ_c will be specified later (in Region 3 in Section B.5.1).

We call this case *main case*. We will first treat the main case in Section B.5.1. The remaining cases will be treated in Section B.5.2 and Section B.6.

B.5.1 Main Case

Region 1: $\overline{p} \in [0, \overline{p}_1]$, where $\overline{p}_1 := p_\ell^B$.

Computation of equilibrium value. Recall that the strategy σ calls for an immediate choice of ℓ for all state $\overline{p} \le p_\ell^B = \overline{p}_1$. It thus immediately follows that the value associated with that strategy is:

$$V^\sigma(p, \overline{p}) := U_\ell(p), \text{ for all } \overline{p} \in [0, \overline{p}_1].$$

Verification of equilibrium conditions. With $V^\sigma(p, \overline{p}) = U_\ell(p)$, $m(\overline{p}) = 1$ and $\rho(\overline{p}) = 0$ imply that $\pi(\overline{p}) = \overline{p}$ is a minimizer in (13). Since $U_\ell'(p) < 0$ by assumption, $\pi(\overline{p}) = \overline{p}$ is the unique minimizer in (7). Substituting $V^\sigma(p, \overline{p}) = U_\ell(p)$, $V_p^\sigma(p, \overline{p}) = U_\ell'(p)$, $V_{\overline{p}}^\sigma(p, \overline{p}) = 0$, and $\pi(\overline{p}) = \overline{p}$, (11) simplifies to

$$\max_{m,\nu,\rho} m\left[U_\rho(\overline{p}) - U_\ell(\overline{p})\right] + (1-m)\left[-c + \nu(U_\rho(\overline{p}) - U_\ell(\overline{p})) + \overline{p}\lambda(u_r^R - U_\ell(\overline{p})) + U_\ell'(\overline{p})\eta(\overline{p})\right] = 0.$$

With $m = 1$ and $\rho = 0$, the LHS is zero. Moreover, for any ρ, ν, simple algebra shows that the coefficient of $(1-m)$ is non-positive for all $\overline{p} \le \overline{p}_1$. Hence $\sigma = (1, 0, 0)$ is a maximizer and both (11) and (12) are satisfied for $V^\sigma(p, \overline{p}) = U_\ell(p)$. Setting $m(\overline{p}) = 1$ and $\rho(\overline{p}) = 0$, the objective in (13) is independent of p and $\pi(\overline{p}) = \overline{p}$ is a minimizer. Hence we have shown that for the posited strategy profile (σ, π) satisfy the equilibrium conditions, (7)–(13) for all $\overline{p} \in [0, \overline{p}_1]$.

Region 2: $\overline{p} \in (\overline{p}_1, \overline{p}_2]$, where $\overline{p}_2 := p_*$.

Computation of equilibrium value. Recall that the strategy σ calls on the DM to experiment until $\overline{p}$ drifts to $\overline{p}_1 = p_\ell^B$. Fix any state $\overline{p}$ in this region. To compute the associated value $V^\sigma(\cdot, \overline{p})$, it is useful to ask: *for which belief $p \in \mathcal{P}(\overline{p})$ is the continuation strategy $\sigma_{\overline{p}}$ Bayesian optimal?* The answer is $p = \overline{p}$, since the strategy is exactly what a Bayesian DM with belief $\overline{p}$ will do—experimenting until $\overline{p}$ reaches the Bayesian optimal stopping belief $p_\ell^B = \overline{p}_1$. Consequently, we must have

$$V^\sigma(\overline{p}, \overline{p}) = \Phi(\overline{p}).$$

Since the strategy needs not be optimal for any DM with belief $p \neq \overline{p}$ (including p outside $\mathcal{P}(\overline{p})$), we must have

$$V^{\sigma}(p, \overline{p}) \leq \Phi(p), \forall p \in [0,1].$$

Finally, since $V^{\sigma}(p, \overline{p})$—the valuation of a fixed action path—must be linear in p, the preceding observations must imply:

$$V^{\sigma}(p, \overline{p}) := \Phi(\overline{p}) + (p - \overline{p})\Phi'(\overline{p}), \text{ for all } \overline{p} \in (\overline{p}_1, \overline{p}_2]. \tag{3}$$

Verification of equilibrium conditions. Since $\Phi'(\overline{p}) \leq 0$ for all $\overline{p} \in (\overline{p}_1, \overline{p}_2]$, $V^{\sigma}_p(p, \overline{p}) \leq 0$ and therefore $\pi(\overline{p}) = \overline{p}$ is a minimizer in (7). With $\pi(\overline{p}) = \overline{p}$ we have

$$V^{\sigma}(\pi(\overline{p}), \overline{p}) = \Phi(\overline{p}), \quad V^{\sigma}_p(\pi(\overline{p}), \overline{p}) = \Phi'(\overline{p}), \quad V^{\sigma}_{\overline{p}}(\pi(\overline{p}), \overline{p}) = 0.$$

Therefore (11) simplifies to

$$\max_{m,\nu,\rho} m\left[U_\rho(\overline{p}) - \Phi(\overline{p})\right] + (1-m)\left[-c + \nu(U_\rho(\overline{p}) - \Phi(\overline{p})) + \overline{p}\lambda(u^R_r - \Phi(\overline{p})) + \Phi'(\overline{p})\eta(\overline{p})\right] = 0.$$

Since $\Phi(\overline{p}) \geq U(\overline{p})$ for any $\overline{p}$, $\nu = 0$ is optimal. Moreover, $m = 0$ is optimal since we have from (2)

$$-c + \overline{p}\lambda(u^R_r - \Phi(\overline{p})) + \Phi'(\overline{p})\eta(\overline{p}) = 0$$

for the Bayesian value function. Therefore (11) and (12) are satisfied for all $\overline{p} \in (\overline{p}_1, \overline{p}_2]$. Condition (13) holds since $m(\overline{p}) = \nu(\overline{p}) = 0$ implies that the objective in (13) is given by

$$\begin{aligned} & -c + p\lambda\left(u^R_r - (\Phi(\overline{p}) + (p - \overline{p})\Phi'(\overline{p}))\right) - \lambda p(1-p)\Phi'(\overline{p}) - \lambda\overline{p}(1-\overline{p})(p-\overline{p})\Phi''(\overline{p}) \\ = & -c + p\lambda\left(u^R_r - (\Phi(\overline{p}) + (1 - \overline{p})\Phi'(\overline{p}))\right) - \lambda\overline{p}(1-\overline{p})(p-\overline{p})\Phi''(\overline{p}). \end{aligned}$$

Differentiating this with respect to p yields

$$\lambda\left(u^R_r - (\Phi(\overline{p}) + (1 - \overline{p})\Phi'(\overline{p}))\right) - \lambda\overline{p}(1-\overline{p})\Phi''(\overline{p}).$$

Differentiating the ODE for the Bayesian Value function (see (2)) with respect to $\overline{p}$ we obtain that this expression is equal to zero. Hence $\pi(\overline{p})$ is a minimizer in (13).

To summarize, we have shown that for the posited profile (σ, π), together with the value function V^σ, satisfy (7)–(13) for all $\overline{p} \in (\overline{p}_1, \overline{p}_2]$.

Region 3: $\overline{p} \in (\overline{p}_2, \overline{p}_3]$**.**

This region exists only when the ambiguity is sufficiently large so that $V^\sigma(\underline{p}(\overline{p}_2), \overline{p}_2) < U_\ell(\underline{p}(\overline{p}_2))$. If this inequality is reversed, then we set $\overline{p}_3 = \overline{p}_2$, and Region 3 is empty. Assuming the inequality, we will specify the upper bound $\overline{p}_3$, the critical cost $\underline{c}$, and ambiguity level Δ_c referred to in the statement of the theorem. Finally, we will specify $\nu(\overline{p})$ and $\pi(\overline{p})$ fully for $\overline{p}$ in this region.

Specification of ν and π and equilibrium value. The strategy σ for this region involves randomization between experimentation and stopping for ℓ, with the latter done at a Poisson rate $\nu(\overline{p})$. Meanwhile, nature chooses belief $\pi(\overline{p}) \in \mathcal{P}(\overline{p})$. Here, we specify (ν, π) precisely, together with the value $V^\sigma(p, \overline{p})$ associated with the strategy. We first construct a value function $V(p, \overline{p})$ that satisfies the equilibrium conditions, (7)-(13), given the candidate strategy σ. We will then establish that the constructed function $V(p, \overline{p})$ indeed coincides with the value of σ.

To begin, fix any $\overline{p} > \overline{p}_2$. First, the fact that the DM randomizes between experimentation and action ℓ means that the coefficient of ν in (12) must vanish, which implies (14). For the required belief $\pi(\overline{p})$ to be nature's choice satisfying (7), it is sufficient, and will be seen also necessary, to have $V_p(p, \overline{p}) = 0$. Namely,

$$V(p, \overline{p}) = \hat{V}(\overline{p}), \forall p. \tag{4}$$

Substituting this into (14), we get (15). The randomization by DM in turn implies that the derivative of the objective in (12) with respect to p must vanish. This fact, together with (4) yields (16).

It now remains to specify the function $\hat{V}(\overline{p})$. To this end, we use (15),(16) and (4) to simplify (11) and obtain the differential equation ($\widehat{\text{ODE}}$). Together with the boundary condition that $\hat{V}(\overline{p}_2) = \Phi(\overline{p}_2)$, ($\widehat{\text{ODE}}$) admits a unique solution:[4]

$$\hat{V}(\overline{p}) := \frac{C}{\left(\frac{1-\overline{p}}{\overline{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}} + C} u_1 + \frac{\left(\frac{1-\overline{p}}{\overline{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}}}{\left(\frac{1-\overline{p}}{\overline{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}} + C} u_2, \tag{5}$$

[4]The derivation of the solution appears in Section B.5.3.

where

$$C := \frac{u_2 - \Phi(\overline{p}_2)}{\Phi(\overline{p}_2) - u_1} \left(\frac{1 - \overline{p}_2}{\overline{p}_2} \right)^{\frac{u_2 - u_1}{\delta_\ell}}, \tag{6}$$

and u_1 and u_2 are lower and higher roots of a quadratic equation; namely,

$$u_{1,2} := \frac{u_r^R + u_\ell^L}{2} \pm \sqrt{\left(\frac{u_r^R - u_\ell^L}{2} \right)^2 + \frac{c}{\lambda} \delta_\ell}. \tag{7}$$

It is instructive to visualize how nature's feasible "set of values" evolves on the equilibrium path as $\overline{p}$ drifts down in this region. Our construction so far indicates that at each state the feasible set forms a flat segment $\{(p, \hat{V}(\overline{p})) : p \in [\underline{p}(\overline{p}), \overline{p}]\}$, and the belief at which the segment crosses ℓ-payoff function $U_\ell(p)$ is precisely nature's choice $\pi(\overline{p})$. As $\overline{p}$ falls in this region, the segment shifts left. The starting state of Region 3, or its right-most boundary, $\overline{p}_3$, is the state such that its associated value segment just meets or "touches" the ℓ-payoff function $U_\ell(p)$ at the former's left-most end. Formally, $\overline{p}_3$ is defined by:

$$\hat{V}(\overline{p}_3) = U_\ell(\underline{p}(\overline{p}_3)). \tag{8}$$

Let $\hat{u} := \min_{p \in [0,1]} U(p)$. The following observations are useful to establish:

Lemma B.3. **(i)** *$\hat{V}(\cdot)$ is strictly decreasing, and $\hat{V}(\overline{p}) \in (u_1, u_2)$ for all $\overline{p} > \overline{p}_2$.*

(ii) *If $V^\sigma(\underline{p}(\overline{p}_2), \overline{p}_2) < U_\ell(\underline{p}(\overline{p}_2))$ (i.e., the condition for Region 3 to exist holds) and $u_1 > u_\ell^R$, then there exists a unique solution $\overline{p}_3 \in (\overline{p}_2, 1)$.*

(iii) *If $c \leq \underline{c}$, then $u_1 \geq \hat{u}$, so $V^\sigma(\cdot, \overline{p}_2) > \hat{u}$ for all $\overline{p} \geq \overline{p}_2$. If $c \in [\underline{c}, \overline{c})$, there exists $\Delta_c > 0$ such that $V^\sigma(\cdot, \overline{p}_3) = \hat{V}(\overline{p}_3) \geq \hat{u}$ if and only if $\Delta < \Delta_c$.*

Proof. We first prove (i). To show $\hat{V}(\overline{p}) \in (u_1, u_2)$, it suffices to show that $C > 0$ since in this case $\hat{V}^\sigma(\overline{p})$ is a convex combination of u_1 and u_2. We show that $\Phi(\overline{p}_2) \in (u_1, u_2)$ which implies that $C > 0$. To see that $\Phi(\overline{p}_2) > u_1$, note that by part (a) of this Lemma, u_1 is the value of a feasible strategy for the DM that differs from the Bayesian strategy. Therefore, the Bayesian value for any belief must strictly exceed u_1. To see that $\Phi(\overline{p}_2) < u_2$, note that $u_2 \geq \frac{1}{2}\left(u_r^R + u_\ell^L\right) + \frac{1}{2}\left|u_r^R - u_\ell^L\right| = u_r^R \vee u_\ell^L > \Phi(\overline{p}_2)$. Since $u_1 < u_2$, the term $\left(\frac{1-\overline{p}}{\overline{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}}$ is decreasing in $\overline{p}$ and therefore $\hat{V}^{\sigma\prime}(\overline{p}) < 0$.

We next prove (ii). By the assumption, we have $V^{\sigma}(\underline{p}(\overline{p}_2), \overline{p}_2) < U_\ell(\underline{p}(\overline{p}_2))$. As $\overline{p} \to 1$, $\hat{V}(\overline{p}) \to u_1$, whereas $U_\ell(\underline{p}(\overline{p})) \to u_\ell^R$, which by assumption is less than u_1.[5] Hence, by the intermediate value theorem, there exists $\overline{p}_3 \in (\overline{p}_2, 1)$ such that $V^{\sigma}(\underline{p}(\overline{p}_3), \overline{p}_3) = \hat{V}(\overline{p}_3) = U_\ell(\underline{p}(\overline{p}_3))$. The fact that such a $\overline{p}_3$ is unique follows from the fact that $\hat{V}(\cdot)$ can only cross $U_\ell(\underline{p}(\cdot))$ from below:

$$
\begin{aligned}
\left.\frac{d\left(\hat{V}^{\sigma}(\overline{p}) - U_\ell(\underline{p}(\overline{p}))\right)}{d\overline{p}}\right|_{U_\ell(\underline{p}(\overline{p})) = \hat{V}(\overline{p})} &= \frac{\left(u_r^R - \underline{p}u_\ell^R - (1-\underline{p})u_\ell^L\right)\left(u_\ell^L - \underline{p}u_\ell^R - (1-\underline{p})u_\ell^L\right)}{u_\ell^L - u_\ell^R} - \frac{c}{\lambda} - \left(u_\ell^R - u_\ell^L\right)\underline{p}(1-\underline{p}) > 0 \\
&\iff \left(u_r^R - \underline{p}u_\ell^R - (1-\underline{p})u_\ell^L\right)\underline{p} - \frac{c}{\lambda} - \left(u_\ell^R - u_\ell^L\right)\underline{p}(1-\underline{p}) > 0 \\
&\iff \left(u_r^R - u_\ell^R\right)\underline{p} - \frac{c}{\lambda} > 0 \\
&\iff \underline{p} > p_\ell^B,
\end{aligned}
$$

where we have used the shorthand $\underline{p} = \underline{p}(\overline{p})$. Note that $U_\ell(\underline{p}(\overline{p})) = \hat{V}(\overline{p})$ implies that $\underline{p} > p_\ell^B$ since $U_\ell(p_\ell^B) = \Phi(p_\ell^B) > \Phi(\overline{p}_2) = \hat{V}(\overline{p}_2) > \hat{V}(\overline{p}) = U_\ell(\underline{p})$ and since $U_\ell(\cdot)$ is decreasing.

Finally, we prove (iii). Simple algebra shows that $u_1 > \hat{u}$ if and only if $c < \underline{c}$:

$$
\begin{aligned}
u_1 > \hat{u} &\iff \frac{u_r^R + u_\ell^L}{2} - \sqrt{\left(\frac{u_r^R - u_\ell^L}{2}\right)^2 + \frac{c}{\lambda}\delta_\ell} > \frac{u_r^R u_\ell^L - u_r^L u_\ell^R}{\delta_r + \delta_\ell} \\
&\iff \frac{\delta^R\delta_r + \delta^L\delta_\ell}{2(\delta_r + \delta_\ell)} > \sqrt{\left(\frac{u_r^R - u_\ell^L}{2}\right)^2 + \frac{c}{\lambda}\delta_\ell} \\
&\iff \left(\frac{\delta^R\delta_r + \delta^L\delta_\ell}{2(\delta_r + \delta_\ell)}\right)^2 > \left(\frac{u_r^R - u_\ell^L}{2}\right)^2 + \frac{c}{\lambda}\delta_\ell \\
&\iff c < \frac{\lambda}{\delta_\ell}\left(\frac{\delta^R\delta_r + \delta^L\delta_\ell}{2(\delta_r + \delta_\ell)} + \frac{u_r^R - u_\ell^L}{2}\right)\left(\frac{\delta^R\delta_r + \delta^L\delta_\ell}{2(\delta_r - \delta_\ell)} + \frac{u_r^R - u_\ell^L}{2}\right) \\
&\iff c < \lambda\frac{\delta^R\delta^L\delta_r}{(\delta_r + \delta_\ell)^2} = \frac{\delta_r}{\delta_r + \delta_\ell}\overline{c} = \underline{c}.
\end{aligned}
$$

For the last statement, fix any $c \in (\underline{c}, \overline{c})$. Then, $u_1 < \hat{u}$, as shown above. Rewrite (8)

$$
\hat{V}(\overline{p}_3(\Delta)) = U_\ell(\underline{p}(\overline{p}_3(\Delta)))
$$

so that the dependence of $\overline{p}_3$ is made explicit. Note first that neither $\hat{V}(\cdot)$ nor $U_\ell(\cdot)$ depends on the size of ambiguity Δ. The only source of dependence arises from the dependence of $\underline{p}(\cdot)$ on

[5]This follows from the fact that $\lim_{\overline{p}\to 1}\left(\frac{1-\overline{p}}{\overline{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}} = 0$. Note that $\underline{p}(\overline{p}) \to 1$ as $\overline{p} \to 1$.

Δ. Clearly $\underline{p}(\overline{p})$ is strictly decreasing in Δ for each $\overline{p}$. Since $\hat{V}(\cdot)$ can only cross $U_\ell(\underline{p}(\cdot))$ from below (as was shown above), it follows that $\overline{p}_3(\Delta)$ is strictly increasing in Δ. Further, as Δ becomes sufficiently small, $\overline{p}_3(\Delta) \to \overline{p}_2$ and as Δ becomes sufficiently large, then $\overline{p}_3(\Delta) \to 1$. Note that $\hat{V}(\overline{p}_2) = \Phi(\overline{p}_2) > \max\{U_r(\overline{p}_2), U_\ell(\overline{p}_2)\} \geq \hat{u}$ since $c < \overline{c}$ and that $\lim_{\overline{p}\to 1} \hat{V}(\overline{p}) = u_1 < \hat{u}$ since $c < \underline{c}$. We thus conclude that there exists a unique Δ_c that satisfies

$$\hat{V}(\overline{p}_3(\Delta_c)) = \hat{u}.$$

Since $\hat{V}(\cdot)$ is strictly decreasing and $\overline{p}_3(\cdot)$ is strictly increasing, it follows that $V^\sigma(\cdot, \overline{p}_3) = \hat{V}(\overline{p}_3) \geq \hat{u}$ if and only if $\Delta < \Delta_c$. $\square$

Lemma B.3-(i) implies that the value segment shifts up (as well as shifts left) as $\overline{p}$ drifts down.[6] Next, Lemma B.3-(ii) implies that $\overline{p}_3$ is well defined and above $\overline{p}_2$ whenever the region exists. Last, Lemma B.3-(iii) ensures that the value segment remains above $\hat{u}$ for all $\overline{p} \in (\overline{p}_2, \overline{p}_3]$, if either $c \leq \underline{c}$ or if $c \in [\underline{c}, \overline{c})$ but $\Delta < \Delta_c$, for some $\Delta_c > 0$. Our verification below will use this fact, assuming the sufficient condition.

So far, our value function $V(p, \overline{p})$ is constructed from HJB equations along with (7). We now claim the value function indeed represents the value of the candidate strategy σ:

Lemma B.4. *For all $\overline{p} \in (\overline{p}_2, \overline{p}_3)$, $V(p, \overline{p}) = V^\sigma(p, \overline{p})$.*

Proof. The value of strategy σ is given by:

$$\begin{aligned} V^\sigma(p, \overline{p}) =& p \int_{\overline{p}_2}^{\overline{p}} e^{-\int_0^{\tau(\overline{p}, \overline{p}')} (\lambda + \nu(p_\tau)) d\tau} \left(-c + \lambda u_r^R + \nu(\overline{p}') u_\ell^R\right) \frac{1}{\lambda \overline{p}'(1 - \overline{p}')} d\overline{p}' \\ &+ (1-p) \int_{\overline{p}_2}^{\overline{p}} e^{-\int_0^{\tau(\overline{p}, \overline{p}')} \nu(p_\tau) d\tau} \left(-c + \nu(\overline{p}') u_\ell^L\right) \frac{1}{\lambda \overline{p}'(1 - \overline{p}')} d\overline{p}' \\ &+ \left(p e^{-\int_0^{\tau(\overline{p}, \overline{p}_2)} (\lambda + \nu(p_\tau)) d\tau} + (1-p) e^{-\int_0^{\tau(\overline{p}, \overline{p}_2)} \nu(p_\tau) d\tau}\right) V^\sigma(p_{\tau(\overline{p}, \overline{p}_2)}(p), \overline{p}_2). \end{aligned}$$

Using the change of variables $\tau = \tau(\overline{p}, \overline{p}'')$ and hence $\overline{p}_\tau = \overline{p}''$ and $d\tau = -\frac{1}{\lambda \overline{p}''(1 - \overline{p}'')} d\overline{p}''$ we can rewrite this as follows:[7]

[6] One can also see that V^σ is continuous, including at the boundary. This follows from the fact that $\hat{V}(\overline{p}_2) = \Phi(\overline{p}_2)$, and that $\overline{p}_2 = p_*$ and $\Phi'(p_*) = 0$.

[7] Remember that $\overline{p}_t = \frac{\overline{p}_0}{\overline{p}_0 + (1 - \overline{p}_0) e^{\lambda t}}$. Setting $\overline{p}_0 = \overline{p}$ and $\overline{p}_t = \overline{p}'$ we get

$$\tau(\overline{p}, \overline{p}') = \frac{1}{\lambda} \left(\log \frac{\overline{p}}{(1 - \overline{p})} - \log \frac{\overline{p}'}{1 - \overline{p}'} \right).$$

$$
\begin{aligned}
V^{\sigma}(p,\overline{p}) =& p\int_{\overline{p}_2}^{\overline{p}} e^{-\int_{\overline{p}'}^{\overline{p}}(\lambda+\nu(\overline{p}''))\frac{d\overline{p}''}{\lambda\overline{p}''(1-\overline{p}'')}}\left(-c+\lambda u_r^R+\nu(\overline{p}')u_\ell^R\right)\frac{1}{\lambda\overline{p}'(1-\overline{p}')}d\overline{p}' \\
&+(1-p)\int_{\overline{p}_2}^{\overline{p}} e^{-\int_{\overline{p}'}^{\overline{p}}\nu(\overline{p}'')\frac{d\overline{p}''}{\lambda\overline{p}''(1-\overline{p}'')}}\left(-c+\nu(\overline{p}')u_\ell^L\right)\frac{1}{\lambda\overline{p}'(1-\overline{p}')}d\overline{p}' \\
&+\left(pe^{-\int_{\overline{p}_2}^{\overline{p}}(\lambda+\nu(\overline{p}''))\frac{d\overline{p}''}{\lambda\overline{p}''(1-\overline{p}'')}}+(1-p)e^{-\int_{\overline{p}_2}^{\overline{p}}\nu(\overline{p}'')\frac{d\overline{p}''}{\lambda\overline{p}''(1-\overline{p}'')}}\right)V^{\sigma}(p_{\tau(\overline{p},\overline{p}_2)}(p),\overline{p}_2).
\end{aligned}
$$

We wish to show that $V(p,\overline{p})$ defined in Region 3 coincides with $V^{\sigma}(p,\overline{p})$ above. To this end, note first that $V(\cdot,\overline{p}_2)=\hat{V}(\overline{p}_2)=V^{\sigma}(\cdot,\overline{p}_2)$, by the definition of $\hat{V}(\overline{p})$ at $\overline{p}=\overline{p}_2$. Recall that we defined $V(p,\overline{p})$ in Region 3 so that $V_p(p,\overline{p})=0$ and that

$$
c=\lambda p\left(u_r^R-V(p,\overline{p})\right)+\nu(\overline{p})\left(U_\ell(\overline{p})-V(p,\overline{p})\right)+\eta(p)V_p(p,\overline{p})+\eta(\overline{p})V_{\overline{p}}(p,\overline{p}) \tag{9}
$$

for all $p\in[0,1]$ and $\overline{p}\in(\overline{p}_2,\overline{p}_3]$.

To show that $V(p,\overline{p})=V^{\sigma}(p,\overline{p})$, substitute $p=0$ and $p=1$ in (9) to obtain two ODEs:

$$
\begin{aligned}
V_{\overline{p}}(0,\overline{p}) &= \frac{c-\nu(\overline{p})u_\ell^L}{\eta(\overline{p})}+\frac{\nu(\overline{p})}{\eta(\overline{p})}V(0,\overline{p}); \\
V_{\overline{p}}(1,\overline{p}) &= \frac{c-\lambda u_r^R-\nu(\overline{p})u_\ell^R}{\eta(\overline{p})}+\frac{\lambda+\nu(\overline{p})}{\eta(\overline{p})}V(1,\overline{p}).
\end{aligned}
$$

Integrating these with the boundary condition given by $V(\cdot,\overline{p}_2)=V^{\sigma}(\cdot,\overline{p}_2)$ we get

$$
\begin{aligned}
V(0,\overline{p}) &= e^{\int_{\overline{p}_2}^{\overline{p}}\frac{\nu(x)}{\eta(x)}dx}V^{\sigma}(0,\overline{p}_2)+e^{\int_{\overline{p}_2}^{\overline{p}}\frac{\nu(x)}{\eta(x)}dx}\int_{\overline{p}_2}^{\overline{p}}e^{-\int_{\overline{p}_2}^{\overline{p}'}\frac{\nu(x)}{\eta(x)}dx}\frac{c-\nu(\overline{p}')u_\ell^R}{\eta(\overline{p}')}d\overline{p}' \\
&= e^{\int_{\overline{p}_2}^{\overline{p}}\frac{\nu(x)}{\eta(x)}dx}V^{\sigma}(0,\overline{p}_2)+\int_{\overline{p}_2}^{\overline{p}}e^{\int_{\overline{p}'}^{\overline{p}}\frac{\nu(x)}{\eta(x)}dx}\frac{c-\nu(\overline{p}')u_\ell^R}{\eta(\overline{p}')}d\overline{p}' \\
&= V^{\sigma}(0,\overline{p}),
\end{aligned}
$$

Hence,

$$
\frac{d\tau(\overline{p},\overline{p}')}{d\overline{p}}=\frac{1}{\lambda\overline{p}(1-\overline{p})} \qquad \text{and} \qquad \frac{d\tau(\overline{p},\overline{p}')}{d\overline{p}'}=-\frac{1}{\lambda\overline{p}(1-\overline{p})}.
$$

and

$$
\begin{aligned}
V(1,\overline{p}) &= e^{\int_{\overline{p}_2}^{\overline{p}} \frac{\lambda+\nu(x)}{\eta(x)}dx} V^\sigma(1,\overline{p}_2) + e^{\int_{\overline{p}_2}^{\overline{p}} \frac{\lambda+\nu(x)}{\eta(x)}dx} \int_{\overline{p}_2}^{\overline{p}} e^{-\int_{\overline{p}_2}^{\overline{p}'} \frac{\lambda+\nu(x)}{\eta(x)}dx} \frac{c-\lambda u_r^R - \nu(\overline{p}')u_\ell^R}{\eta(\overline{p}')} d\overline{p}' \\
&= e^{\int_{\overline{p}_2}^{\overline{p}} \frac{\lambda+\nu(x)}{\eta(x)}dx} V^\sigma(1,\overline{p}_2) + \int_{\overline{p}_2}^{\overline{p}} e^{\int_{\overline{p}'}^{\overline{p}} \frac{\lambda+\nu(x)}{\eta(x)}dx} \frac{c-\lambda u_r^R - \nu(\overline{p}')u_\ell^R}{\eta(\overline{p}')} d\overline{p}' \\
&= V^\sigma(1,\overline{p}).
\end{aligned}
$$

Since $V(p,\overline{p}) = pV(1,\overline{p}) + (1-p)V(0,\overline{p})$ and $V^\sigma(p,\overline{p}) = pV^\sigma(1,\overline{p}) + (1-p)V^\sigma(0,\overline{p})$, we have proven that $V(p,\overline{p}) = V^\sigma(p,\overline{p})$. $\square$

Verification of equilibrium conditions. To verify (11)–(12), note that the definition of $\pi(\overline{p})$ in (15) implies that $U_\ell(\pi(\overline{p})) - V^\sigma(\pi(\overline{p}),\overline{p}) = 0$. We have already observed from Lemma B.3-(iii) that $V^\sigma(\cdot,\overline{p}) \geq \hat{u}$ for all $\overline{p} \in (\overline{p}_2,\overline{p}_3]$. This further implies that $\pi(\overline{p}) \leq \hat{p}$ for $\overline{p} \in (\overline{p}_2,\overline{p}_3]$. Therefore $\rho(\overline{p}) = 0$ is a maximizer in (11)–(12). By (14), $\nu(\overline{p})$ as defined in (16) is a maximizer as well. It remains to show that $m(\overline{p}) = 0$ is a maximizer and the RHS in (11)–(12) is equal to zero, both of which would follow if the coefficient of $(1-m)$ in the objective in (11)–(12) is equal to zero. We can use (14) and (4) to simplify the coefficient of $(1-m)$ to:

$$
-c + \pi(\overline{p})\lambda(u_r^R - \hat{V}(\overline{p})) - \lambda\overline{p}(1-\overline{p})\hat{V}(\overline{p}), \tag{10}
$$

which vanishes precisely because $\hat{V}$ solves ($\widehat{\text{ODE}}$). We thus conclude that $m(\overline{p}) = 0$ is a minimizer in (12). We have thus verified (5) and (12).

To verify (13) note that, with $m(\overline{p}) = \rho(\overline{p}) = 0$ and $\nu(\overline{p})$ defined in (16), the derivative of the objective in (13) with respect to p simplifies to:

$$
\nu(\overline{p})\left(u_\ell^R - u_\ell^L\right) + \lambda(u_r^R - V(\pi(\overline{p}),\overline{p})) = \nu(\overline{p})\left(u_\ell^R - u_\ell^L\right) + \lambda\left(u_r^R - \hat{V}(\overline{p})\right) = 0,
$$

where we have used $V_p(\pi(\overline{p}),\overline{p}) = V_{p\overline{p}}(\pi(\overline{p}),\overline{p}) = 0$. Hence, we conclude that $\pi(\overline{p})$ satisfies (13). It also satisfies (7) since $V^\sigma(p,\overline{p})$ is constant in p.

Region 4: $\overline{p} \in (\overline{p}_3,\overline{p}_4]$.

Here, $\overline{p}_4$ will be defined as part of the analysis.

Computation of equilibrium value. Recall that the strategy σ calls on the DM to experiment fully at each state $\overline{p} > \overline{p}_3$ until state $\overline{p}_3$ is reached. To compute the associated value $V^{\sigma}(p, \overline{p})$, it is useful to consider a (hypothetical) Bayesian DM who would find such a strategy optimal. To this end, we represent the DM's problem as a stopping problem where the stopping payoff is the continuation value $v_{**} := \hat{V}(\overline{p}_3)$:[8]

> *Auxiliary Problem:* Imagine a hypothetical Bayesian DM with any belief $p \geq \underline{p}(\overline{p}_3)$ who at each instant may experiment or stops. She may experiment until either breakthrough occurs or her belief reaches $\underline{p}(\overline{p}_3)$. If she stops at any point, she collects the payoff $\hat{V}(\overline{p}_3)$ (independent of p). The optimal value of this stopping problem is:
>
> $$\Psi(p) = \max_{p' \in [\underline{p}(\overline{p}_3), p]} \phi(p; p', v_{**}). \tag{11}$$
>
> Let p_{**} denote the optimal stopping belief. We note that $\Psi(p_{**}) = v_{**}$ and $\Psi'(p_{**}) \geq 0$, with equality holding whenever $p_{**} > \underline{p}(\overline{p}_3)$ (a consequence of smooth pasting). We can easily see that $p_{**} \in [\underline{p}(\overline{p}_3), \overline{p}_3)$.[9]

Fix any state $\overline{p} > \overline{p}_3$. We ask: *for what belief p is the continuation strategy $\sigma_{\overline{p}}$ optimal over Region 3*? Recall $\sigma_{\overline{p}}$ prescribes: "experiment for the duration of $\tau(\overline{p}, \overline{p}_3)$ and, absent breakthrough by the end of the experimentation, stop and collect $\hat{V}(\overline{p}_3)$," where $\tau(p, p')$ denotes the time it takes for a belief to drift from p to p'. Since, by definition, p_{**} is the optimal stopping belief, the answer to the above question is precisely the belief, $q(\overline{p})$, such that

$$\tau(\overline{p}, \overline{p}_3) = \tau(q(\overline{p}), p_{**}). \tag{12}$$

In words, it is the belief that would be updated to p_{**} after the prescribed duration $\tau(\overline{p}, \overline{p}_3)$ of experimentation.[10] Since p_{**} is the optimal stopping belief in the above Auxiliary Problem, a

[8] Recall that if Region 3 does not exist, then $\overline{p}_3 = \overline{p}_2 = p_*$, so $\hat{V}(\overline{p}_3) = \hat{V}(\overline{p}_2) = \Phi(\overline{p}_2)$, as defined before. All subsequent results hold since $V_p(p, \overline{p}_2) = \Phi'(\overline{p}_2) = \Phi'(p_*) = 0$.

[9] To see that $p_{**} < \overline{p}_3$, it suffices to show that $p_{**} < p_*$ since $p_* \leq \overline{p}_3$. To show $p_{**} < p_*$, suppose otherwise. Then,

$$\lambda p_{**}(1 - p_{**})\Psi'(p_{**}) = \lambda p_{**}(u_r^R - \Psi(p_{**})) - c = \lambda p_{**}(u_r^R - v_{**}) - c > \lambda p_*(u_r^R - \Phi(p_*)) - c = \Psi'(p_*) \geq 0,$$

where the first and second last equalities follow from the fact that both Ψ and Φ solve (2), the strict inequality from $v_{**} < \Phi(p)$ for all p, and the last inequality holds since $\overline{p}_3 \geq p_*$.

[10] Obviously, $q(\overline{p}_3) = p_{**}$.

DM with belief $q(\overline{p})$ finds the prescribed strategy optimal at state $\overline{p}$ and thus will realize the value of $\Psi(q(\overline{p}))$. Consequently,

$$V^{\sigma}(q(\overline{p}), \overline{p}) = \Psi(q(\overline{p})).$$

What about a DM with belief $p \neq q(\overline{p})$? Her value is weakly below $\Psi(q(p))$ and as noted before linear in p. This pins down the value function for belief p at state $\overline{p}$:

$$V^{\sigma}(p, \overline{p}) := \Psi(q(\overline{p})) + (p - q(\overline{p}))\Psi'(q(\overline{p})). \tag{13}$$

We define $\overline{p}_4$ to be such that

$$V^{\sigma}(\underline{p}(\overline{p}_4), \overline{p}_4) = U_r(\underline{p}(\overline{p}_4)). \tag{14}$$

Verification of equilibrium conditions. We first claim that $\pi(\overline{p}) = \underline{p}(\overline{p})$ is a worst-case belief satisfying (7). This follows from the fact, for any $p \in \mathcal{P}(\overline{p}), \overline{p} \geq \overline{p}_3$,

$$V^{\sigma}_p(p, \overline{p}) = \Psi'(q(p)) \geq 0,$$

where the equality follows from (13) and the inequality from the convexity of $\Psi(\cdot)$ and $\Psi'(p_{**}) \geq 0$.

We next verify (12). First, we show that $\nu(\overline{p}) = 0$. To this end, we first observe:

Lemma B.5. *For all $\overline{p} > \overline{p}_3$, $V^{\sigma}(\underline{p}(\overline{p}), \overline{p}) \geq U_\ell(\underline{p}(\overline{p}))$.*

Proof. Note $V^{\sigma}(\underline{p}(\overline{p}_3), \overline{p}_3) = \hat{V}(\overline{p}_3) \geq U_\ell(\underline{p}(\overline{p}_3))$.[11] Next, note that $\underline{p}(\overline{p}_3) > \underline{p}_r^B$ and $V^{\sigma}(\underline{p}(\overline{p}_3), \overline{p}_3) = \phi(\underline{p}(\overline{p}_3); \underline{p}(\overline{p}_3), U_\ell(\underline{p}(\overline{p}_3)))$. Since, by Lemma B.2, $\phi(p; p, U_\ell(p))$ can only cross $U_\ell(p)$ from below, we have

$$\left.\frac{dV^{\sigma}(\underline{p}(\overline{p}), \overline{p})}{d\overline{p}}\right|_{\overline{p}=\overline{p}_3} \geq U_\ell'(\underline{p}(\overline{p}_3)).$$

Observe further $V^{\sigma}(\underline{p}(\overline{p}), \overline{p}) = \phi(\underline{p}(\overline{p}); \underline{p}(\overline{p}_3), U_\ell(\underline{p}(\overline{p}_3)))$ is convex in $\underline{p}(\overline{p})$ whereas $U_\ell(\underline{p}(\overline{p}))$ is linear in $\underline{p}(\overline{p})$. Combining the two facts leads to the desired conclusion. $\square$

Next, we observe that $U_\ell(\underline{p}(\overline{p}_3)) = \hat{V}(\overline{p}_3) > \hat{u}$. This means that

$$V^{\sigma}(\underline{p}(\overline{p}_3), \overline{p}_3) = U_\ell(\underline{p}(\overline{p}_3)) > U_r(\underline{p}(\overline{p}_3)).$$

[11] If $V^{\sigma}(\underline{p}(\overline{p}_2), \overline{p}_2) < U_\ell(\underline{p}(\overline{p}_2))$, then the inequality holds with equality by definition. If $V^{\sigma}(\underline{p}(\overline{p}_2), \overline{p}_2) \geq U_\ell(\underline{p}(\overline{p}_2))$, then the inequality follows from the fact that $\overline{p}_3 = \overline{p}_2$.

By definition of $\overline{p}_4$, we have $V^\sigma(\underline{p}(\overline{p}),\overline{p}) > U_r(\underline{p}(\overline{p}))$ for all $\overline{p} \in (\overline{p}_3, \overline{p}_4)$. Combining this with Lemma B.5, we conclude that $\nu(\overline{p}) = 0$.

Next, we prove $m(\overline{p}) = 0$. Substituting $\nu(\overline{p}) = 0$ and $\pi(\overline{p}) = \underline{p}(\overline{p})$, the objective in (12) becomes

$$m\left[U_\rho(\underline{p}) - V(\underline{p},\overline{p}_-)\right] + (1-m)\left[-c + \lambda\underline{p}(u_r^R - V(\underline{p},\overline{p}_-)) + V_p(\underline{p},\overline{p}_-)\eta(p) + V_{\overline{p}}(\underline{p},\overline{p}_-)\eta(\overline{p})\right],$$

where we suppress the argument of $\underline{p}(\overline{p})$ for notational ease. Note that the coefficient of m is negative (by the same argument as for $\nu = 0$). The coefficient of $(1-m)$ can be written

$$\begin{aligned}
&-c + \lambda\underline{p}(u_r^R - V(\underline{p},\overline{p}_-)) + V_p(\underline{p},\overline{p}_-)\eta(\underline{p}) + V_{\overline{p}}(\underline{p},\overline{p}_-)\eta(\overline{p}) \\
=&-c + \lambda\underline{p}\left(u_r^R - \Psi(q(\overline{p})) - (\underline{p} - q(\overline{p}))\Psi'(q(\overline{p}))\right) \\
&\quad + \eta(\underline{p})\Psi'(q(\overline{p})) + \eta(\overline{p})\left[(\underline{p} - q(\overline{p}))\Psi''(q(\overline{p}))q'(\overline{p})\right] \\
=&-c + \lambda q(\overline{p})\left(u_r^R - \Psi(q(\overline{p}))\right) + \eta(q(\overline{p}))\Psi'(q(\overline{p})) \\
&\quad + \lambda(\underline{p} - q(\overline{p}))\left[u_r^R - \Psi(q(\overline{p})) - (1-q(\overline{p}))\Psi'(q(\overline{p})) - q(\overline{p})(1-q(\overline{p}))\Psi''(q(\overline{p}))\right] \\
=&0,
\end{aligned}$$

where we have used (13) for the first equality and $q'(\overline{p}) = \eta(q(\overline{p}))/\eta(\overline{p})$ for the second equality.[12] The last equality follows since Ψ satisfies (2) and its derivative vanishes. We have shown that the coefficient of m in the objective in (11) and (12) is negative and the coefficient of $(1-m)$ is zero. Therefore, $m = 0$ is a maximizer and (11) holds.

It remains to verify (13). Substituting $m(\overline{p}) = \nu(\overline{p}) = 0$, the objective becomes the same as the coefficient of $1-m$ above, except that $\underline{p}$ is replaced by p. Differentiating the expression with respect to p yields:

$$\lambda\left(u_r^R - \Psi(q(\overline{p})) - (1-q(\overline{p}))\Psi'(q(\overline{p}))\right) - \lambda q(\overline{p})(1-q(\overline{p}))\Psi''(q(\overline{p})) = 0,$$

where the equality follows from the fact that Ψ satisfies (2), so its derivative, which coincides

[12] This holds since $\tau(q(\overline{p}), q(\overline{p}_3)) = \tau(\overline{p}, \overline{p}_3)$ implies that $\ln(q(\overline{p})/(1-q(\overline{p}))) - \ln(p_*/(1-p_*)) = \ln(\overline{p}/(1-\overline{p})) - \ln(\overline{p}_3/(1-\overline{p}_3))$. Hence the difference in log-likelihood ratios of $q(\overline{p})$ and $\overline{p}$ is constant

$$\ln\frac{q(\overline{p})}{1-q(\overline{p})} - \ln\frac{\overline{p}}{1-\overline{p}} = K,$$

for some K. Differentiating this with respect to $\overline{p}$, we obtain $q'(\overline{p}) = \eta(q(\overline{p}))/\eta(\overline{p})$.

with the LHS of the above equation, must vanish.

We have thus shown that the objective in (13) is independent of p and hence the requirement that $p = \underline{p}(\overline{p})$ is a minimizer in (13) is satisfied.

Region 5: $\overline{p} \in [\overline{p}_4, 1]$.

Computation of equilibrium value. Recall that the strategy σ calls for an immediate choice of r for all states $\overline{p} \in [\overline{p}_4, 1]$. It thus immediately follows that the value associated with that strategy is:

$$V^\sigma(p, \overline{p}) := U_r(p), \text{ for all } \overline{p} \in [\overline{p}_4, 1].$$

Verification of equilibrium conditions. Since $U_r(\cdot)$ is increasing, $\pi(\overline{p}) = \underline{p}(\overline{p})$ satisfies (7); given $m(\overline{p}) = \rho(\overline{p}) = 1$, the coefficient of p in the objective of (13) vanishes, so $\pi(\overline{p}) = \underline{p}(\overline{p})$ is also a minimizer in (13).

Substituting $\pi(\overline{p}) = \underline{p}(\overline{p}) = \underline{p}$ and $V^\sigma(p, \overline{p}) = U_r(\overline{p})$ in (11)–(12) we get the following expression for $\overline{p} > \overline{p}_4$

$$\begin{aligned}
&m\left[U_\rho(\underline{p}) - U_r(\underline{p})\right] + (1-m)\left[-c + \lambda \underline{p}\left(u_r^R - U_r(\underline{p})\right) - \lambda \underline{p}(1-\underline{p})U_r'(\underline{p})\right] \\
=&m\left[U_\rho(\underline{p}) - U_r(\underline{p})\right] + (1-m)\left[-c + \lambda \underline{p}\left(U_r(1) - U_r(\underline{p}) - (1-\underline{p})U_r'(\underline{p})\right)\right] \\
=&m\left[U_\rho(\underline{p}) - U_r(\underline{p})\right] + (1-m)(-c),
\end{aligned}$$

where the last line follows since $U_r(\overline{p})$ is linear. Recall that $U_\ell(\underline{p}(\overline{p}_4)) \le V^\sigma(\underline{p}(\overline{p}_4), \overline{p}_4) = U_r(\underline{p}(\overline{p}_4))$, so we note that $U_\ell(\underline{p}) \le U_r(\underline{p})$ for $\underline{p} \le \overline{p}_4$. Therefore, $m(\overline{p}) = \rho(\overline{p}) = 1$ satisfies (12). Substituting these, we also have (11).

B.5.2 Case 1 ($\Phi'(p_*) = 0$) with $c \ge \underline{c}$ and $\Delta > \Delta_c$

In this case, Region 4, as well as its boundaries $\overline{p}_3$ and $\overline{p}_4$, need to be modified. Theorem 1 specifies

$$(m(\overline{p}), \nu(\overline{p}), \rho(\overline{p})) = \left(1, 0, \tfrac{\delta_\ell}{\delta_r + \delta_\ell}\right), \tag{15}$$

and $\pi(\overline{p}) = \hat{p}$, for $\overline{p} \in [\overline{p}_3, \overline{p}_4)$, where $\overline{p}_3$ is now set at $\overline{p}_3'$, which satisfies $\hat{V}(\overline{p}_3') = \hat{u}$. We note that this new $\overline{p}_3'$ is smaller than the original $\overline{p}_3$ and is still larger than $\overline{p}_2$.[13] $\overline{p}_4$ is now set at $\overline{p}_4'$, which

[13] Recall $\hat{V}(\cdot)$ is strictly decreasing. Since $\hat{V}(\overline{p}_3) < \hat{u}$ when $c \in (\underline{c}, \overline{c})$ and $\Delta > \Delta_c$, we have $\overline{p}_3' < \overline{p}_3$. We also have $\overline{p}_3' > \overline{p}_2$ since $\hat{V}(\overline{p}_2) = \Phi(\overline{p}_2) > \max\{U_\ell(\overline{p}_2), U_r(\overline{p}_2)\} \ge \hat{u}$.

uniquely satisfies $\underline{p}(\overline{p}_4') = \hat{p}$. One can see that $\overline{p}_4' > \overline{p}_3'$. Given $\sigma(\overline{p})$ for this region, the value of the strategy is given by $V^\sigma(p, \overline{p}) = \hat{u}$.

It is straightforward to verify that $V^\sigma(p, \overline{p}) = \hat{u}$ together with $\pi(\overline{p}) = \hat{p}$ and (15) satisfy (7)–(13) for $\overline{p} \in (\overline{p}_3, \overline{p}_4)$. Indeed, inserting $\pi(\overline{p}) = \hat{p}$ and $V^\sigma(p, \overline{p}) = \hat{u}$ as well as $V_p^\sigma(p, \overline{p}) = 0$ and $V_{\overline{p}}^\sigma(p, \overline{p}) = 0$, the objective in (11) becomes $-(1-m)c$. Hence $m = 1$ is optimal and (11) holds. Since the objective is independent of ν and ρ, (12) also holds. At $\overline{p} = \overline{p}_3$, we have $V^\sigma(p, \overline{p}_-) = \hat{V}^\sigma(\overline{p}_3) = \hat{u}$ and $V_p^\sigma(p, \overline{p}_-) = 0$, and $V_{\overline{p}}^\sigma(p, \overline{p}_-) < 0$, but we can employ the same argument as in Region 3 to show that the coefficients of m and $(1-m)$, in the objective is zero and hence $m = 1$ is optimal.

Finally, we verify (13). Inserting (15), and $V^\sigma(p, \overline{p}_-) = V^\sigma(p, \overline{p}) = \hat{u}$ in the objective in (13), we see that the objective is equal to zero and hence $\pi(\overline{p}) = \hat{p}$ is a minimizer. Finally, $V_p^\sigma(p, \overline{p}) = 0$ implies that (7) holds.

B.5.3 Solution to $(\widehat{\text{ODE}})$

With the boundary condition $\hat{V}^\sigma(\overline{p}_2) = \Phi(\overline{p}_2)$ we obtain a candidate value function $\hat{V}^\sigma(\overline{p})$.

$$\overline{p}(1-\overline{p})\hat{V}^{\sigma\prime}(\overline{p}) - \frac{\left(u_r^R - \hat{V}^\sigma(\overline{p})\right)\left(u_\ell^L - \hat{V}^\sigma(\overline{p})\right)}{\delta_\ell} + \frac{c}{\lambda} = 0$$

$$\overline{p}(1-\overline{p})\hat{V}^{\sigma\prime}(\overline{p}) - \frac{u_r^R u_\ell^L - \left(u_r^R + u_\ell^L\right)\hat{V}^\sigma(\overline{p}) + \left(\hat{V}^\sigma(\overline{p})\right)^2}{\delta_\ell} + \frac{c}{\lambda} = 0$$

$$\overline{p}(1-\overline{p})\hat{V}^{\sigma\prime}(\overline{p}) + \frac{u_r^R + u_\ell^L}{\delta_\ell}\hat{V}^\sigma(\overline{p}) - \frac{1}{\delta_\ell}\left(\hat{V}^\sigma(\overline{p})\right)^2 + \frac{c}{\lambda} - \frac{u_r^R u_\ell^L}{\delta_\ell} = 0.$$

Letting $\overline{\xi} := \ln \frac{\overline{p}}{1-\overline{p}}$ and $\tilde{V}(\overline{\xi}) = \hat{V}^\sigma(\overline{p})$ we have $\overline{p}(1-\overline{p})\hat{V}^{\sigma\prime}(\overline{p}) = \tilde{V}'(\overline{\xi})$ so that the ODE can be written as

$$\tilde{V}'(\overline{\xi}) + \frac{u_r^R + u_\ell^L}{\delta_\ell}\tilde{V}(\overline{\xi}) - \frac{1}{\delta_\ell}\left(\tilde{V}(\overline{\xi})\right)^2 + \frac{c}{\lambda} - \frac{u_r^R u_\ell^L}{\delta_\ell} = 0$$

$$\tilde{V}'(\overline{\xi}) = \frac{u_r^R u_\ell^L}{\delta_\ell} - \frac{c}{\lambda} - \frac{u_r^R + u_\ell^L}{\delta_\ell}\tilde{V}(\overline{\xi}) + \frac{1}{\delta_\ell}\left(\tilde{V}(\overline{\xi})\right)^2.$$

Let $t(\xi) = \tilde{V}(\xi)/\delta_\ell$. Then we have

$$t'(\xi) = \frac{\tilde{V}'(\xi)}{\delta_\ell} = \frac{u_r^R u_\ell^L}{\delta_\ell^2} - \frac{c}{\lambda\delta_\ell} - \frac{u_r^R + u_\ell^L}{\delta_\ell^2}\tilde{V}(\overline{\xi}) + \left(\frac{\tilde{V}(\overline{\xi})}{\delta_\ell}\right)^2$$

$$
\begin{aligned}
&= \frac{u_r^R u_\ell^L}{\delta_\ell^2} - \frac{c}{\lambda\delta_\ell} - \frac{u_r^R + u_\ell^L}{\delta_\ell} t(\xi) + (t(\xi))^2 \\
t'(\xi) - (t(\xi))^2 &= \frac{u_r^R u_\ell^L}{\delta_\ell^2} - \frac{c}{\lambda\delta_\ell} - \frac{u_r^R + u_\ell^L}{\delta_\ell} t(\xi)
\end{aligned}
$$

. Next suppose there is a function $\varphi(\xi)$ such that

$$
t(\xi) = -\frac{\varphi'(\xi)}{\varphi(\xi)}.
$$

Then

$$
t'(\xi) = -\frac{\varphi''(\xi)\varphi(\xi) - (\varphi'(\xi))^2}{(\varphi(\xi))^2} = -\frac{\varphi''(\xi)}{\varphi(\xi)} + (t(\xi))^2 .
$$

Hence

$$
\begin{aligned}
-\frac{\varphi''(\xi)}{\varphi(\xi)} &= \frac{u_r^R u_\ell^L}{\delta_\ell^2} - \frac{c}{\lambda\delta_\ell} + \frac{u_r^R + u_\ell^L}{\delta_\ell} \frac{\varphi'(\xi)}{\varphi(\xi)} \\
\varphi''(\xi) + \frac{u_r^R + u_\ell^L}{\delta_\ell} \varphi'(\xi) + \left(\frac{u_r^R u_\ell^L}{\delta_\ell^2} - \frac{c}{\lambda\delta_\ell} \right) \varphi(\xi) &= 0
\end{aligned}
$$

. The roots of $a^2 + \frac{u_r^R + u_\ell^L}{\delta_\ell} a + \left(\frac{u_r^R u_\ell^L}{\delta_\ell^2} - \frac{c}{\lambda\delta_\ell} \right)$ are

$$
\begin{aligned}
a_{1,2} &= -\frac{u_r^R + u_\ell^L}{2\delta_\ell} \pm \sqrt{\left(\frac{u_r^R + u_\ell^L}{2\delta_\ell} \right)^2 - \frac{u_r^R u_\ell^L}{\delta_\ell^2} + \frac{c}{\lambda\delta_\ell}} \\
&= -\frac{u_r^R + u_\ell^L}{2\delta_\ell} \pm \sqrt{\left(\frac{u_r^R - u_\ell^L}{2\delta_\ell} \right)^2 + \frac{c}{\lambda\delta_\ell}}.
\end{aligned}
$$

Note that this implies $a_1 = -u_2/\delta_\ell$ and $a_2 = -u_1/\delta_\ell$ where we use the subscript 1 for the lower root in both a_i and u_i. Given that the roots are real we have a general solution

$$
\varphi(\xi) = C_1 e^{a_1 \xi} + C_2 e^{a_2 \xi},
$$

where C_1, C_2 are constants of integration. Setting $C = C_2/C_1$ we get the general solution for $t(\xi)$, $\tilde{V}(\xi)$, and $\hat{V}^\sigma(\overline{p})$:

$$t(\xi) = -\frac{C_1 a_1 e^{a_1 \xi} + C_2 a_2 e^{a_2 \xi}}{C_1 e^{a_1 \xi} + C_2 e^{a_2 \xi}} = -\frac{a_1 e^{a_1 \xi} + C a_2 e^{a_2 \xi}}{e^{a_1 \xi} + C e^{a_2 \xi}} = -\frac{a_1 e^{(a_1 - a_2)\xi} + C a_2}{e^{(a_1 - a_2)\xi} + C}$$

$$\tilde{V}(\xi) = -\delta_\ell \frac{a_1 e^{(a_1 - a_2)\xi} + C a_2}{e^{(a_1 - a_2)\xi} + C}$$

$$\hat{V}^\sigma(\overline{p}) = -\delta_\ell \frac{a_1 \left(\frac{\overline{p}}{1-\overline{p}}\right)^{a_1 - a_2} + C a_2}{\left(\frac{\overline{p}}{1-\overline{p}}\right)^{a_1 - a_2} + C},$$

where

$$a_1 - a_2 = -\frac{1}{\delta_\ell} \sqrt{\left(u_r^R - u_\ell^L\right)^2 + 4\frac{c}{\lambda}\delta_\ell} < 0.$$

We note that

$$\lim_{\overline{p} \to 1} \hat{V}^\sigma(\overline{p}) = -\delta_\ell a_2 = u_1.$$

Boundary condition:

$$\begin{aligned}
&\hat{V}^\sigma(\overline{p}_2) = \Phi(\overline{p}_2) \\
\iff\ & -\delta_\ell \frac{a_1 \left(\frac{\overline{p}}{1-\overline{p}}\right)^{a_1 - a_2} + C a_2}{\left(\frac{\overline{p}}{1-\overline{p}}\right)^{a_1 - a_2} + C} = \Phi(\overline{p}_2) \\
\iff\ & -\delta_\ell \left(a_1 \left(\frac{\overline{p}}{1-\overline{p}}\right)^{a_1 - a_2} + C a_2\right) = \Phi(\overline{p}_2) \left(\left(\frac{\overline{p}}{1-\overline{p}}\right)^{a_1 - a_2} + C\right) \\
\iff\ & \left(\Phi(\overline{p}_2) + \delta_\ell a_2\right) C = -\left(\Phi(\overline{p}_2) + \delta_\ell a_1\right) \left(\frac{\overline{p}}{1-\overline{p}}\right)^{a_1 - a_2} \\
& C = -\frac{\Phi(\overline{p}_2) + \delta_\ell a_1}{\Phi(\overline{p}_2) + \delta_\ell a_2} \left(\frac{\overline{p}_2}{1-\overline{p}_2}\right)^{a_1 - a_2} \\
& \quad = \frac{u_2 - \Phi(\overline{p}_2)}{\Phi(\overline{p}_2) - u_1} \left(\frac{\overline{p}_2}{1-\overline{p}_2}\right)^{a_1 - a_2}.
\end{aligned}$$

B.6 Case 2: $\Phi'(p_*) < 0$

Theorem B.1. *Suppose $c < \bar{c}$ and $\Phi'(p_*) < 0$. For each $c \in (0, \bar{c})$ there is a threshold $\Delta_c \in (0, +\infty]$ with $\Delta_c = +\infty$ if and only if $c \le \underline{c}$ such that the intrapersonal equilibrium is described by:*

(a) If $\Delta \le \Delta_c$, there exist cutoffs $0 < \overline{p}_1 < \overline{p}_2 \le \overline{p}_3 < \overline{p}_4 < 1$ with $\overline{p}_1 = p_\ell^B$ and $\overline{p}_2 = p_*$ such that

$$(m(\overline{p}), \nu(\overline{p}), \rho(\overline{p})) = \begin{cases} (1,0,0) \\ (0,0,\cdot) \\ (m(\overline{p}_2),0,1) \\ (0,\nu^*(\overline{p}),0) \\ (0,0,\cdot) \\ (1,0,1) \end{cases} \qquad \pi(\overline{p}) = \begin{cases} \overline{p} & \text{if } \overline{p} \in [0, \overline{p}_1] \\ \overline{p} & \text{if } \overline{p} \in (\overline{p}_1, \overline{p}_2) \\ \overline{p} & \text{if } \overline{p} = \overline{p}_2 \\ \pi^*(\overline{p}) & \text{if } \overline{p} \in (\overline{p}_2, \overline{p}_3) \\ \underline{p}(\overline{p}) & \text{if } \overline{p} \in [\overline{p}_3, \overline{p}_4) \\ \underline{p}(\overline{p}) & \text{if } \overline{p} \in [\overline{p}_4, 1] \end{cases}$$

with $m(\overline{p}_2) \in (0,1)$, $\nu^*(\overline{p}) > 0$ and $\pi^*(\overline{p}) \in (\underline{p}(\overline{p}), \overline{p})$, for $\overline{p} \in (\overline{p}_2, \overline{p}_3)$, and $\overline{p}_2 < \overline{p}_3$ if and only if $\Phi(p^*) < U_\ell(\underline{p}(p^*))$.

(b) If $\Delta > \Delta_c$, the equilibrium has the same structure as in (a) except that for $\overline{p} \in [\overline{p}_3, \overline{p}_4)$, where

$$(m(\overline{p}), \nu(\overline{p}), \rho(\overline{p})) = (1, 0, \hat{\rho}) \quad \text{and} \quad \pi(\overline{p}) = \hat{p}.$$

In this case, the characterization remains the same except at one state $\overline{p} = \overline{p}_2$. In the main case, recall that $\overline{p}_2$ was part of Region 2. When $\Phi'(p_*) < 0$, Region 2 strategy and verification apply only for $\overline{p} \in (\overline{p}_1, \overline{p}_2)$. (The strategies and verification for all other regions remain valid.) The case of $\overline{p} = \overline{p}_2$ will change as follows.

When $\Phi'(p_*) < 0$, we have $\overline{p}_2 = p_* = p_r^B$, where recall p_r^B is the stopping boundary for the Bayesian DM. Therefore, $\Phi(\overline{p}_2) = U_r(\overline{p}_2)$. The strategy $\sigma(\overline{p}_2)$ specifies $\nu(\overline{p}_2) = 0$ and $\rho(\overline{p}_2) = 1$. We now specify $m(\overline{p}_2)$ below. Given $\nu(\overline{p}_2) = 0$ and $\rho(\overline{p}_2) = 1$, we have

$$\begin{aligned} V^\sigma(p, \overline{p}_2) &= U_r(\overline{p}_2) + (p - \overline{p}_2)\left[m(\overline{p}_2)U_r'(\overline{p}_2) + (1 - m(\overline{p}_2))V_p^\sigma(p, \overline{p}_-^2)\right] \\ &= U_r(\overline{p}_2) + (p - \overline{p}_2)\left[m(\overline{p}_2)U_r'(\overline{p}_2) + (1 - m(\overline{p}_2))\Phi'(\overline{p}_2)\right]. \end{aligned}$$

To understand the first line, note that conditional on stopping, the DM gets a payoff of $U_r(p) =$

$U_r(\overline{p}_2) + (p - \overline{p}_2)U_r'(\overline{p}_2)$, and conditional on not stopping, she gets $V^\sigma(p, \overline{p}_{2-}) = V^\sigma(\overline{p}_{2-}, \overline{p}_{2-}) + (p - \overline{p}_2)V_p^\sigma(p, \overline{p}_{2-}) = U_r(\overline{p}_2) + (p - \overline{p}_2)\Phi'(\overline{p}_2)$.[14]

We set $m(\overline{p}_2)$ so that the term in the square brackets vanishes:

$$m(\overline{p}_2) = \frac{-\Phi'(\overline{p}_2)}{U_r'(\overline{p}_2) - \Phi'(\overline{p}_2)} \in (0,1),$$

where $m(\overline{p}_2) \in (0,1)$ holds since $\Phi'(\overline{p}_2) < 0$. With this definition of $m(\overline{p}_2)$, we have $V_p^\sigma(p, \overline{p}_2) = 0$ so that $\pi(\overline{p}_2) = \overline{p}_2$ is a minimizer in (7).

To verify (11) at $\overline{p}_2$, we substitute $\pi(\overline{p}_2) = \overline{p}_2$, $V(\pi(\overline{p}_2), \overline{p}_{2-}) = \Phi(\overline{p}_2) = U_r(\overline{p}_2)$, $V_p^\sigma(\pi(\overline{p}_2), \overline{p}_{2-}) = \Phi'(\overline{p}_2)$, and $V_{\overline{p}}^\sigma(\pi(\overline{p}_2), \overline{p}_{2-}) = 0$. Since $\overline{p}_2 > \hat{p}$ if $\Phi(\overline{p}_2) = U_r(\overline{p}_2)$, $\rho = 1$ is optimal, and hence the objective is independent of ν, we obtain the following simplified version of (11):

$$\max_m (1-m)\left[-c + p\lambda(u_r^R - \Phi(\overline{p}_2)) + \Phi'(\overline{p}_2)\eta(\overline{p}_2)\right] = 0.$$

The terms inside the square brackets are equal to zero (from (1)), so (11) holds, and $\sigma(\overline{p}_2) = (m(\overline{p}_2), 0, 1)$ is a maximizer of (12). Finally, note that the objective in (13) is

$$\begin{aligned}
&m(\overline{p}_2)\left(U_r(p) - V^\sigma(p, \overline{p}_{2-})\right) \\
&\quad + (1 - m(\overline{p}_2))\left[-c + p\lambda\left(u_r^R - V^\sigma(p, \overline{p}_{2-})\right) + V_p^\sigma(p, \overline{p}_{2-})\eta(p) + V_{\overline{p}}^\sigma(p, \overline{p}_{2-})\eta(\overline{p}_2)\right] \\
=& m(\overline{p}_2)\left(U_r(p) - (\Phi(\overline{p}_2) + (p - \overline{p}_2)\Phi'(\overline{p}_2))\right) - (1 - m(\overline{p}_2))c \\
&\quad + (1 - m(\overline{p}_2))\begin{bmatrix} p\lambda\left(u_r^R - (\Phi(\overline{p}_2) + (p - \overline{p}_2)\Phi'(\overline{p}_2))\right) - \lambda p(1-p)\Phi'(\overline{p}_2) \\ -\lambda\overline{p}_2(1 - \overline{p}_2)(p - \overline{p}_2)\Phi''(\overline{p}_2) \end{bmatrix} \\
=& m(\overline{p}_2)\left(U_r(p) - (\Phi(\overline{p}_2) + (p - \overline{p}_2)\Phi'(\overline{p}_2))\right) - (1 - m(\overline{p}_2))c \\
&\quad + (1 - m(\overline{p}_2))\left[p\lambda\left(u_r^R - \Phi(\overline{p}_2)\right) + p\lambda(1 - \overline{p}_2)\Phi'(\overline{p}_2) - \lambda\overline{p}_2(1 - \overline{p}_2)(p - \overline{p}_2)\Phi''(\overline{p}_2)\right],
\end{aligned}$$

where we have used (3) to obtain the second line. Differentiating this with respect to p yields

$$\begin{aligned}
&m(\overline{p}_2)\left(U_r'(p) - \Phi'(\overline{p}_2)\right) \\
&\quad + (1 - m(\overline{p}_2))\left[\lambda\left(u_r^R - (\Phi(\overline{p}_2) + (1 - \overline{p}_2)\Phi'(\overline{p}_2))\right) - \lambda\overline{p}_2(1 - \overline{p}_2)\Phi''(\overline{p}_2)\right] \\
=& -\Phi'(\overline{p}_2) + (1 - m(\overline{p}_2))\left[\lambda\left(u_r^R - \Phi(\overline{p}_2)\right) + \lambda(1 - \overline{p}_2)\Phi'(\overline{p}_2) - \lambda\overline{p}_2(1 - \overline{p}_2)\Phi''(\overline{p}_2)\right],
\end{aligned}$$

[14] Here we have to distinguish carefully between $V^\sigma(p, \overline{p}_{2-})$, $V_p^\sigma(p, \overline{p}_{2-})$ and $V^\sigma(p, \overline{p}_2)$, $V_p^\sigma(p, \overline{p}_2)$.

where we have used the definition of $m(\overline{p}_2)$ to obtain the second line. Since the derivative of (2) vanishes, the terms inside the brackets vanish. Given $\Phi'(\overline{p}_2) < 0$, $\pi(\overline{p}) = \overline{p}$ is thus the unique maximizer in (13).

To summarize, we have shown that for the posited σ, (7)–(13) hold for $\overline{p} = \overline{p}_2$.

B.7 Uniqueness

Here we prove the intrapersonal equilibrium, $\sigma = (\nu, \mu, \rho)$, together with nature's choice π, described in Theorem 1 and Theorem B.1, is unique for the case $c < \underline{c}$ and $c \geq \overline{c}$. Towards a contradiction, suppose there is a different equilibrium $\tilde{\sigma} = (\tilde{\nu}, \tilde{m}, \tilde{\rho})$, with $\tilde{\pi}$ describing nature's choice. Let

$$\tilde{p} := \inf\{\overline{p} \in (0,1) : \tilde{\sigma}(\overline{p}) \neq \sigma(\overline{p})\}.$$

Assuming $c < \underline{c}$, we will start with Case 1 and move through the different parameter regions, excluding each of them.

1. $\tilde{p} \in [0, \overline{p}_1)$: Recall $m(\overline{p}) = 1$ and $\rho(\overline{p}) = 0$ for all $\overline{p} \in [0, \overline{p}_1)$. Assuming $\tilde{p} < \overline{p}_1$, the admissibility of $\tilde{\sigma}$ means that there must exist an open interval $(\tilde{p}, \tilde{p} + \varepsilon) \subset [0, \overline{p}_1]$ such that $\tilde{m}(\overline{p}) = 0$ for all $\overline{p} \in (\tilde{p}, \tilde{p} + \varepsilon)$.[15] Consider $\overline{p} \in (\tilde{p}, \tilde{p} + \varepsilon)$. Since $V^{\tilde{\sigma}}(p, \overline{p}) < \Phi^*(p)$ for all $p \in [\underline{p}(\overline{p}), \overline{p}]$ and since $\Phi^*(p) \leq U_\ell(p)$ for all $p \leq \overline{p}_1 = p_r^B$, we have $V^{\tilde{\sigma}}(p, \overline{p}) < U_\ell(p), \forall p \in [\underline{p}(\overline{p}), \overline{p}]$ and hence $V^{\tilde{\sigma}}(\tilde{\pi}(\overline{p}), \overline{p}) < U_\ell(\tilde{\pi}(\overline{p}))$, which violates condition (8) of Definition 1.

2. $\tilde{p} \in [\overline{p}_1, \overline{p}_2)$: Recall $m(\overline{p}) = \nu(\overline{p}) = 0$ for all $\overline{p} \in (\overline{p}_1, \overline{p}_2)$.

   (a) Suppose there is some $\varepsilon > 0$ such that $\tilde{m}(\overline{p}) = 0$ for all $\overline{p} \in [\tilde{p}, \tilde{p} + \varepsilon)$. We then have $V^{\tilde{\sigma}}(p, \tilde{p}) = V^{\sigma}(p, \tilde{p})$ for all p, with the value $V^{\tilde{\sigma}}(p, \overline{p})$ being continuous in $\overline{p}$ at $\overline{p} = \tilde{p}$. Since $V_p^{\sigma}(p, \overline{p}) < 0$, it follows $V_p^{\tilde{\sigma}}(p, \overline{p}) < 0$ and hence $\tilde{\pi}(\overline{p}) = \overline{p}$ for $\overline{p} \geq \tilde{p}$ sufficiently close to $\tilde{p}$. Next, since $V^{\sigma}(\overline{p}, \overline{p}) > U_\rho(\overline{p})$ for all $\overline{p} \in (\overline{p}_1, \overline{p}_2)$ and all $\rho \in [0,1]$, the same property holds under strategy $\tilde{\sigma}$ for $\overline{p} \geq \tilde{p}$ sufficiently close to $\tilde{p}$, that is, $V^{\tilde{\sigma}}(\overline{p}, \overline{p}) > U_\rho(\overline{p})$ for all $\rho \in [0,1]$. Condition (9) of Definition 1 then implies $\tilde{\nu}(\overline{p}) = \nu(\overline{p}) = 0$ for all $\overline{p}$ belonging to a right neighborhood of $\tilde{p}$, which is a contradiction.

[15] Our admissibility restrictions rule out the possibilities of $m(\overline{p}) < 1$ on an isolated point as well as $m(\overline{p}) \in (0,1)$ for all $\overline{p}$ belonging to an interval of states.

(b) Suppose next that for every $\varepsilon > 0$ there is some $\overline{p} \in [\tilde{p}, \tilde{p} + \varepsilon)$ such that $\tilde{m}(\overline{p}) > 0$. Since $\tilde{p} < \overline{p}_2 = p_* < \hat{p}$, we have $\overline{p} < \hat{p}$ for all $\overline{p} \geq \tilde{p}$ sufficiently close to $\tilde{p}$. Hence, by condition (10), if $\tilde{m}(\overline{p}) > 0$ for such values of $\overline{p}$, then $\tilde{\rho}(\overline{p}) = 0$. We start by considering the DM's strategy at state $\tilde{p}$. With probability $\tilde{m}(\tilde{p})$ the DM stops and collects payoff $U_\ell(p)$ for some $p \in \mathcal{P}(\tilde{p})$; and with the complementary probability she obtains the continuation payoff of strategy σ. Her value at state $\tilde{p}$ is thus

$$V^{\tilde{\sigma}}(p, \tilde{p}) = \tilde{m}(\tilde{p}) U_{\tilde{\rho}(\tilde{p})}(p) + (1 - \tilde{m}(\tilde{p})) V^{\sigma}(p, \tilde{p}), \tag{16}$$

with $U_{\tilde{\rho}(\tilde{p})}(p) = U_\ell(p)$. Since both $U_\ell(p)$ and $V^\sigma(p, \tilde{p})$ are strictly decreasing in p, we have $\tilde{\pi}(\tilde{p}) = \tilde{p}$ by condition (7). Since $U_\ell(\tilde{p}) < V^\sigma(\tilde{p}, \tilde{p}) = \Phi(\tilde{p})$, $V^{\tilde{\sigma}}(\tilde{\pi}(\tilde{p}), \tilde{p})$ is strictly decreasing in $\tilde{m}(\tilde{p})$ and hence $\tilde{m}(\tilde{p}) = 0$. Since we assumed that for every $\varepsilon > 0$ there exists some $\overline{p} \in (\tilde{p}, \tilde{p} + \varepsilon)$ such that $\tilde{m}(\overline{p}) > 0$ and since we just showed that $\tilde{m}(\tilde{p}) = 0$, the remaining possibility is to have an interval to the right of $\tilde{p}$ on which $\tilde{m}(\overline{p}) = 1$ (and $\tilde{\rho}(\overline{p}) = 0$). For all states $\overline{p}$ belonging to this interval, we thus have $V_p^{\tilde{\sigma}}(p, \overline{p}) = U_\ell'(p) < 0$ and hence $\tilde{\pi}(\overline{p}) = \overline{p}$. The HJB functional then simplifies to

$$G(m, \nu, 0, \tilde{\pi}(\overline{p}), \overline{p}, V^{\tilde{\sigma}}, dV^{\tilde{\sigma}}) = (1 - m)\left(-c + \overline{p}\lambda\delta^R\right).$$

Since $\overline{p} > p_\ell^B = \frac{c}{\lambda\delta^R}$, the functional is strictly decreasing m. Condition (12) thus requires $\tilde{m}(\overline{p}) = 0$ for all $\overline{p} \geq \tilde{p}$ sufficiently close to $\tilde{p}$, which yields a contradiction.

3. $\tilde{p} \in [\overline{p}_2, \overline{p}_3)$: Recall $m(\overline{p}) = 0$, $\nu(\overline{p}) > 0$ and $\rho(\overline{p}) = 0$ for all states $\overline{p}$ in Region 3.

(a) Suppose again there is some $\varepsilon > 0$ such that $\tilde{m}(\overline{p}) = 0$ for all $\overline{p} \in [\tilde{p}, \tilde{p} + \varepsilon)$. For every $\varepsilon' \in (0, \varepsilon)$, we can then find some $\overline{p} \in [\tilde{p}, \tilde{p} + \varepsilon')$ such that $\tilde{\nu}(\overline{p}) \neq \nu(\overline{p})$.[16] As before, we have $V^{\tilde{\sigma}}(p, \tilde{p}) = V^\sigma(p, \tilde{p})$, with $V^{\tilde{\sigma}}(p, \overline{p})$ being continuous in $\overline{p}$ at $\overline{p} = \tilde{p}$. If $\tilde{\nu}(\overline{p}) > 0$ on a right neighborhood of $\tilde{p}$, the same HJB conditions pinning down $\nu(\overline{p})$ must hold. Uniqueness of the solution of these conditions then implies $\tilde{\nu}(\overline{p}) = \nu(\overline{p})$ for all $\overline{p}$ belonging to this neighborhood. We are left with the following possibility: for every $\varepsilon' \in (0, \varepsilon)$, we can find some $\overline{p} \in [\tilde{p}, \tilde{p} + \varepsilon')$ such that $\tilde{\nu}(\overline{p}) = 0$. As $\varepsilon' \to 0$, the probability that R-evidence arrives within time $\tau(\overline{p}, \tilde{p})$ vanishes, so both $V_p^{\tilde{\sigma}}(p, \overline{p})$

[16] The possibility where there is an $\varepsilon' \in (0, \varepsilon)$ such that $\tilde{\nu}(\overline{p}) = \nu(\overline{p})$ for all $\overline{p} \in [\tilde{p}, \tilde{p} + \varepsilon')$, but then $\tilde{\rho}(\overline{p}) \neq \rho(\overline{p}) = 0$ for $\overline{p}$ close or equal to $\tilde{p}$ would clearly violate the HJB conditions.

and $V^{\tilde{\sigma}}_{p\overline{p}}(p,\overline{p})$ vanish (recall that at $\overline{p}=\tilde{p}$, the value segment $\{V^{\sigma}(p,\tilde{p})\}_{p\in\mathcal{P}(\tilde{p})}$ is flat.) Hence, for ε' sufficiently small, the right-hand side of the HJB functional is strictly increasing in p for all $\overline{p}$ such that $\tilde{\nu}(\overline{p})=0$. Nature thus chooses $\tilde{\pi}(\overline{p})=\underline{p}(\overline{p})$. Since however $U_\ell(\underline{p}(\overline{p}))>V^{\tilde{\sigma}}(\underline{p}(\overline{p}),\overline{p})$ for $\overline{p}$ sufficiently close to $\tilde{p}$, this violates condition (9).

(b) Consider next the case where for every $\varepsilon>0$ there is some $\overline{p}\in[\tilde{p},\tilde{p}+\varepsilon)$ such that $\tilde{m}(\overline{p})>0$. As before, we start by showing $\tilde{m}(\tilde{p})=0$. At state $\tilde{p}$ the DM's value is described by (16).

- Suppose $\tilde{\pi}(\tilde{p})>\pi(\tilde{p})$. Then, since $V^{\sigma}(\tilde{\pi}(\tilde{p}),\tilde{p})>U(\tilde{\pi}(\tilde{p}))$, $V^{\tilde{\sigma}}(p,\tilde{p})$ strictly decreases in $\tilde{m}(\tilde{p})$, which implies $\tilde{m}(\tilde{p})=0$.
- Suppose $\tilde{\pi}(\tilde{p})\leq\pi(\tilde{p})$. Since $\pi(\tilde{p})<\hat{p}$, condition (10) implies $\tilde{\rho}(\tilde{p})=0$.[17] For nature to optimally choose $\tilde{\pi}(\tilde{p})<\tilde{p}$, the DM's value $V^{\tilde{\sigma}}(p,\tilde{p})$ must be weakly increasing in p, that is, $V^{\tilde{\sigma}}_p(p,\overline{p})\geq 0$. However, $V^{\sigma}_p(p,\tilde{p})=0$ and $U'_\ell(p)<0$, so $V^{\tilde{\sigma}}_p(p,\overline{p})\geq 0$ requires $\tilde{m}(\tilde{p})=0$, as can be seen from (16).

Having shown $\tilde{m}(\tilde{p})=0$, this leaves the possibility that there is an interval $(\tilde{p},\tilde{p}+\varepsilon],\varepsilon>0$, such that $\tilde{m}(\overline{p})=1$ for all $\overline{p}\in(\tilde{p},\tilde{p}+\varepsilon]$. If $\overline{p}\leq\hat{p}$ for all $\overline{p}$ in that interval, then (10) implies $\tilde{\rho}(\overline{p})=0$, in which case we have $V^{\tilde{\sigma}}(p,\overline{p})=U_\ell(p)$ and hence $\tilde{\pi}(\overline{p})=\overline{p}$. We can then follow the same argument as in point 2(b) to show that $\tilde{m}(\overline{p})=1$ violates the HJB condition (12). Suppose instead $\tilde{p}\geq\hat{p}$. Given $\underline{p}(\tilde{p})<\underline{p}(\overline{p}_3)<\hat{p}$, the saddle point conditions (7) and (10) of Definition 1 imply $\tilde{\rho}(\overline{p})=\hat{\rho}$ and $\tilde{\pi}(\overline{p})=\hat{p}$ for all $\overline{p}>\tilde{p}$ sufficiently close to $\tilde{p}$ such that $\tilde{m}(\overline{p})=1$, with $V^{\tilde{\sigma}}(p,\overline{p})=\hat{u}$ for all $p\in\mathcal{P}(\overline{p})$ as the DM's value. Substituting for nature's choice and $V^{\tilde{\sigma}}$,[18] the HJB functional simplifies to

$$G(m,\nu,0,\tilde{\pi}(\overline{p}),\overline{p},V^{\tilde{\sigma}},dV^{\tilde{\sigma}})=(1-m)\left(-c+\hat{p}\lambda(u^R_r-\hat{u})\right).$$

The coefficient of $(1-m)$ is strictly positive if

$$\hat{p}>\frac{c}{\lambda(u^R_r-\hat{u})}\quad\Longleftrightarrow\quad\frac{\delta^L}{\delta^R+\delta^L}>\frac{c}{\lambda\left(u^R_r-\frac{u^R_r u^L_\ell-u^L_r u^R_\ell}{\delta^R+\delta^L}\right)}$$

[17] Unless $\tilde{m}(\tilde{p})=\tilde{\nu}(\tilde{p})=0$, in which case $\tilde{\rho}(\tilde{p})$ can take an arbitrary value.
[18] Note that, given $V^{\tilde{\sigma}}(p,\overline{p})=\hat{u}$, we have $V^{\tilde{\sigma}}_p(p,\overline{p})=V^{\tilde{\sigma}}_{\overline{p}}(p,\overline{p})=0$.

$$\iff \quad c < \lambda \frac{\delta^L}{\delta^R + \delta^L} \frac{u_r^R(\delta^R + \delta^L) - u_r^R u_\ell^L + u_r^L u_\ell^R}{\delta^R + \delta^L}$$
$$\iff \quad c < \lambda \frac{\delta^L}{\delta^R + \delta^L} \frac{\delta^R \delta_r}{\delta_r + \delta_\ell} = \underline{c}$$

Since the latter inequality is satisfied by assumption, we ruled out the possibility $\tilde{m}(\overline{p}) = 1$ on an interval $(\tilde{p}, \tilde{p} + \varepsilon)$. Taken together, this shows that $\tilde{p}$ cannot belong to $[\overline{p}_2, \overline{p}_3)$.

4. $\tilde{p} \in [\overline{p}_3, \overline{p}_4)$: Recall $m(\overline{p}) = \nu(\overline{p}) = 0$ for all $\overline{p} \in [\overline{p}_3, \overline{p}_4)$. The argument is analogous to that in Point 2. In the current case, $U_{\tilde{\rho}(\tilde{p})}(p) < V^\sigma(p)$ for all $p \in \mathcal{P}(\tilde{p})$ if $\tilde{\rho}(\tilde{p}) \leq \hat{\rho}$. If instead $\tilde{\rho}(\tilde{p}) > \hat{\rho}$, then $\tilde{\pi}(\tilde{p}) = \underline{p}(\tilde{p})$ and the argument follows from $U_{\tilde{\rho}(\tilde{p})}(\underline{p}(\tilde{p})) < V^\sigma(\underline{p}(\tilde{p}), \tilde{p})$.

5. $\tilde{p} \in [\overline{p}_4, 1]$: Recall $m(\overline{p}) = 1$ and $\rho(\overline{p}) = 1$ for all $\overline{p} \in [\overline{p}_4, 1]$. For every $\overline{p} \geq \overline{p}_4$, we have $\underline{p}(\overline{p}) > \hat{p}$ and hence $U_\rho(p) < U_r(p)$ for all $\rho < 1$ and all $p \in \mathcal{P}(\overline{p})$. Condition (12) thus requires $\tilde{\rho}(\overline{p}) = 1$ for all $\overline{p} \geq \tilde{p}$. At $\overline{p} = \overline{p}_4$, $U_r(\underline{p}(\overline{p})) = V^\sigma(\underline{p}(\overline{p}), \overline{p})$, so any $\tilde{m}(\overline{p}_4) \in [0, 1]$ solves the DM's optimization problem. Leaving aside this point of indifference, we want to show that for any $\overline{p} > \overline{p}_4$, $\tilde{m}(\overline{p}) = m(\overline{p}) = 1$ must hold. Assuming this property fails, we must have an open interval $(\tilde{p}, \tilde{p} + \varepsilon), \varepsilon > 0$ such that $\tilde{m}(\overline{p}) = 0$ for all $\overline{p} \in (\tilde{p}, \tilde{p} + \varepsilon)$.[19] If $\tilde{p} > \overline{p}_4$, then the DM always ends up choosing r, whether a breakthrough occurs or not. This strategy is thus dominated by an immediate stopping with action r. Suppose instead $\tilde{p} = \overline{p}_4$ and $\tilde{m}(\overline{p}) = 0$ for all $\overline{p} \in (\overline{p}_4, \overline{p}_4 + \varepsilon)$ and some $\varepsilon > 0$. If $\tilde{\nu}(\overline{p}) = 0$ on a right neighborhood of $\tilde{p}$, the DM's value on that neighborhood is described by

$$V^{\tilde{\sigma}}(p, \overline{p}) = \Psi(q(\overline{p})) + (p - q(\overline{p}))\Psi'(q(\overline{p})),$$

with Ψ and q as defined in (11) and (12). But, by definition of $\overline{p}_4$ in (14), $V^{\tilde{\sigma}}(p, \overline{p}) < U_r(\overline{p})$ holds for all $\overline{p} > \overline{p}_4 = \tilde{p}$ and $p \in \mathcal{P}(\overline{p})$, so the strategy violates condition (9). For interior stopping rates $\tilde{\nu}(\overline{p}) > 0$ on a right neighborhood of $\tilde{p}$, (12) requires $V^{\tilde{\sigma}}(\underline{p}(\overline{p}), \overline{p}) = U_r(\underline{p}(\overline{p}))$ for all $\overline{p}$ belonging to that neighborhood, which in turn requires $\tilde{m}(\overline{p}) = 1$ on that neighborhood, a contradiction.

Case 2. Case 2 is distinguished from Case 1 only by the DM's strategy at state p_* (the boundary between Regions 2 and 3), where the DM mixes instantaneously between experimentation and

[19] See Footnote 15.

action r. The arguments above clearly apply to Regions 1 and 2. They also apply to Regions 3-5 if we can show $\tilde{\sigma}(p_*) = \sigma(p_*)$, so that the DM's value at the left boundary of Region 3 has the same properties as in Case 1. To this end, we distinguish two possibilities based on whether the value segment is upward sloping or downward sloping at p_*.

1. Suppose $\tilde{m}(p_*)$ and $\tilde{\rho}(p_*)$ are such that $V_p^{\tilde{\sigma}}(p, p_*) < 0$. Then $\tilde{\pi}(p_*) = p_*$. Since $\hat{p} < p_*$ in Case 2, we have $U_\rho(p_*) < U_r(p_*)$ for all $\rho < 1$ and hence $\tilde{\rho}(p_*) = 1$. Consider a right neighborhood of p_* and suppose $\tilde{m}(\overline{p}) = 0$ for all $\overline{p}$ belonging to the neighborhood. Under this assumption the value $V^{\tilde{\sigma}}(p, \overline{p})$ is right-continuous in $\overline{p}$ at $\overline{p} = p_*$.[20] For $\overline{p} > p_*$ sufficiently close to p_* we thus have $\tilde{\pi}(\overline{p}) = \overline{p}$. Since stopping and taking action r is strictly Bayesian optimal for all $\overline{p} > p_*$, we then have $U_r(\tilde{\pi}(\overline{p})) > V^{\tilde{\sigma}}(\tilde{\pi}(\overline{p}), \overline{p})$ for all $\overline{p} > p_*$ sufficiently close to p_*, thus violating (9) of Definition 1. Suppose next $\tilde{m}(\overline{p}) = 1$ for all $\overline{p}$ belonging to an interval $(p_*, p_* + \varepsilon), \varepsilon > 0$. We distinguish two cases. If $\underline{p}(p_*) \geq \hat{p}$, then (10) implies $\tilde{\rho}(\overline{p}) = 1$ for all $\overline{p} \in (p_*, p_* + \varepsilon)$ and thus $V_p^{\tilde{\sigma}}(p, p_*) > 0$, a contradiction. If $\underline{p}(p_*) < \hat{p}$, then there exists an $\varepsilon > 0$ such that $\tilde{\rho}(\overline{p}) = \hat{\rho}$ and $\tilde{\pi}(\overline{p}) = \hat{p}$ for all $\overline{p} \in (p_*, p_* + \varepsilon)$. This possibility is ruled out by the argument in point 3(b) above.

2. Suppose $\tilde{m}(p_*)$ and $\tilde{\rho}(p_*)$ are such that $V_p^{\tilde{\sigma}}(p, p_*) > 0$. Nature's choice thus satisfies $\tilde{\pi}(p_*) = \underline{p}(p_*)$. If $\tilde{\pi}(p_*) = \underline{p}(p_*) \leq p_\ell^B$, then, given nature's choice, it is Bayesian optimal to stop and take action ℓ, so we have $\tilde{m}(p_*) = 1$ and $\tilde{\rho}(p_*) = 0$. This implies $V^{\tilde{\sigma}}(p, p_*) = U_\ell(p)$ and thus $V_p^{\tilde{\sigma}}(p, p_*) < 0$, a contradiction. Suppose instead $\tilde{\pi}(p_*) = \underline{p}(p_*) > p_\ell^B$. The DM's value at p_* under strategy $\tilde{\sigma}$ is now described by

$$V^{\tilde{\sigma}}(\tilde{\pi}(p_*), p_*) = \tilde{m}(p_*) U_{\tilde{\rho}(p_*)}(\tilde{\pi}(p_*)) + (1 - \tilde{m}(p_*)) \Phi(\tilde{\pi}(p_*)).$$

   Since $U_{\tilde{\rho}(p_*)}(\tilde{\pi}(p_*)) < \Phi(\tilde{\pi}(p_*))$ for all $\tilde{\rho}(p_*) \in [0, 1]$, this value is strictly decreasing in $\tilde{m}(p_*)$. We thus have $\tilde{m}(p_*) = 0$, which in turn implies $V^{\tilde{\sigma}}(p, p_*) = \Phi(p)$ and thus $V_p^{\tilde{\sigma}}(p, p_*) < 0$, again a contradiction.

Taken together, we have shown that $\tilde{\sigma}$ must be such that

$$V_p^{\tilde{\sigma}}(p, p_*) = \tilde{m}(p_*) U'_{\tilde{\rho}(p_*)}(p) + (1 - \tilde{m}(p_*)) \Phi'(p) = 0. \tag{17}$$

[20]As we approach p_* from the right, the probability of stopping before p_* vanishes.

Since $\Phi'(p) < 0$ for all $p \in \mathcal{P}(p_*)$, this equality requires $\tilde{\rho}(p_*) \in [\hat{\rho}, 1]$ and $\tilde{m}(p_*) \in [m(p_*), 1]$.[21] For the DM to optimally choose $\tilde{m}(p_*) \geq m(p_*)$, it must further hold $U_{\tilde{\rho}(p_*)}(\tilde{\pi}(p_*)) = \Phi(\tilde{\pi}(p_*))$. Since for all $p \in \mathcal{P}(p_*)$ and all $\tilde{\rho}(p_*) \geq \hat{\rho}$, we have $U_{\tilde{\rho}(p_*)}(p) \leq \max\{\hat{u}, U_r(p)\} \leq \Phi(p)$, with the second equality being strict unless $p = p_*$ and the first inequality being strict at $p = p_*$ unless $\tilde{\rho}(p_*) = 1$. Hence, $U_{\tilde{\rho}(p_*)}(\tilde{\pi}(p_*)) = \Phi(\tilde{\pi}(p_*))$ holds if and only if $\tilde{\rho}(p_*) = \rho(p_*) = 1$ and $\tilde{\pi}(p_*) = \pi(p_*) = p_*$. The stopping probability $\tilde{m}(p_*)$ is then pinned down by (17) and equal to $m(p_*)$. We thus have $\tilde{\sigma}(p_*) = \sigma(p_*)$. Given this property, the arguments used to prove uniqueness in Case 1, in particular those regarding Region 3, apply to the current case.

High experimentation costs. When $c \geq \bar{c}$, uniqueness of the intrapersonal equilibrium follows from the argument in the main text following Theorem 1, namely the facts that the commitment solution involves no experimentation and that implementing this solution requires no commitment. Combining these properties with the fact that the commitment solution is unique (see Section B.2) implies that under any alternatively strategy there must be a state at which the DM can profitably deviate to the commitment solution.

B.8 Proof of Proposition 1

Consider first $c < \underline{c}$. The candidate strategy profile has stationary actions $\tilde{\sigma} = (0, \tilde{\nu}, 0)$ and $\pi = \tilde{p} := \frac{u_\ell^L - \tilde{u}}{u_\ell^L - u_\ell^R}$, where $\tilde{u}$ is defined in the statement. The associated value of the stationary strategy is $V^{\tilde{\sigma}}(p) = \tilde{u}$ for all p.[22] We observe that $\tilde{u} > U_r(\hat{p})$ and $\tilde{p} < \hat{p}$ if and only if $c < \underline{c}$. Hence, for $c < \underline{c}$, we have $U_\ell(\tilde{p}) = \tilde{u} \geq U_r(\tilde{p})$. It then follows

$$(0, \tilde{\nu}, 0) \in \arg\max_{(m,\nu,\rho)} m(U_\rho(\tilde{p}) - \tilde{u}) + (1-m)\left[-c + \nu(U_\rho(\tilde{p}) - \tilde{u}) + \lambda\tilde{p}(u_r^R - \tilde{u})\right],$$

which proves that both (11) and (12) are satisfied. Next, given $\tilde{\nu}$, the derivative of

$$\tilde{\nu}(U_\rho(p) - \tilde{u}) + \lambda p(u_r^R - \tilde{u})$$

[21] Recall that $m(p_*)$ is determined by

$$m(p_*)U_r'(p_*) + (1 - m(p_*))\Phi'(p_*) = 0.$$

[22] That this is the correct value of $\tilde{\sigma}$ can be seen by confirming that $\tilde{u} = \frac{\tilde{\nu}U_\ell(\tilde{p}) + \tilde{p}\lambda u_r^R}{\tilde{\nu} + \tilde{p}\lambda}$.

with respect to p vanishes, so

$$\tilde{p} \in \arg\min_p \tilde{m}(U_{\tilde{\rho}}(p) - \tilde{u}) + (1 - \tilde{m}) \left[-c + \tilde{\nu}(U_{\tilde{\rho}}(p) - \tilde{u}) + \lambda p(u_r^R - \tilde{u}) \right],$$

where $(\tilde{m}, \tilde{\rho}) = (0, 0)$. We have thus verified (13).

Consider next $c \geq \underline{c}$. As noted above, we have $\tilde{u} \leq \hat{u}$ in this case. The proof is then identical to that of Section B.5.2, which applies here since $\Delta = \infty$ if $[\underline{p}_0, \overline{p}_0] = [0, 1]$.

To see that the equilibrium is the limit of the equilibrium strategies as $(\underline{p}_0, \overline{p}_0) \to (0, 1)$, note that for $c < \underline{c}$ and $(\underline{p}_0, \overline{p}_0)$ sufficiently close to $(0, 1)$, $\overline{p}_0$ falls into Region 3. Indeed, as $(\underline{p}_0, \overline{p}_0) \to (0, 1)$, $\overline{p}_3$ rises to 1 faster than $\overline{p}_0$ does. To see this, suppose $(\underline{p}_0, \overline{p}_0)$ is sufficiently close to $(0, 1)$ so that $\overline{p}_0 > p_* = \overline{p}_2$ and $\underline{p}_0 < p_\ell^B$. Recall $\overline{p}_3$ is given by $V^\sigma(\underline{p}(\overline{p}_3), \overline{p}) = U_\ell(\underline{p}(\overline{p}_3))$, that is, as the value of $\overline{p}$ for which the left end of the value segment touches U_ℓ. Since $V^\sigma(\underline{p}(\overline{p}_3), \overline{p}_3)$ lies below the Bayesian value function, we clearly have $\underline{p}(\overline{p}_3) \geq p_\ell^B$. Since $\underline{p}_0 = \underline{p}(\overline{p}_0) < p_\ell^B$, $\overline{p}_0$ is strictly smaller than $\overline{p}_3$ (and of course strictly greater than p_*). Given that $\overline{p}_0$ falls into Region 3, it is then routine to check that $(\nu(\overline{p}_0), \pi(\overline{p}_0)) \to (\tilde{\nu}, \tilde{p})$ as $\overline{p}_0 \to 1$. Similarly, when $c \geq \underline{c}$, the same conclusion follows since the condition for Section B.5.2 is satisfied for $\overline{p}_0$ sufficiently close to 1.

B.9 Proof of Proposition 2

In the symmetric case, the Bayesian value function can be written as

$$\Phi(p) = \delta - \frac{c}{\lambda} - (1 - p) \ln \left(\frac{p}{1 - p} \frac{\delta - c/\lambda}{c/\lambda} \right) \frac{c}{\lambda}$$

with derivative

$$\Phi'(p) = \left(\ln \left(\frac{p}{1 - p} \frac{\delta - c/\lambda}{c/\lambda} - \frac{1}{p} \right) \right) \frac{c}{\lambda}.$$

It is easy to see that $\lim_{p \uparrow 1} \Phi'(\overline{p}) = +\infty$. For all states $\overline{p} \in (p_\ell^B, p_*]$ experimentation until the Bayesian update of $\overline{p}$ reaches the Bayesian stopping boundary p_ℓ^B (or R-evidence arrives) is clearly optimal: by the assumption $c < \overline{c}$, the worst-case payoff associated to this strategy, $\Phi(\overline{p})$, is strictly greater than the worst-case payoff for action ℓ, given by $U_\ell(\underline{p}(\overline{p}))$ and strictly greater than the worst-case payoff for action r, given by $U_r(\overline{p})$.

Fixing this part of the strategy and moving backwards in time from $\overline{p} = p_*$, the value segment becomes upward sloping and the experimentation strategy is evaluated at the left-most belief.

The DM's value at this belief lies on the tangency line of Φ touching at $\overline{p}$, given by

$$l(p;\overline{p}) = \Phi(\overline{p}) + \Phi'(\overline{p})(p - \overline{p}).$$

Convexity of Φ implies that the value $l(0,\overline{p})$ is strictly decreasing in $\overline{p}$ with $\lim_{\overline{p}\uparrow 1} l(0;\overline{p}) = -\infty < U_\ell(p)$ for all p. Given $l(0,p_*) = \Phi(p_*) > U_\ell(p_*)$, the difference $U_\ell(\overline{p}) - l(0;\overline{p})$ as a function of $\overline{p}$ must then have an intersection with zero. Notice next that for any given value of $\overline{p}$, we can find a Δ sufficiently large such that $\underline{p}(\overline{p})$ is arbitrarily close to zero. Together with the previous property, this implies that for Δ sufficiently large, there is a state $\overline{p} > p_*$ such that $l(\underline{p}(\overline{p});\overline{p}) = U_\ell(\overline{p})$. Let $\overline{p}_s$ denote the smallest of such states. For all $\overline{p} \in (\pi_\ell^B, \overline{p}_s)$, the DM will then experiment, whereas at $\overline{p} = \overline{p}_s$ she stops to take action ℓ. [23]

What remains to be shown is that there is an $\varepsilon > 0$ sufficiently small such that the DM prefers experimentation over stopping at all states $\overline{p} \in (\overline{p}_s, \overline{p}_s + \varepsilon)$. The payoff from experimentation at $\overline{p} > \overline{p}_s$ evaluated at belief p is[24]

$$p\left(1 - e^{-\lambda\tau(\overline{p},\overline{p}_s)}\right)\delta + (1-p)\delta - c\tau(\overline{p},\overline{p}_s).$$

As $\overline{p} \downarrow \overline{p}_s$, this payoff converges to $(1-p)\delta$. Hence, for $\overline{p}$ sufficiently close to $\overline{p}_s$, the worst-case belief is $\overline{p}$. According to belief $\overline{p}$ experimentation for time $\tau(\overline{p},\overline{p}_s)$ followed by action ℓ is preferred to taking action ℓ immediately if and only if

$$\overline{p}\left(1 - e^{-\lambda\tau(\overline{p},\overline{p}_s)}\right)\delta - c\tau(\overline{p},\overline{p}_s) > 0, \tag{18}$$

where

$$\tau(\overline{p},\overline{p}_s) = \frac{1}{\lambda}\ln\left(\frac{\overline{p}}{1-\overline{p}}\frac{1-\overline{p}_s}{\overline{p}_s}\right).$$

It can be easily verified that the left-hand side of (18) approaches zero from above as $\overline{p} \downarrow \overline{p}_s$ if and only if $\overline{p}_s > p_\ell^B$, which is satisfied since $\overline{p}_s > p_* > p_\ell^B$.

[23] For $\underline{p}(\overline{p}_s)$ sufficiently close to zero, we must have $U_\ell(\overline{p}_s) \geq U_r(\underline{p}(\overline{p}_s))$, so action ℓ is indeed optimal.

[24] Recall that $\tau(\overline{p},\overline{p}_s)$ denotes the time is takes for belief $\overline{p} > \overline{p}_s$ to be updated to $\overline{p}_s$ in the absence of a breakthrough.

B.10 Proof of Proposition 3

It is useful to characterize the learning time through the following recursion equation.[25]

$$\lambda\theta(1-\theta)T'_\Delta(\theta) = 1 - T_\Delta(\theta)(\theta\lambda + \nu_\Delta(\theta)),$$

where $\nu_\Delta(\theta)$ is the stopping rate for ambiguity level Δ when the state $\overline{p}$ is such that $\ln\left(\frac{\overline{p}}{1-\overline{p}}\right) = \ln\left(\frac{\theta}{1-\theta}\right) + \frac{\Delta}{2}$. Using the corresponding equation for $\Delta' > \Delta$, one can write:

$$\lambda\theta(1-\theta)[T'_\Delta(\theta) - T'_{\Delta'}(\theta)] = T_{\Delta'}(\theta)(\theta\lambda + \nu_{\Delta'}(\theta)) - T_\Delta(\theta)(\theta\lambda + \nu_\Delta(\theta)).$$

For any θ such that $\ln\left(\frac{\theta}{1-\theta}\right) + \frac{\Delta'}{2} \leq p_*$, we have $\nu_{\Delta'}(\theta) = \nu_\Delta(\theta) = 0$, and $T_{\Delta'}(\theta) \geq T_\Delta(\theta)$, with the inequality being strict whenever $T_{\Delta'}(\theta) > 0$. This proves that there exists a $\hat{\theta}$ such that $T_{\Delta'}(\theta) \geq T_\Delta(\theta)$ if $\theta < \hat{\theta}$. The above equation implies that for such θ, we have $T'_{\Delta'}(\theta) \leq T'_\Delta(\theta)$.

To see the only if part, observe first that $\nu_\Delta(\theta) \leq \nu_{\Delta'}(\theta)$ for all θ, which follows from the fact that $\ln\left(\frac{\theta}{1-\theta}\right) + \frac{\Delta'}{2} > \ln\left(\frac{\theta}{1-\theta}\right) + \frac{\Delta}{2}$ and from the equation determining ν in Region 3. This means that whenever $T_{\Delta'}(\theta) = T_\Delta(\theta)$, $T'_\Delta(\theta) \geq T'_{\Delta'}(\theta)$; namely, $T_{\Delta'}$ crosses T_Δ at most once and from above.

B.11 Proof of Proposition 4

We indicate by $\overline{p}_n$ and $\overline{q}_n$ for $n = 1, ..., 4$ the boundaries of the different regions when the initial sets of priors are, respectively, $\mathcal{P}$ and $\mathcal{Q}$. Let $\bar{t}_\mathcal{P}$ denote the supremum of Supp$(F_\mathcal{P})$, i.e., the latest time the DM stops.

Since $F_\mathcal{P}(t) = 1$ for all $t > \bar{t}_\mathcal{P}$, the property $F_\mathcal{Q}(t) \leq F_\mathcal{P}(t)$ for all $t > \bar{t}_\mathcal{P}$ is trivially satisfied. What remains to be shown is thus $F_\mathcal{Q}(t) \geq F_\mathcal{P}(t)$ for all $t < \bar{t}_\mathcal{P}$. Clearly, if $\mathcal{P}$ is such that the DM stops deterministically, this property is satisfied, as $F_\mathcal{P}(t) = 0$ for all $t < \bar{t}_\mathcal{P}$. We thus focus on the case where the DM randomizes under $\mathcal{P} = [\underline{p}_0, \overline{p}_0]$, which requires $\overline{p}_0 \in (p_*, \overline{p}_4)$ as well as $\Phi(p_*) < U_\ell(\hat{p}(\tau(\overline{p}_0, p_*), \underline{p}_0))$.

[25] This equation is derived as follows. Take a short length $dt > 0$ of time, then

$$\begin{aligned} T(\theta;\Delta) &= (\lambda\theta + \nu_\Delta(\theta))dt \cdot dt + (1 - (\lambda\theta + \nu_\Delta(\theta))dt)(dt + T(\theta_{dt})) + o(dt) \\ &= (\lambda\theta + \nu_\Delta(\theta))dt \cdot dt + (1 - (\lambda\theta + \nu_\Delta(\theta))dt)(dt + T(\theta;\Delta) + T'(\theta;\Delta)\dot{\theta}dt) + o(dt). \end{aligned}$$

Collecting terms, and letting $dt \to 0$ while using $\dot{\theta} = -\lambda\theta(1-\theta)$, we obtain the equation.

We show first $\tau(\overline{p}_0, \overline{p}_3) \geq \tau(\overline{q}_0, \overline{q}_3)$. Towards a contradiction, suppose $\tau(\overline{p}_0, \overline{p}_3) < \tau(\overline{q}_0, \overline{q}_3)$, so the intrapersonal; equilibrium under $\mathcal{P}$ prescribes randomization at time $\tau(\overline{q}_0, \overline{q}_3)$. By Lemma B.3, $\hat{V}(\overline{p})$ is strictly decreasing in $\overline{p}$ in Region 3. Recalling that p^δ denotes the δ update of p,[26] we have

$$U_\ell(\underline{p}_0^{\tau(\overline{q}_0,\overline{q}_3)}) \;\; \geq \;\; \hat{V}(\overline{p}_0^{\tau(\overline{q}_0,\overline{q}_3)}) > \hat{V}(\overline{p}_0^{\tau(\overline{p}_0,\overline{p}_3)}) > \hat{V}(\overline{q}_0^{\tau(\overline{q}_0,\overline{q}_3)}) = U_\ell(\underline{q}_0^{\tau(\overline{q}_0,\overline{q}_3)}).$$

The first inequality follows from the optimality of randomization at time $\tau(\overline{q}_0, \overline{q}_3)$, the second inequality is due to $\tau(\overline{p}_0, \overline{p}_3) < \tau(\overline{q}_0, \overline{q}_3)$, and the third inequality follows from $\hat{V}(\overline{p}_3) > \hat{V}(\overline{q}_3)$ (since $\overline{p}_3 < \overline{q}_3$). Since, however, $U_\ell(\cdot)$ is strictly decreasing and $\underline{p}_0 \geq \overline{q}_0$, $U_\ell(\underline{p}_0^{\tau(\overline{q}_0,\overline{q}_3)}) \leq U_\ell(\underline{q}_0^{\tau(\overline{q}_0,\overline{q}_3)})$ must hold, which yields a contractions. Hence, $\tau(\overline{p}_0, \overline{p}_3) \geq \tau(\overline{q}_0, \overline{q}_3)$.

Next, we prove that for each $t < \bar{t}_{\mathcal{P}}$, the DM stops at a smaller rate when her initial set of priors is $\mathcal{P}$ than when it is $\mathcal{Q}$. For each $t \geq 0$, denote by $\hat{\nu}(t;\mathcal{P}) := \nu(\overline{p}^t;\mathcal{P})$ the DM's stopping rate as a function of time. We want to show for all $t \leq \bar{t}_{\mathcal{P}}$,

$$\hat{\nu}(t;\mathcal{P}) \leq \hat{\nu}(t;\mathcal{Q}). \tag{19}$$

When $t < \tau(\overline{p}_0, \overline{p}_3)$ or $t \in (\tau(\overline{p}_0, \overline{p}_*), \bar{t}_{\mathcal{P}}]$, $\hat{\nu}(t;\mathcal{P}) = 0$, so (19) trivially holds. Considering $t \in [\tau(\overline{p}_0, \overline{p}_3), \tau(\overline{p}_0, p_*)]$, we have

$$\hat{\nu}'(t;\mathcal{Q}) = \nu'(\hat{p}(t,\mathcal{Q}))\frac{d\hat{p}(t,\overline{q}_0)}{dt} < 0, \tag{20}$$

since for all $\overline{p}$ belonging to Region 3, $\nu'(\overline{p}) = \lambda \hat{V}'(\overline{p})/\delta^\ell < 0$ (see (16)). (20) directly implies

$$\hat{\nu}(t;\mathcal{P}) = \nu(\hat{p}(t,\overline{p}_0)) = \nu(\hat{p}(t + \tau(\overline{q}_0,\overline{p}_0),\overline{q}_0)) = \hat{\nu}(t + \tau(\overline{q}_0,\overline{p}_0);\mathcal{Q}) \leq \hat{\nu}(t;\mathcal{Q})$$

and thus (19). Since (19) holds for all $t < \bar{t}_{\mathcal{P}}$, it follows

$$F_{\mathcal{P}}(t) \;\; = \;\; 1 - e^{-\int_0^t \hat{\nu}(u;\mathcal{P})du} \leq 1 - e^{-\int_0^t \hat{\nu}(u;\mathcal{Q})du} = F_{\mathcal{Q}}(t)$$

for all $t < \bar{t}_{\mathcal{P}}$, as desired.

[26] That is, $p^\delta = \frac{pe^{-\lambda\delta}}{pe^{-\lambda\delta}+1-p}$

C General Poisson Learning

In this extension, we consider a general Poisson model introduced by Che and Mierendorff (2019): at each instant, the DM may seek either R-evidence (as before) or L-evidence. One interpretation is that there are two news sources emitting the two types of evidence. The DM may then allocate a share $\alpha \in [0, 1]$ of her attention to the R-evidence news source and a share $1 - \alpha$ to the L-evidence news source, and receive the evidence proportionately at rates $\alpha\lambda$ in state R and at rates $(1-\alpha)\lambda$ in state L. For instance, a theorist may try either to "prove" a theorem (R-evidence), to find a "counter-example" disproving it (L-evidence), or to divide effort between the two endeavors. We assume throughout that the stopping payoffs are symmetric: $u_r^R = u_\ell^L = \delta > 0$ and $u_\ell^R = u_r^L = 0$.

C.1 Bayesian optimal strategy.

Consider the Bayesian optimal strategy, assuming $c < \overline{c} = \frac{\lambda\delta}{2}$ so that experimentation is optimal for some beliefs.[27] The optimal strategy, characterized in Che and Mierendorff (2019), is succinctly described in Figure 1 for two different ranges of learning costs. Suppose first the cost is *intermediate*; i.e., $c \in [\underline{c}^*, \overline{c})$, where $\underline{c}^* := \frac{\lambda\delta}{1+e^2}$. Then, the DM seeks contradictory evidence; namely she chooses $\alpha = 0$ if $p \in (1/2, p_r^B)$ and $\alpha = 1$ if $p \in (p_\ell^B, 1/2)$. See panel (a) of Figure 1.

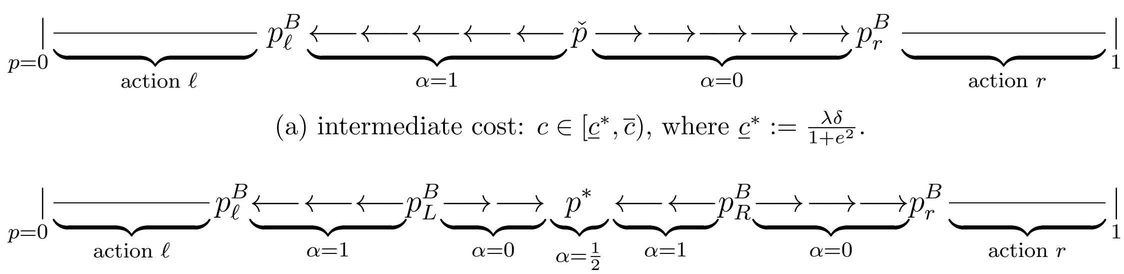

(a) intermediate cost: $c \in [\underline{c}^*, \overline{c})$, where $\underline{c}^* := \frac{\lambda\delta}{1+e^2}$.

(b) low cost: $c < \underline{c}^*$.

Figure 1: Bayesian optimal strategy.

For a *low cost* $c < \underline{c}^*$, the contradictory-evidence seeking is still optimal near the stopping boundaries, as before, but a new learning strategy emerges in the middle region. In that region, the DM seeks *confirmatory evidence*, L-evidence for $p \in (p_L^B, 1/2)$ and R-evidence for $p \in (1/2, p_R^B)$; see panel (b) of Figure 1. When such evidence does not arrive, the DM's belief

[27]Otherwise, the DM simply chooses action $a = r, \ell$ that maximizes $U(p)$.

drifts inwardly toward $p^* = 1/2$. Once $p^* = 1/2$ is reached, the DM splits her attention between the two news sources with $\alpha = 1/2$; her belief then never moves, and learning continues indefinitely until either evidence is obtained. Intuitively, at that belief, the DM finds both types of experimentation equally tempting and acts like a "Buridan's donkey," unable to drift away from the belief that causes the dilemma. The value of this *split-attention* learning can be computed:

$$u^* = \frac{1}{2}u_r^R + \frac{1}{2}u_\ell^L - \frac{2c}{\lambda} = \delta - \frac{2c}{\lambda}.$$

With a little abuse of notation, let $\Phi^*(p)$ denote the value of the Bayesian optimal strategy for belief p. Given the symmetry of payoffs, Φ^* is symmetric around, and attains its minimum at, 1/2. One implication of the above characterization is that $\Phi^*(1/2) = u^*$ for $c \leq \underline{c}^*$ and $\Phi^*(1/2) > u^*$ for $c > \underline{c}^*$, which explains why splitting attention is part of Bayesian optimal strategy if and only if $c < \underline{c}^*$.

C.2 Ambiguity aversion.

Suppose now the DM is ambiguity averse, endowed with an interval of priors. How would the presence of the additional news source affect her behavior? Will see that the solution exhibits several patterns of behavior established from the baseline model as well as some additional features that are linked to the DM's imformation choice in the generalized model.

Low cost or small ambiguity. Consider first the case in which either the cost is low with $c < \underline{c}^*$ or the ambiguity is small. As before, a "small" ambiguity means that $\Phi^*(\frac{1}{2}) \geq U_\ell(\underline{p}(\frac{1}{2}))$, where we recall that $\underline{p}(\overline{p})$ denotes the lowest possible belief when the state is $\overline{p}$. Below, we suppress the dependence of $\underline{p}$ on $\overline{p}$ when no confusion arises. The equilibrium for this case is characterized as follows.

Theorem C.1. *Suppose the cost is low or the ambiguity is small. The intrapersonal equilibrium of the ambiguity-averse DM is as follows.*

- *If $\overline{p} \leq 1/2$, the DM employs the Bayesian optimal strategy of type $\overline{p}$, with the worst-case belief $\pi(\overline{p}) = \overline{p}$.*
- *If $\underline{p} \geq 1/2$, the DM employs the Bayesian optimal strategy of type $\underline{p}$, with the worst-case belief $\pi(\overline{p}) = \underline{p}$.*

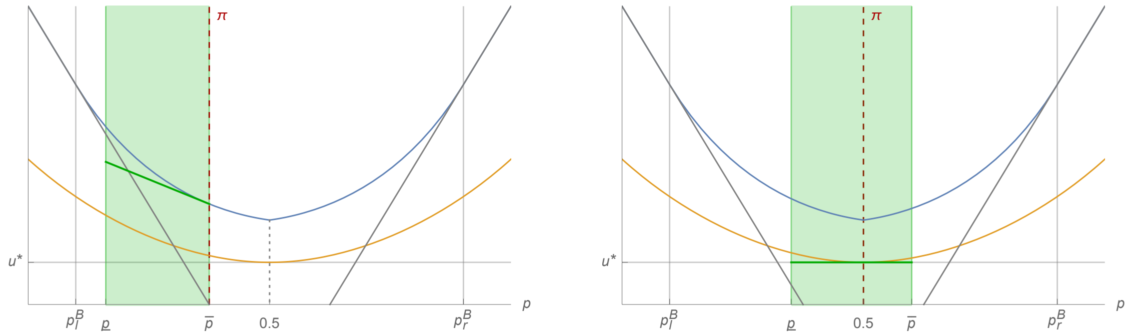

Figure 2: The case of intermediate cost and small ambiguity.

- *If $\underline{p} < 1/2 < \overline{p}$, then the DM employs the split-attention learning strategy or a hedged action, whichever gives a higher payoff, with the worst-case belief $\pi(\overline{p}) = 1/2$.*

Proof. The proof is completely analogous to the baseline case for either $\overline{p} < p_* = 1/2$ or $\underline{p} > p_*$. Specifically, consider the former case. With $\pi(\overline{p}) = \overline{p}$, the DM's HJB coincides with that of the Bayesian DM for belief $\overline{p}$, so the result follows from Che and Mierendorff (2019). For nature's optimality, since the Bayesian value function is decreasing, clearly $\pi(\overline{p}) = \overline{p}$ is optimal. For nature's HJB, its first-order condition with respect to p reduces to the second derivative of the Bayesian HJB, which of course is zero, so the condition holds.

Next, consider the case in which $\underline{p} < 1/2 < \overline{p}$. Nature chooses $\pi(\overline{p}) = 1/2$. If $u^* < \max\{U_\ell(1/2), U_r(1/2)\}$, then randomizing between r and ℓ with equal probabilities satisfies the HJB conditions. Given this, the value function $V(p, \overline{p}) = \hat{u} = \frac{1}{2}\delta$ for all p, so nature does not wish to deviate from $\pi(\overline{p}) = 1/2$. If $u^* \geq \max\{U_\ell(1/2), U_r(1/2)\}$, the HJB conditions hold with $\nu = 0$ and $\alpha = 1/2$, so split-attention learning is optimal. Given this choice, since $V(p, \overline{p}) = u^*$ for all p, nature does not wish to deviate from $\pi(\overline{p}) = 1/2$. □

To explain, suppose first the highest belief $\overline{p}$ is less than $1/2$. Then, as before, the worst-case belief is the right-most belief $\overline{p}$, and the DM adopts the Bayesian optimal strategy for that belief. See the left panel of Figure 2. In case $c \in [\underline{c}^*, \overline{c})$, the worst-case belief $\overline{p}$ drifts leftward, and the experimentation lasts until the belief reaches the left stopping boundary. In other words, Theorem 1 applies precisely in this case with the same implication—including prolonged learning—as before. In case $c < \underline{c}^*$, the belief $\overline{p}$, actually drifts rightward when $\overline{p} > p_L^B$. Once it reaches $p^* = 1/2$, split-attention learning begins, and it lasts until the state is learned.

The case where the lowest belief $\underline{p}$ is greater than $1/2$ is the mirror image of the preceding case. The worst-case belief is now the left-most belief $\underline{p}$, and the DM adopts the Bayesian optimal belief. Equilibrium is precisely the same as above, except that the type of evidence the DM seeks is the opposite to the above case.

Finally, suppose $\underline{p} < 1/2 < \overline{p}$. In this case, the Bayesian optimal strategy for $p \neq 1/2$ cannot occur. To see this, suppose the worst-case belief is $\pi > 1/2$. Then, the value of the Bayesian optimal strategy for belief π—i.e., the value segment tangent to $\Phi^*(\cdot)$ at $p = \pi$—is increasing in the belief, so the worst-case belief is $\underline{p}$, which, however, is strictly less than $1/2$. The worst-case belief cannot be $\underline{p}$ either, since, if so, the value of the Bayesian optimal strategy for $\underline{p}$ is decreasing in the belief, suggesting that the worst-case belief is $\overline{p}$, which now exceeds $1/2$.

This dilemma is resolved by a form of hedging. If $u^* > \hat{u}$, then split-attention learning serves that purpose, with the worst-case belief $\pi = 1/2$. As can be seen in the right panel of Figure 2, the value segment touches the Baysian value function associated with confirmatory evidence seeking (orange curve) and is flat, justifying the interior choice of the worst-case belief. Note that the belief is absorbing since no updating from $p = 1/2$ can occur; in short, once the ambiguity set of beliefs $\mathcal{P}$ contains $1/2$, the set never moves, and the DM is destined forever to split her attention between two news sources.

If the cost is low, then the split-attention learning is Bayesian optimal at belief $p = 1/2$, so the strategy is relatively cheap; this is why it prevails as the dominant form of hedging no matter how large the ambiguity is. If the cost is intermediate with $c \in [\underline{c}^*, \overline{c})$, then the split-attention payoff u^* is strictly below the Bayesian value function, so hedging in this way is relatively costly. In fact, the Bayesian DM would never adopt split-attention learning in this case (see Figure 1-(a)). This explains the discontinuity in the right panel of Figure 2.

Intermediate cost and large ambiguity. Suppose next that the cost is *intermediate*, namely $c \in [\underline{c}^*, \overline{c})$, and the ambiguity is *large* in the sense that $U_\ell(\underline{p}(\frac{1}{2})) > \Phi^*(\frac{1}{2})$. Recall the value $\hat{V}(\overline{p})$ defined in Theorem 1 for the randomized stopping region (Region 3). This value solves $(\widehat{\text{ODE}})$ with the boundary condition $(\overline{p}_2, \Phi^*(\overline{p}_2)) = (\frac{1}{2}, \Phi^*(\frac{1}{2}))$ and is explicitly expressed in (5) (with $\overline{p}_2 = \frac{1}{2}$). Let $p_+ := \hat{V}^{-1}(u^* \vee \hat{u})$ and $p_- := 1 - p_+$. We note that $0 < p_- < 1/2 < p_+ < 1$. The intra-personal equilibrium is then characterized as follows.

Theorem C.2. *Suppose the cost is intermediate, and the ambiguity is large. Then, the following is an equilibrium.*

- *For any $\overline{p} \leq 1/2$, the DM employs the Bayesian optimal strategy for belief $\pi(\overline{p}) = \overline{p}$.*

- *For any $\underline{p} \geq 1/2$, the DM employs the Bayesian optimal strategy for belief $\pi(\overline{p}) = \overline{p}$.*
- *Suppose $[\underline{p}, \overline{p}] \supset [p_-, p_+]$. Then, the DM hedges either by employing the split-attention learning or by mixing between r and ℓ, with the worst-case belief $\pi(\overline{p}) = 1/2$.*
- *Suppose $\underline{p} < 1/2$ and $\overline{p} \in (1/2, p_+]$. Then, the DM randomizes between ℓ and R-evidence seeking. Namely, she stops according to a Poisson rate or else she seeks R-evidence, as characterized in Theorem 1 and Section B.5.*
- *Suppose $\underline{p} \in [p_-, 1/2)$ and $\overline{p} > 1/2$. Then, the DM randomizes between r and L-evidence seeking. Namely, she stops according to a Poisson rate or else she seeks L-evidence, as characterized in Theorem 1 and Section B.5.*[28]

Proof. The proof for all cases, except for the last two are the same as above. Of the latter two, since the case of $\underline{p} \in [p_-, 1/2)$ is an exact mirror image of the case $\overline{p} \in (1/2, p_+]$, we simply focus on that case. This case in turn corresponds precisely to Region 3 of Theorem 1. The construction of $(\nu(\overline{p}), \pi(\overline{p}), \hat{V}(\overline{p}))$ follows exactly without any modification. The verification of HJB is also the same except that $\alpha(\overline{p}) = 1$ needs to be verified now; namely, we additionally need to show that the DM wishes to experiment by confirming state R (rather than state L). To this end, we take the derivative of G with respect to α:

$$\left.\frac{\partial G}{\partial \alpha}\right|_{\alpha=1} = \pi(u_r^R - \hat{V}(\overline{p})) - (1-\pi)(u_\ell^L - \hat{V}(\overline{p})) - 2\hat{V}'(\overline{p})\overline{p}(1-\overline{p}).$$

It suffices to show that this derivative is nonnegative when evaluated at $(\nu(\overline{p}), \pi(\overline{p}), \hat{V}(\overline{p}))$. Indeed, when substituting $(\nu(\overline{p}), \pi(\overline{p}), \hat{V}(\overline{p}))$ from (16), (15), and ($\widehat{\text{ODE}}$), this expression is nonnegative if and only if

$$\hat{V}(\overline{p}) \geq u^*,$$

which holds since $\overline{p} \leq p_+$. □

If $\overline{p} \leq 1/2$ or $\underline{p} \geq 1/2$, then just as before, the DM employs the Bayesian optimal strategy for the inner-most belief, which is to seek contradictory evidence until the most skeptical belief reaches a stopping boundary.

[28]Given the symmetric payoffs, the characterization is a mirror image of the preceding case. More precisely, with the state now indexed by $\underline{p}$, the stopping rate is $\nu^*(1-\underline{p})$ and the worst-case belief is $\pi^*(1-\underline{p})$, where ν^* and π^* are defined respectively in (16) and (15), respectively.

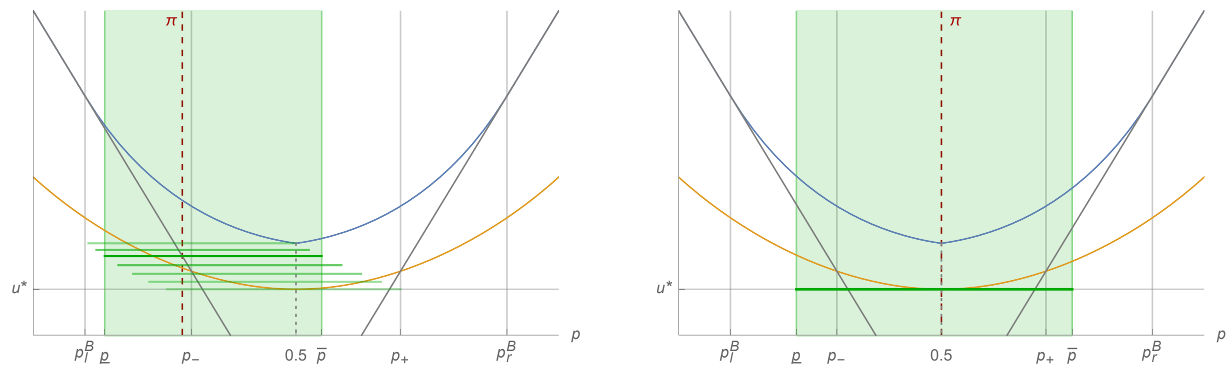

Figure 3: The case of intermediate c and large Δ.

If $\underline{p} < 1/2 < \overline{p}$, the DM employs a hedging strategy. If $[\underline{p}, \overline{p}] \supset [p_-, p_+]$ the DM either adopts split attention learning (if $u^* > \hat{u}$) or randomizes between ℓ and r (if $u^* \leq \hat{u}$). As mentioned above, since the cost is intermediate, split-attention learning involves a strict welfare loss, so the Bayesian DM would never adopt that strategy. See the right panel of Figure 3. If instead $\overline{p} \in (1/2, p_+]$, the DM randomizes between R-evidence seeking and immediate action ℓ. Namely, she stops according to a Poisson rate or else she seeks R-evidence, as characterized in Theorem 1. See the left panel of Figure 3. Symmetrically if $\underline{p} \in [p_-, 1/2)$, then the DM adopts randomized stopping, now, between L-evidence seeking and action r.

In sum, the equilibrium exhibits the behavioral patterns observed in the baseline model, such as *excessive learning* and *premature stopping*, but it also features a new strategy, namely split attention.