
ECONtribute
Discussion Paper No. 194

**Prolonged Learning and Hasty Stopping: the
Wald Problem with Ambiguity**

Sarah Auster

Yeon-Koo Che

Konrad Mierendorff

October 2023 (updated Version)

www.econtribute.de



Prolonged Learning and Hasty Stopping: the Wald Problem with Ambiguity*

Sarah Auster Yeon-Koo Che Konrad Mierendorff[†]

September 2023

Abstract

This paper studies sequential information acquisition by an ambiguity-averse decision maker (DM), who decides how long to collect information before taking an irreversible action. The agent optimizes against the worst-case belief and updates prior by prior. We show that the consideration of ambiguity gives rise to rich dynamics: compared to the Bayesian DM, the DM here tends to experiment excessively when facing modest uncertainty and, to counteract it, may stop experimenting prematurely when facing high uncertainty. In the latter case, the DM's stopping rule is non-monotonic in beliefs and features randomized stopping.

Keywords: Wald problem, ambiguity aversion, prolonged learning, preemptive stopping

JEL Classification Numbers: C61, D81, D83, D91

*We thank David Ahn, Nageeb Ali, Renee Bowen, Martin Cripps, Mark Dean, Francesc Dilmé, Daniel Gottlieb, Will Grimme, Yoram Halevy, Nenad Kos, Stephan Laueremann, Bart Lipman, Georg Noeldeke, Jawwad Noor, Pietro Ortoleva, Wolfgang Pesendorfer, Larry Samuelson, Ludvig Sinander, Philipp Strack, Bruno Strulovici, Tomasz Strzalecki, Juuso Toikka, Yangfan Zhou and audiences in various seminars and conferences for helpful comments and discussions. Sarah Auster gratefully acknowledges funding from the German Research Foundation (DFG) through Germany's Excellence Strategy—EXC 2126/1—390838866 and CRC TR 224 (Project B02).

[†]Auster: Department of Economics, University of Bonn (email: auster.sarah@gmail.com); Che: Department of Economics, Columbia University (email: yeonkooche@gmail.com); Mierendorff: Department of Economics, University College London

1 Introduction

The problem of making a decision on an action after deliberating on its merits is ubiquitous in many situations of life. Examples include a firm contemplating a potential new project, a jury deliberating on a verdict, a citizen pondering over his/her political vote, or a scientist evaluating a hypothesis. The pervasiveness of such a problem makes the framework of central importance, as seen by the extensive research on the subject matter; see Wald (1947), Arrow et al. (1949), Moscarini and Smith (2001), Fudenberg et al. (2018), Che and Mierendorff (2019), among others. With the exception of a few recent papers, the vast literature studying the problem considers a Bayesian framework. The Bayesian model misses, however, a crucial aspect of the problem: the decision maker (DM) often lacks a clear model or beliefs about the uncertainty she is facing. In fact, Wald, often regarded as the first to formally introduce the framework, was mindful of the ambiguity and suggested a robust approach Wald (1947), which anticipates the maxmin decision rule axiomatized by Gilboa and Schmeidler (1989).

We adopt the model in Gilboa and Schmeidler (1989) by considering a DM who has ambiguous beliefs about the state of the world and evaluates her choices against the worst-case scenario in a set of plausible priors. Specifically, the state is either L or R and the DM must decide between two irreversible actions, ℓ and r . At each point, the DM can stop and choose an action. Alternatively, she can postpone her action and learn about the state. In the baseline model, the learning takes the form of a Poisson breakthrough that confirms state R . Imagine, for instance, a CEO of a biotech company seeking to establish the efficacy of a new drug or a theorist seeking to prove a theorem. The Bayesian optimal strategy in such a situation calls for looking for evidence of state R —efficacy of the drug or validity of the proof—for a duration of time until a long stretch of unsuccessful attempts convinces her to declare R to be unlikely and stop for action ℓ .

Introducing ambiguity in a dynamic setting raises conceptual issues about updating and dynamic consistency. To avoid the latter, the existing literature often follows a recursive approach, requiring the DM to adopt a prior-by-prior updating rule on a *rectangular* set of priors (see Epstein and Schneider 2003). Rectangularity requires that the set of priors is rich enough to contain a single worst-case belief that guides the DM’s conditional choices. We follow a different approach: rather than assuming away dynamic inconsistency, we seek to understand its behavioral implications for a “sophisticated” DM who recognizes potential preference reversals and deals with them in a forward-looking and rational manner. To this end, we maintain

prior-by-prior updating but allow for changes in the conditional worst-case scenario.

The paper shows that the consideration of ambiguity and time inconsistency gives rise to rich dynamics in sequential learning. Late in the process, ambiguity increases the DM’s incentives to confirm the state, thereby prolonging experimentation. After a period of unsuccessful search for R -evidence, the DM contemplates two possibilities: no such evidence exists because the state is L or instead the state is R and she has simply been unlucky. As time progresses, the DM becomes increasingly convinced of the former. Following a sufficiently long period of futile search, the Bayesian DM would thus stop to take action ℓ . The ambiguity-averse DM, on the other hand, continues to worry about the second possibility and, as a consequence, keeps looking for more proof. Ambiguity aversion thus causes the DM to continue experimenting well past the time at which the Bayesian counterpart would stop. In the context of our example, the biotech CEO or the theorist refuses to see the “writing on the wall,” and gets hung up on proving the efficacy of the drug or the validity of the theorem. From the perspective of an outsider, the ambiguity-averse DM would thus be seen to exhibit prolonged indecision.

Early on, however, the DM is concerned about the opposite scenario, namely that R is relatively unlikely and the search for evidence will be in vain. From this vantage point, her later self’s refusal to discount R sufficiently enough to stop experimentation in time is seen as misguided. Recognizing the peril of continued indecision, the DM must then intervene and stop before it is too late, namely before the change in the worst-case belief occurs. The DM adopts a mixed strategy in this case, randomizing between experimentation and stopping. Mixing here hedges against ambiguity, as is well-known in the literature. What is novel here is the form it takes; she randomizes in stopping according to a Poisson process at a rate suitably calibrated to hedge against both events. The stopping occurs around the point of belief switch, which is typically well before the Bayesian DM would find it reasonable to stop. Hence, to an outside observer, the DM would be seen to exhibit premature decisiveness.

Comparative statics yields further implications of ambiguity. An increase in ambiguity—as measured by the size of the initial set of priors—increases the tendency for the DM to prolong her experimentation later in the game, which increases her need to preemptively stop earlier in the process and, naturally, the rate at which she does so. These effects, taken together, imply that increased ambiguity “blunts” the DM’s ability to adapt her learning strategy to the true likelihood of states. In particular, she fails to stop learning when learning becomes unprofitable (i.e., when the true likelihood of R falls low), so that given a sufficiently large ambiguity, the

DM’s expected experimentation time *increases* rather than decreases (as would occur under a Bayesian DM) when R becomes unlikely. The combined effects also imply that the stopping time is more dispersed with larger ambiguity.

Finally, we extend the baseline model to the case of incremental learning with diffusion signals. Although some details of the solution change, we show that the main features of the model—excessive learning and preemption via randomized stopping—persist in this case.

In sum, our analysis captures two behavioral traits—protracted indecision and hasty, often unwarranted, decisiveness. The other main contribution of the paper is methodological. Analyzing dynamic decision-making under ambiguity is a priori difficult because the associated time inconsistency problem makes the dynamic programming machinery inapplicable. We characterize the solution to our problem via a set of novel Hamilton-Jacobi-Bellman (HJB) conditions that account for potential differences in preferences of current and future selves. Our HJB characterization thereby restores the dynamic programming logic in the time-inconsistent environment and hence serves as a powerful toolkit that can make analysis tractable without compromising sharp prediction. We expect the approach, and the micro-foundation justifying it, to be valuable beyond our setting.

Related Literature. We incorporate ambiguity into an optimal stopping framework à la [Wald \(1947\)](#). Optimal stopping problems in the Bayesian framework with Poisson signals have been analyzed in [Peskir and Shiryaev \(2006, ch.VI\)](#), with economic applications including [Che and Mierendorff \(2019\)](#) and [Nikandrova and Panes \(2018\)](#). Other applications adopt a drift-diffusion model, e.g. [Moscarini and Smith \(2001\)](#), [Fudenberg et al. \(2018\)](#), while [Zhong \(2019\)](#) studies the problem under flexible choice of information.

The current paper is part of a growing literature on optimal stopping problems under ambiguity. Most of the existing literature adopts a recursive approach, thereby precluding the possibility of changes in the worst-case scenario and time inconsistency; examples include [Epstein and Schneider \(2007\)](#), [Riedel \(2009\)](#), [Miao and Wang \(2011\)](#), and [Cheng and Riedel \(2013\)](#). Closest to our paper in this literature is [Epstein and Ji \(2022\)](#), which studies a Wald problem in a drift-diffusion framework with recursive preferences. Under the recursive approach, there is a single prior in the initial set that minimizes the DM’s expected payoff at every point in time. Due to this feature, standard dynamic programming methods are applicable and the optimal stopping rule continues to be described by two stopping boundaries, as in the Bayesian bench-

mark. By contrast, dynamic inconsistency figures prominently in our model, which we analyze by developing a saddle-point version of the dynamic programming technique.

There is also a growing literature looking at the implication of dynamic inconsistency due to ambiguity in different applications. For example, [Bose and Daripa \(2009\)](#), [Bose and Renou \(2014\)](#), [Ghosh and Liu \(2021\)](#) and [Auster and Kellner \(2022\)](#) consider auctions/mechanisms with ambiguity-averse agents, while [Kellner and Le Quement \(2018\)](#) and [Beauchêne et al. \(2019\)](#) study ambiguity and updating in communication and persuasion settings.

Finally, there is a literature studying dynamic inconsistent stopping in other behavioral frameworks. For instance, [Barberis \(2012\)](#), [Xu and Zhou \(2013\)](#), [Ebert and Strack \(2015, 2018\)](#), and [Henderson et al. \(2017\)](#) analyze stopping problems under prospect theory. [Christensen and Lindensjö \(2018; 2020\)](#) study stopping problems for a class of time-inconsistent models, which can capture endogenous habit formation, non-exponential discounting, and mean-variance utility.

2 Model

A DM must choose between two actions, r and ℓ , whose payoffs depend on an unknown state $\omega \in \{R, L\}$. Specifically, the DM’s payoffs satisfy $u_r^R > u_\ell^R$ and $u_\ell^L > u_r^L$, where $u_a^\omega \in \mathbb{R}$ denotes her payoff from choosing action $a \in \{r, \ell\}$ in state $\omega \in \{R, L\}$. The DM thus wants to “match” the action with the state. We further assume $u_r^R > u_r^L$ and $u_\ell^R < u_\ell^L$. This means that whether a state is desirable or not depends on the action that the DM takes, a feature that makes ambiguity nontrivial.¹ Let $U_a(p) = pu_a^R + (1 - p)u_a^L$ denote the DM’s payoff associated with action a when the probability of state R is p .

The DM may delay her action and acquire information. We model the DM’s information acquisition problem as a stopping problem in continuous time. At each point in time $t \geq 0$, the DM decides whether to wait for more information or whether to stop and take an immediate action, ℓ or r . If the DM waits (or “experiments”), she incurs a flow cost $c > 0$ per unit of time. Payoffs are realized when the DM stops and takes an action. Information takes the form of R -evidence which is received only in state R with arrival rate $\lambda > 0$. Observing news thus conclusively reveals state R , while the absence of news is a signal favoring state L .

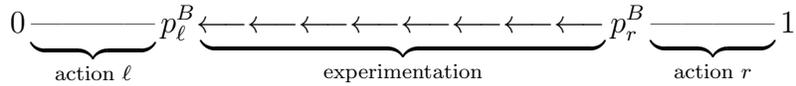
¹If this assumption does not hold, the ambiguity-averse DM minimizing over a set of probabilities is behaviorally indistinguishable from a Bayesian DM, whose prior is equal to the element of the set that maximizes the likelihood of the “bad” state.

Bayesian benchmark. Considering the benchmark case of a Bayesian DM, let p_t denote the probability that the DM assigns to state R at time t . In the absence of news, the DM’s belief drifts “leftward” according to the law of motion:

$$\dot{p}_t = \eta(p_t) := -\lambda p_t(1 - p_t). \quad (1)$$

Intuitively, the DM becomes pessimistic toward R when she looks for R -evidence but is unable to find any.

As is well known, the Bayesian optimal information acquisition rule can be described by two thresholds, p_ℓ^B and p_r^B . If $p_t \leq p_\ell^B$ or $p_t \geq p_r^B$, the DM is sufficiently confident in the state and takes the appropriate action without gathering any additional information. If instead p_t belongs to (p_ℓ^B, p_r^B) , the DM experiments until either she receives R -evidence or the Bayesian update of p_t reaches the threshold p_ℓ^B .



Let $\phi(p; p', U_\ell(p'))$ denote the value to the Bayesian DM of experimenting until the belief p reaches $p' < p$, assuming that if she receives R -evidence during the experimentation, she realizes u_r^R , while if she does not receive such evidence until her belief reaches p' , she stops and takes action ℓ . The standard HJB conditions characterize the DM’s value function as a solution to an ODE:

$$c = \lambda p(u_r^R - \phi(p)) + \phi'(p)\eta(p), \forall p > p'. \quad (2)$$

Intuitively, the flow cost of experimentation (LHS) equals the rate of value increase due to possible breakthrough news and belief updating in its absence (RHS). Optimal stopping calls for p' to be chosen optimally, so the value of the stopping problem is

$$\Phi(p) = \max_{p' \in [0, p]} \phi(p; p', U_\ell(p')), \quad (3)$$

and the left stopping boundary p_ℓ^B is the associated maximizer.² Finally, since the DM has the

²One can characterize the optimal stopping boundary p_ℓ^B by invoking value matching and smooth pasting. Note also that if c is sufficiently large, then the optimal solution is to stop immediately; that is, $p' = p$.

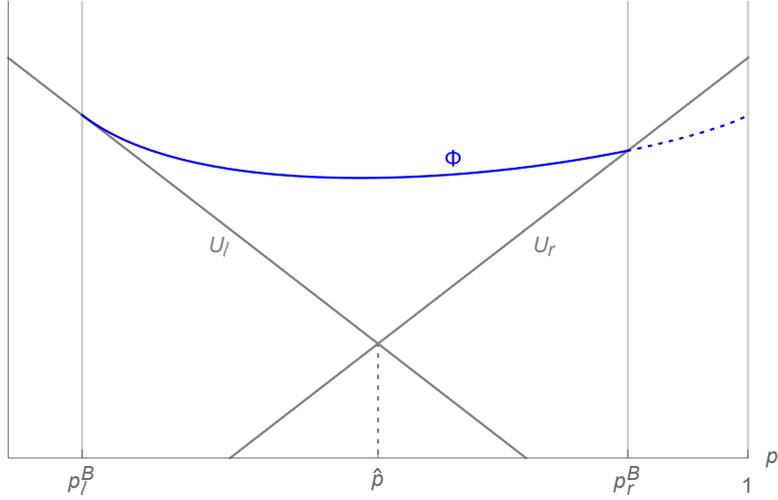


Figure 1: Bayesian optimal stopping rule.

option to choose r at any point in time, the DM with belief p enjoys a payoff of

$$\Phi^*(p) = \max\{\Phi(p), U_r(p)\}, \quad (4)$$

with the DM indifferent between experimentation and action r at the right stopping boundary p_r^B . We will call $\Phi^*(p)$ the *Bayesian value function* throughout. Its characterization is relegated to [Appendix B.1](#).

[Figure 1](#) illustrates graphically the value of the Bayesian DM in a typical situation. The downward sloping line shows the DM's expected payoff from taking action ℓ as a function of the probability p she assigns to state R . Similarly, the upward-sloping line represents the DM's expected payoff from taking action r . The crossing point \hat{p} is the belief at which the DM is indifferent between actions ℓ and r . The blue curve depicts Φ , i.e. the value from experimenting with stopping belief p_ℓ^B . The optimal value Φ^* can be seen as the upper envelope of Φ and $U := U_\ell \vee U_r$. Naturally, the experimentation range (p_ℓ^B, p_r^B) depends on the flow cost c ; the range shrinks as c increases and disappears when

$$c \geq \bar{c} := \lambda \frac{(u_r^R - u_\ell^R)(u_\ell^L - u_r^L)}{(u_r^R - u_\ell^R) + (u_\ell^L - u_r^L)}. \quad (5)$$

The threshold \bar{c} is the value of c at which $\Phi(\hat{p}) = U(\hat{p})$.

Ambiguity and Updating. We assume that the DM faces ambiguity over the likelihood of state R and that her preferences are represented by the maxmin expected utility model (Gilboa and Schmeidler, 1989). According to this model, the DM seeks to maximize her payoff guarantee across a set of priors, described by an interval of probabilities on state R : $\mathcal{P}_0 = [\underline{p}_0, \bar{p}_0]$. We assume $0 < \underline{p}_0 \leq \bar{p}_0 < 1$.

A key question is how the DM updates her ambiguous beliefs as information arrives. We assume that the DM uses *Full Bayesian Updating*: she updates each element of \mathcal{P}_0 according to Bayes rule and considers the worst case over the set of posterior beliefs obtained in this way. The set of posteriors at time $t > 0$ is thus given by an interval $\mathcal{P}_t := [\underline{p}_t, \bar{p}_t]$, where \underline{p}_t is the Bayesian update of the initial lower bound \underline{p}_0 and \bar{p}_t is the Bayesian update of the initial upper bound \bar{p}_0 .

As time progresses and no breakthrough occurs, the interval \mathcal{P}_t drifts to the left. In the process, the interval expands in size as it moves away from the right boundary and shrinks as it approaches the left boundary (ultimately, both \underline{p}_t and \bar{p}_t converge to zero). Rescaling the belief in terms of a log-likelihood ratio keeps the interval “size” constant. Hence, from now on, we shall refer to

$$\Delta := \ln \left(\frac{\bar{p}_0}{1 - \bar{p}_0} \right) - \ln \left(\frac{\underline{p}_0}{1 - \underline{p}_0} \right)$$

as *the degree of ambiguity*. It is easy to verify that Δ does not vary as one updates beliefs over time.

Preference reversals. As is well known, the Full Bayesian Updating rule can lead to violations of dynamic consistency.³ This feature will play a crucial role in our model. To explain the source of dynamic inconsistency in our setting, it will be useful to first consider the DM’s optimal information acquisition rule from the perspective of time 0; we will then show why, after updating the belief set, the DM may have the incentive to deviate from it.

The ex-ante optimal rule solves the DM’s optimization problem in a hypothetical setting where she can commit to a dynamic strategy that maximizes her expected payoff against the worst prior in \mathcal{P}_0 . This optimization problem is clearly independent of the assumed updating

³Beyond the issue of dynamic inconsistency, Full Bayesian Updating, also referred to as *prior-by-prior updating*, has been criticized in the literature for implying a possible situation in which “all news is bad news” (see Gul and Pesendorfer (2021) and Shishkin and Ortoleva (2023)). Since in our setting R -evidence is always good news, this type of situation does not arise here.

rule and thus constitutes an important benchmark. Finding the solution amounts to analyzing a zero-sum game between the DM and adversarial nature choosing $p \in \mathcal{P}_0$ to minimize the DM's expected payoff. Its equilibrium—a saddle point—describes the DM's maxmin optimal value obtained from any information acquisition rule. [Appendix B.2](#) establishes the existence of a saddle point. Given a saddle point, its value also equals the value of the minmax problem in which nature first chooses a prior in \mathcal{P}_0 and the DM then selects the best strategy for the chosen prior (see [Osborne and Rubinstein \(1994\)](#) Proposition 22.2, for example). It follows that the minmax value is simply the lowest Bayesian value Φ^* within \mathcal{P}_0 , or

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

This value coincides with the Bayesian value at the left-most belief \underline{p}_0 if \mathcal{P}_0 is sufficiently far to the right, at the right-most belief \bar{p}_0 if \mathcal{P}_0 is sufficiently far to the left, and at an interior belief if \mathcal{P}_0 contains the global minimizer of Φ^* (see [Figure 2](#) for the first two cases). For later purposes, we denote the minimizer of Φ^* by

$$p_* := \arg \min_{p \in [0,1]} \Phi^*(p).$$

While the maxmin commitment value is easily obtained, the saddle point strategy on the part of the DM requires a little care. As we show in [Appendix B.2](#), if $p_* < p_r^B$ or $p_r^B \notin \mathcal{P}_0$, then the maxmin strategy is indeed the Bayesian optimal strategy for the worst belief in \mathcal{P}_0 . If $p_* = p_r^B \in \mathcal{P}_0$, however, nature's indifference, which is needed for the saddle point requirement, calls for the DM to randomize. If $c < \bar{c}$, then the DM must randomize between r and experimentation. If instead $c \geq \bar{c}$, experimentation is too costly and the DM mixes between actions r and ℓ with probabilities $\hat{\rho}$ and $1 - \hat{\rho}$, respectively, where $\hat{\rho}$ satisfies

$$\hat{\rho}U'_r(p) + (1 - \hat{\rho})U'_\ell(p) = 0. \tag{6}$$

To see now why the described solution may be incompatible with the preferences of the DM's later selves, suppose the interval \mathcal{P}_0 is sufficiently far to the right so that \underline{p}_0 minimizes Φ^* over \mathcal{P}_0 (see the left panel of [Figure 2](#)). If $\underline{p}_0 < p_r^B$, the DM's optimal (commitment) plan at time 0 is to experiment until either she receives R -evidence or the Bayesian update \underline{p}_t of \underline{p}_0 reaches p_ℓ^B . Once the intended stopping point t where $\underline{p}_t = p_\ell^B$ is reached, however, the worst-case belief is

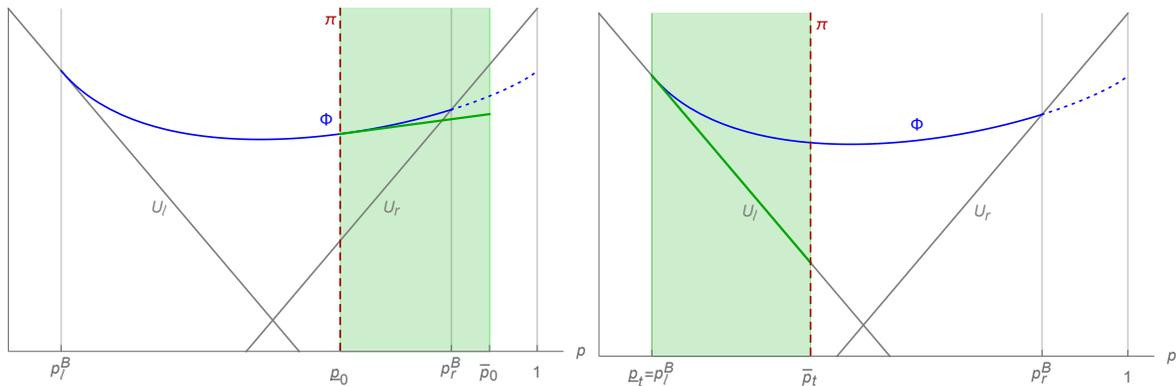


Figure 2: Worst-case scenario (indicated by the red dashed line) at the start of experimentation (left panel) and at the ex-ante optimal stopping time (right panel).

no longer given by \underline{p}_t but instead by \bar{p}_t , as illustrated in the right panel of Figure 2. According to \bar{p}_t , the optimal plan (as prescribed by the commitment solution) is to acquire more information, until the Bayesian update of \bar{p}_t reaches the stopping boundary p_ℓ^B . Hence, later selves of the DM will renege on the original plan prescribed by the commitment solution and *experiment excessively* from the perspective of the original self.

A dynamic inconsistency arises as a result of a switch in the worst-case scenario. When the DM experiments, she is concerned about two events: 1) not receiving R -evidence in state L and 2) taking action ℓ in state R . At the start of the experimentation phase, pessimism about the former event dominates and the worst-case belief is given by \underline{p}_t , which minimizes the likelihood of R . When the DM gets closer to the intended stopping point, however, the concern about taking the wrong action looms large and the worst-case belief switches to \bar{p}_t , i.e. the belief maximizing the likelihood of R . The new worst-case belief calls for more experimentation, thereby creating a conflict with the preferences of earlier selves. A sophisticated DM understands that later selves deviate from the commitment plan by experimenting for too long and optimizes accordingly. We next analyze the behavior of such a DM.

3 Sophisticated Stopping Rule

We now analyze the dynamic behavior of a sophisticated DM. The dynamic inconsistency problem requires us to analyze the DM's stopping problem as an intrapersonal game in which a current self plays against future selves.

3.1 Intrapersonal game.

At each point in time $t \geq 0$, the DM decides between continuing to gather information and stopping, taking as given the choices of her future selves. If a breakthrough occurs, all posteriors in the updated set are equal to one, so it is optimal for the DM to stop and take action r . Taking this part of the strategy as given, it is without loss of generality to focus on the Markov strategy that depends on the set of posteriors rather than time. Indeed, in the absence of R -evidence, there is a one-to-one mapping between calendar time and the set of posteriors. Having fixed Δ as a primitive, we can index the possible sets of posteriors by their upper bound $\bar{p} \in (0, 1)$. The right-most belief in the set, \bar{p} , is thus our state variable. We denote the set of posteriors at state \bar{p} by

$$\mathcal{P}(\bar{p}) := \left\{ p \leq \bar{p} : \ln \left(\frac{\bar{p}}{1 - \bar{p}} \right) - \ln \left(\frac{p}{1 - p} \right) \leq \Delta \right\},$$

and define $\underline{p}(\bar{p}) := \min \mathcal{P}(\bar{p})$ as the lower bound of this set. Clearly, it follows that $\underline{p}_t = \underline{p}(\bar{p}_t)$.

A Markov strategy is a measurable function $\sigma : (0, 1) \rightarrow [0, \infty) \times [0, 1] \times [0, 1]$, with $\sigma(\bar{p}) = (\nu(\bar{p}), m(\bar{p}), \rho(\bar{p}))$. The strategy specifies a *stopping rate* $\nu(\bar{p})$, an *instantaneous stopping probability* $m(\bar{p})$, and the *probability of taking action r conditional on stopping* $\rho(\bar{p})$. Intuitively, densities in the distribution of stopping times are accommodated by interior stopping rates $\nu(\bar{p})$, while atoms in the stopping time distribution are captured by positive values of $m(\bar{p})$.⁴ To ensure that a strategy σ induces a well-defined stopping time, we impose the following *admissibility* restrictions on strategies: (1) every connected set $\mathcal{P} \subseteq [0, 1]$ with $m(\bar{p}) = 1$ for all $\bar{p} \in \mathcal{P}$ contains its supremum; (2) there is a countable set $M^\sigma \subset (0, 1)$ such that $m(\bar{p}) \in (0, 1)$ only if $\bar{p} \in M^\sigma$.⁵

Fixing a strategy σ , we denote by $V^\sigma(p, \bar{p})$ the DM's expected payoff at state \bar{p} when the strategy is evaluated at belief p . The detailed expression of $V^\sigma(p, \bar{p})$ can be found in [Appendix A.1](#); see [\(23\)](#).

Nature's goal is to minimize $V^\sigma(p, \bar{p})$ at each state \bar{p} . Since $V^\sigma(p, \bar{p})$ is a convex combination of the payoffs that strategy σ yields in state R and in state L with coefficient p , the DM's value

⁴One can show that our Markov strategy is equivalent to an alternative formulation that specifies the DM's strategy by a family of stopping times, or more precisely by a family $\{F_t\}_{t \geq 0}$ of distributions satisfying some restrictions such as right-differentiability, where $F_t(\tau)$ is the probability that the DM stops by time $t + \tau$ conditional on not having received any breakthrough news by time t . See [Appendix B.3](#) for details.

⁵These conditions correspond to the requirement that the stopping-time distribution is right-continuous and differentiable at the points of continuities. Specifically, condition (1) ensures that the stopping time is well-defined when the belief drifts toward a (left) stopping boundary. Meanwhile, Condition (2) corresponds to the fact that a distribution function can have at most countably many points of discontinuities.

$V^\sigma(p, \bar{p})$ is linear in the first argument. Nature's problem thus has a corner solution, except when $V_p^\sigma(\cdot, \bar{p}) = 0$. Letting $\pi(\bar{p}) \in \mathcal{P}(\bar{p})$ denote nature's choice at state \bar{p} , we define a solution concept as follows.

Definition 1. A strategy $\sigma = (\nu, m, \rho)$ constitutes a *pseudo equilibrium* if for each state $\bar{p} \in (0, 1)$, there exists some $\pi = \pi(\bar{p})$ with

$$\pi(\bar{p}) \in \arg \min_{p \in \mathcal{P}(\bar{p})} V^\sigma(p, \bar{p}) \quad (7)$$

such that⁶

$$U(\pi) \leq V^\sigma(\pi, \bar{p}); \quad (8)$$

$$U(\pi) = V^\sigma(\pi, \bar{p}) \text{ if } \nu(\bar{p}) \text{ or } m(\bar{p}) > 0; \quad (9)$$

$$\rho(\bar{p}) \in \arg \max_{\rho \in [0,1]} \rho U_r(\pi) + (1 - \rho) U_\ell(\pi). \quad (10)$$

A strategy σ is a pseudo equilibrium if, given nature's choice π , no self of the DM has the incentive to deviate.⁷ Due to the continuous-time nature of our game, many pseudo-equilibria exist. For instance, it is a pseudo equilibrium for the DM to stop at every state and choose the optimal action against the worst-case belief. If some self of the DM were to deviate and refuse to stop, her infinitesimally close future selves would stop immediately. Since experimentation for an instant yields a zero probability of a breakthrough, the deviation is never profitable. However, the prescribed strategy is implausible and simply an artifact of the continuous time nature of our game.

To rule out such implausible strategies, we focus on the pseudo equilibria in which the DM can control her action for a vanishingly small amount of time. Formally, for any pseudo equilibrium σ we imagine a DM who can commit at each time t her action for the period $[t, t + \epsilon)$ against adversarial nature, anticipating that σ will follow after $t + \epsilon$. We then require σ to be obtained as a limit of such ϵ -commitment solutions as $\epsilon \rightarrow 0$.⁸ Such a vanishing control over future behavior

⁶Recall that $U = U_\ell \vee U_r$ is the maximum stopping payoff from actions ℓ and r .

⁷Note that the conditions in [Definition 1](#) and later in [Definition 2](#) make no reference to the initial state \bar{p}_0 . This means that the same strategy should be employed, *regardless of the initial state*. While this is, in principle, a restriction, this restriction entails no loss except for one knife-edge case. See [Footnote 15](#) for details.

⁸This refinement is similar in effect to taking a limit of the behavior in discrete times models as the time interval shrinks to zero; in such models, the DM's experimentation during a unit period has a non-negligible probability of producing a breakthrough. The current refinement accomplishes the same effect within a continuous time

avoids total coordination failure among the different selves. As justified in [Appendix A.2](#) more precisely, the refinement introduces a version of dynamic programming conditions.

To begin, for a value function $V : [0, 1]^2 \rightarrow \mathbb{R}$, define a functional

$$G(m, \nu, \rho, p, \bar{p}, V, dV) = m [U_\rho(p) - V(p, \bar{p}_-)] + (1 - m) [-c + \nu(U_\rho(p) - V(p, \bar{p}_-)) + p\lambda(u_r^R - V(p, \bar{p}_-)) + V_p(p, p_-)\eta(p) + V_{\bar{p}}(p, \bar{p}_-)\eta(\bar{p})],$$

where dV is the gradient of the value function, $V(p, \bar{p}_-) := \lim_{\bar{p}' \uparrow \bar{p}} V(p, \bar{p}')$, and $V_x(p, p_-) := \lim_{\bar{p}' \uparrow \bar{p}} \partial V(p, \bar{p}') / \partial x$.⁹ Recall $\eta(p) = -\lambda p(1 - p)$ denotes the law of motion. The functional G has a familiar interpretation from the dynamic programming literature. Its first term captures a value increase when the DM stops with a point mass m ; she collects $U_\rho(p) - V(p, \bar{p}_-)$. The second term (inside the square brackets) captures a value increase when the DM “continues.” In this case, at the flow cost c , the DM engages in a flow stopping at rate ν and experimentation yielding a breakthrough at rate λp : her value increases by $U_\rho(p) - V(p, \bar{p}_-)$ in the former case and by $u_r^R - V(p, \bar{p}_-)$ in the latter case. The last two terms track the change of value arising from the updating of the state \bar{p} and the updating of the current worst-case belief. This part of the HJB is non-standard and accounts for dynamically-inconsistent belief changes. To compare, in the standard case without ambiguity, the DM’s belief p is the only payoff relevant variable. Hence, the first term $V_p^\sigma(p, \bar{p})$ would be present, while the second-term $V_{\bar{p}}^\sigma(p, \bar{p})$ would not. In the current case, the worst-case belief tomorrow may differ from tomorrow’s update of today’s worst-case belief,¹⁰ making it necessary to account for changes both in the state \bar{p} , which determines the continuation strategy, and in the belief p with which the strategy is evaluated.

For each $\bar{p} \in (0, 1)$, consider a saddle-point version of Hamilton-Jacobi-Bellman equations:¹¹

$$0 = G(m(\bar{p}), \nu(\bar{p}), \rho(\bar{p}), \pi(\bar{p}), \bar{p}, V, dV), \quad (11)$$

framework.

⁹We use the left limits of the value function and its partial derivatives (a) to make sure the conditions are well defined even at points where the derivatives do not exist and (b) since the belief drifts downwards and the value function in the HJB equation reflects the continuation value.

¹⁰Specifically, a dynamic inconsistency arises when $\pi(\bar{p} + \eta(\bar{p})dt) \neq \pi(\bar{p}) + \eta(\pi(\bar{p}))dt$.

¹¹Similar to the commitment solution, a saddle point here characterizes the maxmin optimality of the DM’s strategy (for a vanishingly near future). Equations (11)-(13) resemble Isaac’s condition, which characterizes the equilibrium of dynamic zero-sum games played by two long-run players (see [Bardi and Capuzzo-Dolcetta \(1997\)](#), page 446, and [Laraki and Solan \(2005\)](#)). The fact that our DM is not a time-consistent long-run player differentiates our condition in its function and purpose from Isaac’s condition.

$$\sigma(\bar{p}) = (m(\bar{p}), \nu(\bar{p}), \rho(\bar{p})) \in \arg \max_{m, \nu, \rho} G(m, \nu, \rho, \pi(\bar{p}), \bar{p}, V, dV), \quad (12)$$

$$\pi(\bar{p}) \in \arg \min_{p \in \mathcal{P}(\bar{p})} G(m(\bar{p}), \nu(\bar{p}), \rho(\bar{p}), p, \bar{p}, V, dV). \quad (13)$$

Condition (11) captures the optimality requirement embodied in the value function: the variation from the optimal value for the DM is precisely zero at the equilibrium. Condition (12) requires the DM’s flow decision to be optimal against nature’s choice π , taking her future strategy as given. This is not required in Definition 1, as is seen in the implausible “always stopping” behavior. Finally, (13) chooses the belief to be worst for the DM against the flow action, taking the continuation strategy into consideration. Effectively, equations (11), (12), and (13) represent the saddle-point version of HJB equations. We define our equilibrium as follows:

Definition 2. A pseudo equilibrium $\sigma = (\nu, m, \rho)$, together with π , is an *intrapersonal equilibrium*, or simply an *equilibrium*, if it satisfies (11), (12), and (13) for each $\bar{p} \in (0, 1)$.

The precise microfoundation for this solution concept based on vanishing commitment power is provided in Appendix A.2. It is important to note that the purpose of our HJB equations differs from the standard one, which is to characterize the long-term optimal strategy. Recall that our DM cannot control her future selves, so her strategy is typically not long-term optimal (even given her adversarial belief). Instead, our HJB equations characterize the (vanishingly) near-term maxmin-optimality of her strategy, as explained above, thereby refining the set of potential equilibria. This use of HJB conditions for dynamic decisions under ambiguity, and the microfoundation justifying them, are of independent interest that we believe are applicable beyond the particular setting we consider here. While time inconsistency makes the dynamic programming machinery inapplicable, the HJB characterization restores the dynamic programming logic in a saddle-point and near-term optimality sense. This characterization serves as a powerful toolkit that can make analysis tractable without compromising sharp prediction, as we argue is the case here.

3.2 Main Result

To state our main result, we introduce an additional threshold for the flow cost c :

$$\underline{c} := \frac{\delta_r}{\delta_r + \delta_\ell} \bar{c},$$

where $\delta_a = |u_a^R - u_a^L|$ denotes the payoff difference across the two states given the DM's action $a \in \{\ell, r\}$. As can be easily seen, \underline{c} is strictly smaller than \bar{c} , i.e. the Bayesian threshold defined in Section 2. We will focus on the case in which the Bayesian value function Φ^* attains its minimum at $p_* < p_r^B$ such that $\Phi'(p_*) = 0$, as graphed in Figure 3. This case, labeled *Case 1*, occurs when the value of outright action r is not too high. The opposite case, labeled *Case 2*, in which $p_* = p_r^B$, or $\Phi'(p_*) < 0$ (as graphed in Figure 7) will be treated at the end of this section. This is also where we discuss the conditions under which the equilibrium we describe is unique (and under which it is not).

Theorem 1. *The following intrapersonal equilibrium exists.*

- (i) *Suppose $c \geq \bar{c}$. The DM does not acquire information. If $\bar{p} \leq \hat{p}$, she takes action ℓ ; if $\underline{p}(\bar{p}) \geq \hat{p}$, she takes action r ; if $\hat{p} \in (\underline{p}(\bar{p}), \bar{p})$, she randomizes between ℓ and r with probability \hat{p} .*
- (ii) *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 $c \leq \underline{c}$ such that the intrapersonal equilibrium is described as follows:*

- (a) *If $\Delta \leq \Delta_c$, there exist cutoffs $0 < \bar{p}_1 < \bar{p}_2 \leq \bar{p}_3 < \bar{p}_4 < 1$ with $\bar{p}_1 = p_\ell^B$ and $\bar{p}_2 = p_*$ such that¹²*

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

where $\nu^*(\bar{p}) > 0$ and $\pi^*(\bar{p}) \in (\underline{p}(\bar{p}), \bar{p})$, for $\bar{p} \in (\bar{p}_2, \bar{p}_3)$, and $\bar{p}_2 < \bar{p}_3$ 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*

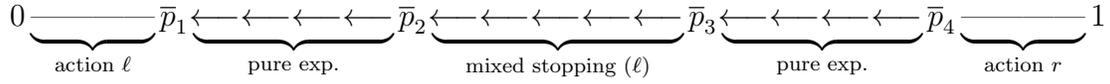
$$(m(\bar{p}), \nu(\bar{p}), \rho(\bar{p})) = (1, 0, \hat{p}) \quad \text{and} \quad \pi(\bar{p}) = \hat{p},$$

¹²We leave ρ unspecified if $m = \nu = 0$.

for $\bar{p} \in [\bar{p}_3, \bar{p}_4)$.

We structure the discussion of [Theorem 1](#) according to the level of experimentation costs. If $c \geq \bar{c}$, the DM does not experiment, and her equilibrium strategy coincides with the commitment solution described in [Proposition B.2](#). This is a consequence of the facts that 1) given $c \geq \bar{c}$, the DM's commitment solution involves no experimentation, as we have seen in [Section 2](#), and 2) implementing this solution requires no commitment. The DM's value in this case is $V^\sigma(\pi(\bar{p}), \bar{p}) = \min_{p \in \mathcal{P}(\bar{p})} U(p)$.

Our main focus will lie on the case $c \leq \underline{c}$. Under this restriction, the equilibrium is always described by [Theorem 1](#)-(ii)(a) and can be illustrated as follows:



When \bar{p} is sufficiently extreme ($\bar{p} \leq \bar{p}_1$ or $\bar{p} \geq \bar{p}_4$), the DM takes immediate action without gathering additional information, just like the Bayesian DM. For the remaining states, however, the characterization differs qualitatively from the benchmark case. In particular, if Δ is sufficiently large, there is a middle region of states where the DM randomizes between experimentation and action ℓ at a positive rate. As a result, the region of pure experimentation is no longer convex.

To derive the solution we leverage the fact that in the absence of breakthrough news, the set of beliefs drifts to the left. Hence, we explain the logic of the equilibrium strategy via backward induction, starting from a far left set of beliefs.

Region 1: $\bar{p} \leq \bar{p}_1$. The first threshold \bar{p}_1 coincides with the Bayesian stopping boundary p_ℓ^B . If $\bar{p} \leq p_\ell^B$, then stopping followed by action ℓ is optimal for all beliefs in $[\underline{p}(\bar{p}), \bar{p}]$ and thus constitutes the equilibrium strategy for our ambiguity-averse DM. The worst-case belief under this strategy is \bar{p} .

Region 2: $\bar{p} \in (\bar{p}_1, \bar{p}_2]$. In Region 2, the right-most belief \bar{p} continues to be the worst-case belief for all $\bar{p} \in (\bar{p}_1, \bar{p}_2]$, and the DM follows the Bayesian optimal stopping rule with respect to \bar{p} . Since $\bar{p} > \bar{p}_1 = p_\ell^B$, this strategy entails experimentation until either a breakthrough occurs or the Bayesian update of \bar{p} reaches p_ℓ^B . The right boundary of this region is $\bar{p}_2 := p_*$, the belief at which the Bayesian value function attains its minimum.

Fix any state $\bar{p} \in (\bar{p}_1, \bar{p}_2]$. We can evaluate the expected payoff $V^\sigma(p, \bar{p})$ of strategy σ

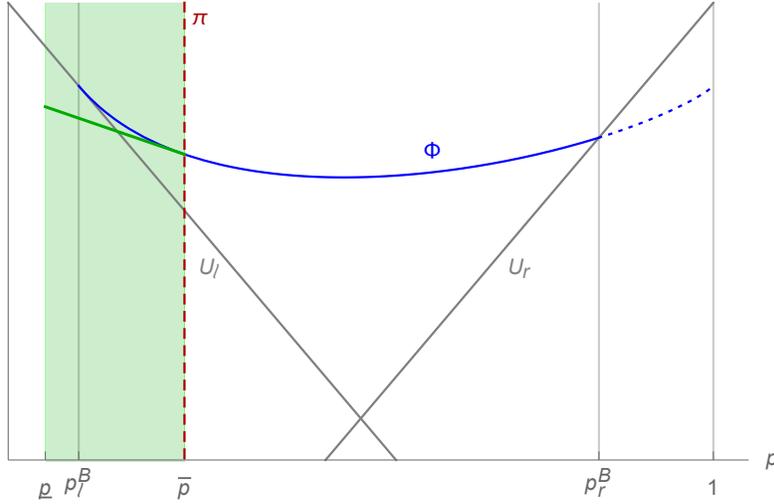


Figure 3: The value segment (in green) and worst-case belief for $\bar{p} \in (\bar{p}_1, \bar{p}_2)$.

as a function of the belief p . The set of values corresponding to possible alternative beliefs, $\{(p, V^\sigma(p, \bar{p})) : p \in \mathcal{P}(\bar{p})\}$, forms a linear segment, which we will call *value segment* throughout (see for instance Figure 3). It is useful to track the movement of the value segment, as time moves backward or forward. Note that the value segment is tangent to the value function Φ at \bar{p} , due to the fact that the stopping rule is Bayesian optimal with respect to \bar{p} .¹³ Tangency at this point implies that the value segment is downward sloping and, hence, that $V^\sigma(p, \bar{p})$ is indeed minimized at $p = \bar{p}$. Hence, $\pi(\bar{p}) = \bar{p}$ is nature's optimal choice against $\sigma(\bar{p})$, and the prescribed $\sigma(\bar{p})$, being Bayesian optimal for $p = \bar{p}$, is a best response against $\pi(\bar{p})$. The profile $(\sigma(\bar{p}), \pi(\bar{p}))$ thus forms a saddle point for state \bar{p} .

At $\bar{p} = p_* =: \bar{p}_2$, the value segment becomes flat (see Figure 4), as we assumed $\Phi'(p_*) = 0$, so all posteriors in $[\underline{p}(\bar{p}), \bar{p}]$ minimize the DM's expected payoff.

Regions 3: $\bar{p} \in (\bar{p}_2, \bar{p}_3]$. Moving further back in time, suppose the state \bar{p} is slightly higher than p_* and the DM contemplates experimenting according to the strategy prescribed in Region 2. The value segment becomes now upward sloping, implying that the worst-case scenario is described by $\underline{p}(\bar{p})$ rather than \bar{p} . We then distinguish two cases. If the degree of ambiguity Δ is sufficiently small and, hence, the value segment is sufficiently “short,” experimentation is

¹³Bayesian optimality of the strategy with respect to \bar{p} clearly implies $V^\sigma(\bar{p}, \bar{p}) = \Phi(\bar{p})$. This, together with linearity of $V^\sigma(p, \bar{p})$ in the first argument and the fact that $\Phi(p) \geq V^\sigma(p, \bar{p})$ holds for all $p \in (p_\ell^B, p_r^B)$, implies that the value segment is tangent to Φ at \bar{p} .

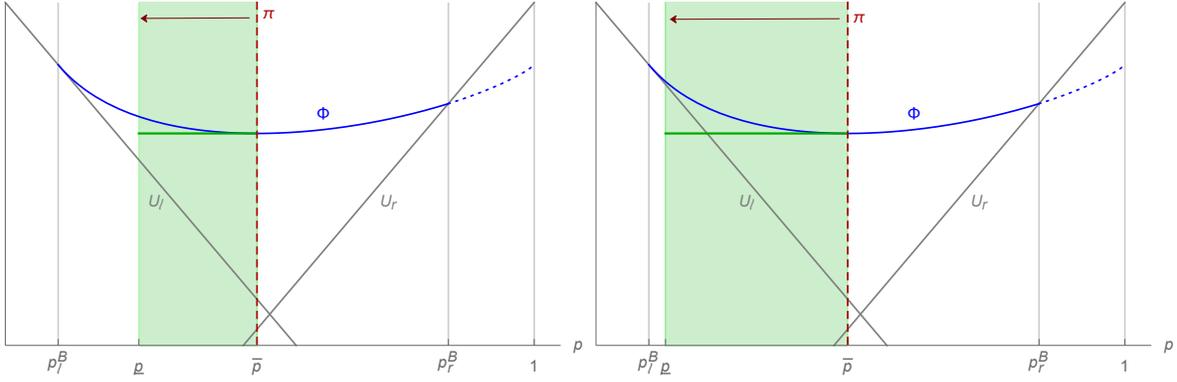


Figure 4: The value segment and worst-case belief for $\bar{p} = \bar{p}_2$ with $V^\sigma(\underline{p}(p_*), p_*) \geq U_\ell(\underline{p}(p_*))$ (left panel) and $V^\sigma(\underline{p}(p_*), p_*) < U_\ell(\underline{p}(p_*))$ (right panel).

optimal at the worst-case belief $\underline{p}(\bar{p})$ (see Figure 4, left panel).¹⁴ The randomization region, Region 3, is empty here, and we directly transition to Region 4 at $\bar{p} = p_*$.

More interestingly, suppose Δ is sufficiently large such that the value segment crosses the stopping payoff line U_ℓ from below. Then, at state $\bar{p} = p_*$ experimentation is dominated by stopping followed by ℓ according to the worst-case belief $\underline{p}(p_*)$ (see Figure 5, right panel). Choosing ℓ , however, cannot be an equilibrium either, as nature would now choose \bar{p} , which makes it optimal for the DM to experiment. The intrapersonal equilibrium is thus in mixed strategies, where the DM randomizes between experimentation and action ℓ .

The DM's value and the equilibrium stopping rate in the mixed stopping region can be derived via our HJB conditions. Randomization between experimentation and action ℓ requires that the coefficient of ν in (12) must vanish and thus:

$$V(\pi(\bar{p}), \bar{p}) = U_\ell(\pi(\bar{p})). \quad (14)$$

Graphically, the stopping payoff U_ℓ and the value segment must cross at the worst-case belief $\pi(\bar{p})$, making the DM indifferent between experimentation and action ℓ . Since the required belief $\pi(\bar{p})$ will be in the interior of $[\underline{p}(\bar{p}), \bar{p}]$, for nature to optimally choose it, the value segment must remain flat so that all posteriors in $[\underline{p}(\bar{p}), \bar{p}]$ minimize the DM's expected payoff. We can thus write $V(p, \bar{p}) = \hat{V}(\bar{p}), \forall p$. Substituting this into the DM's indifference condition (14), we obtain

¹⁴Formally, we have $V^\sigma(\underline{p}(p_*), p_*) \geq U_\ell(\underline{p}(p_*))$.

nature's choice

$$\pi(\bar{p}) = \frac{u_\ell^L - \hat{V}(\bar{p})}{u_\ell^L - u_\ell^R}. \quad (15)$$

The randomization by the DM in turn implies that the derivative of the objective in (12) with respect to p must vanish. This fact, together with $V(p, \bar{p}) = \hat{V}(\bar{p}), \forall p$, yields the equilibrium stopping rate

$$\nu(\bar{p}) = \frac{\lambda(u_r^R - \hat{V}(\bar{p}))}{u_\ell^L - u_\ell^R}. \quad (16)$$

Using (15) and (16), the value function $\hat{V}(\bar{p})$ is derived by substituting for π and ν in the HJB condition (11), which gives rise to the following differential equation:

$$\bar{p}(1 - \bar{p})\hat{V}'(\bar{p}) = \frac{(u_r^R - \hat{V}(\bar{p})) (u_\ell^L - \hat{V}(\bar{p}))}{u_\ell^L - u_\ell^R} - \frac{c}{\lambda}. \quad (\widehat{\text{ODE}})$$

Together with the boundary condition $\hat{V}(\bar{p}_2) = \Phi(\bar{p}_2)$, ($\widehat{\text{ODE}}$) admits a unique solution, describing the DM's value in Region 3.

The reason that the DM stops according to a Poisson process, and thus with vanishing probability (rather than positive probability), is because when $\bar{p} \approx p_*$, the value segment is already arbitrarily close to flat, so it takes a vanishingly small stopping probability to turn the value segment completely flat. The same situation and intuition are repeated as we move backward in time: for a positive range of states, labeled Region 3, the DM repeatedly hedges between experimentation and Poisson stopping at a positive rate $\nu(\bar{p})$.

Intuitively, the repeated hedging by the DM can be seen as a resolution of conflicts between current and future selves. On the one hand, the DM is highly uncertain about the state and thus finds it optimal to acquire information. In particular, the DM worries about action ℓ being the wrong choice (state R) and would thus like to gather more information before taking it. On the other hand, the DM is aware that the journey of experimentation will take on its own life, with the future selves doing too much of it from the current self's perspective. Prolonged experimentation is particularly detrimental in state L where the costly search for evidence is futile. The randomized stopping hedges the DM against both events: at each state in the mixing

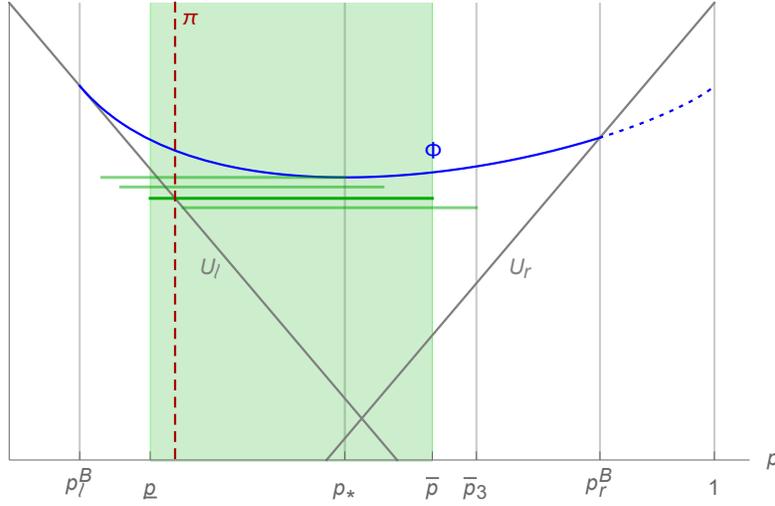


Figure 5: The value segment and worst-case belief for $\bar{p} \in (\bar{p}_2, \bar{p}_3)$.

region, the DM's continuation value is independent of the underlying state. This resolves the conflict and restores the delicate balance between the multiple selves.

While at first glance it may seem odd that the randomized stopping arises in the middle region around p_* , this is no coincidence. As noted earlier, it is precisely where preference reversal and dynamic inconsistency arise. As will be seen, this seemingly peculiar feature of the equilibrium proves quite robust across different specifications of the learning environment and technologies.

Regions 4: $\bar{p} \in (\bar{p}_3, \bar{p}_4]$. As already mentioned, when ambiguity Δ is small, Region 3 is empty, and Regions 2 and 4 collapse to a single experimentation region. Its right boundary, \bar{p}_4 , occurs at the point where, according to the worst-case belief $\underline{p}(\bar{p})$, action r yields the same expected payoff as experimentation, i.e. $U_r(\underline{p}(\bar{p}_4)) = V^\sigma(\underline{p}(\bar{p}_4), \bar{p}_4)$. Graphically, \bar{p}_4 can be seen as the point where the left endpoint of the value segment crosses the stopping payoff U_r (see right panel of **Figure 6**). It is important to notice that at $\bar{p} = \bar{p}_4$, the value of $\underline{p}(\bar{p})$ lies strictly to the left of the Bayesian stopping boundary p_r^B . This follows from the fact that the DM's future selves follow a rule that is optimal with respect to \bar{p} but not with respect to $\underline{p}(\bar{p})$. From the perspective of $\underline{p}(\bar{p})$, the value of experimentation is thus strictly lower than the Bayesian value $\Phi(\underline{p}(\bar{p}))$. This means that the DM with worst-case belief $\underline{p}(\bar{p})$ is more inclined to stop than the Bayesian DM with the corresponding belief.

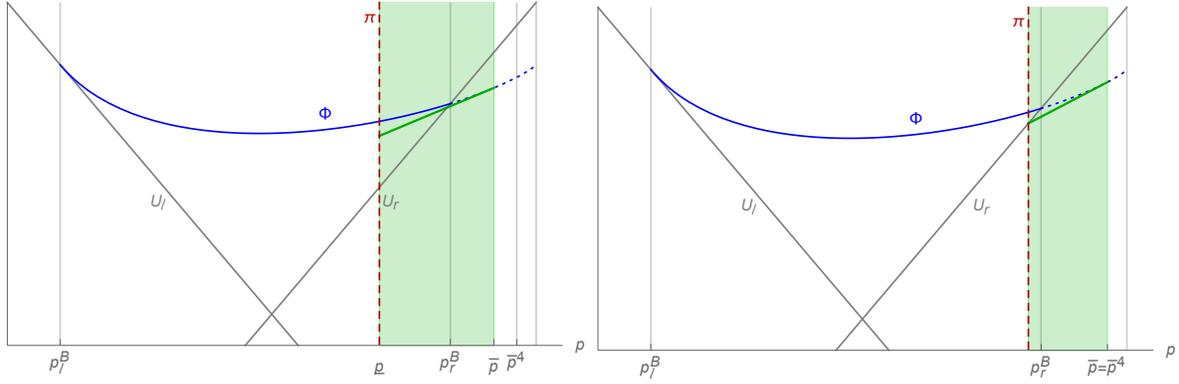


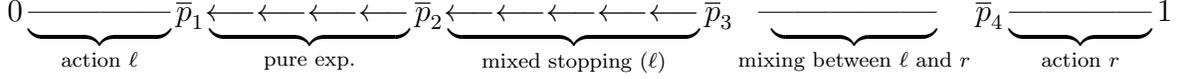
Figure 6: The value segment and worst-case belief for $\bar{p} \in (\bar{p}_3, \bar{p}_4)$ in the case of Δ small.

Next, consider the case in which Δ is sufficiently large so that Region 3 with mixed stopping is nonempty. The right boundary of the mixed stopping region, \bar{p}_3 , is given by the state \bar{p} at which the value segment separates from U_ℓ and experimentation becomes again optimal for all posteriors in $[\underline{p}(\bar{p}), \bar{p}]$ (see Figure 5). To the right of \bar{p}_3 , we then have a second region of pure experimentation, Region 4. The value segment becomes upward-sloping in this region, so $\underline{p}(\bar{p})$ becomes the worst-case belief. As with the case of small Δ , the right boundary of Region 4 is determined as the point where the left end of the value segment reaches the stopping payoff U_r . Notice, however, that the value segment in the experimentation region remains below the Bayesian value function. Compared to the case of small Δ , the possibility of randomization in future states reduces the value of experimentation further, which pushes the stopping boundary \bar{p}_4 towards the middle. Intuitively, the DM here is discouraged from experimentation, and is thus more inclined to stop, not only because it leads to excessive experimentation in the distant future, but also because it leads to inefficient stopping in the intermediate future.

We now complete the characterization by addressing several remaining cases and issues. The reader interested in the behavioral implications of the baseline characterization may skip to the next section.

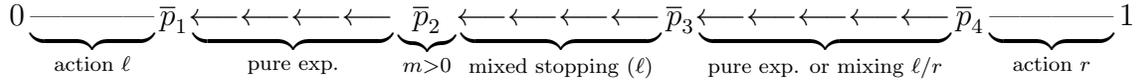
Intermediate experimentation costs. Consider next the case in which $c \in (\underline{c}, \bar{c})$. Theorem 1 shows that for each $c \in (\underline{c}, \bar{c})$ there exists a threshold $\Delta_c \in \mathbb{R}$ such that for $\Delta \leq \Delta_c$ the equilibrium continues to be described by case (a) of Theorem 1. If instead $\Delta > \Delta_c$, the fourth region features

randomization between actions ℓ and r rather than experimentation.



The equilibrium behavior for the first three regions is the same as in the case with low experimentation costs. In the third region, however, where the DM randomizes between experimentation and action ℓ , the value segment may drop as low as the point at which the two stopping payoff lines, U_ℓ and U_r , cross. The state at which this occurs constitutes our new boundary \bar{p}_3 . The DM's value at this point is equal to $\min_p U(p) = U(\hat{p})$, the minimal expected payoff that the DM can guarantee to herself by randomizing between actions r and ℓ with probabilities \hat{p} and $1 - \hat{p}$, respectively. At states to the right of this point, randomizing between immediate actions with \hat{p} strictly dominates experimentation with randomized stopping, giving rise to the fourth region $[\bar{p}_3, \bar{p}_4)$. Its right boundary is the state \bar{p} that satisfies $\underline{p}(\bar{p}) = \hat{p}$. Further to the right, where $\underline{p}(\bar{p}) > \hat{p}$, the DM optimally takes action r . The solution for intermediate costs $c \in (\underline{c}, \bar{c})$ and large ambiguity $\Delta > \Delta_c$ can thus be viewed as a combination of the previous two solutions.

Case 2. We now turn to Case 2, which is characterized by $\Phi'(p_*) < 0$. The formal characterization of the intrapersonal equilibrium for this case is relegated to [Appendix B.8](#). Compared with the case $\Phi'(p_*) = 0$, there is only one significant change. At the right boundary of the first experimentation region (\bar{p}_1, \bar{p}_2) , the stopping distribution has an atom: the DM takes action r with positive probability $m(\bar{p}_2)$ and experiments with the complementary probability.



Moving backwards in time through the experimentation region (\bar{p}_1, \bar{p}_2) , the value segment remains downward sloping as the state \bar{p} approaches $\bar{p}_2 = p_*$. The unique worst-case belief in the limit as $\bar{p} \nearrow p_*$ is thus \bar{p} . For states to the right of p_* , the Bayesian optimal strategy under \bar{p} is to take action r rather than experimenting. Taking action r in those states does, however, not constitute an intrapersonal equilibrium, as the worst-case belief under this strategy would be $\underline{p}(\bar{p})$ rather than \bar{p} . The DM must therefore experiment with positive probability in states to the right of p_* . To make experimentation optimal, the value segment must be weakly upward

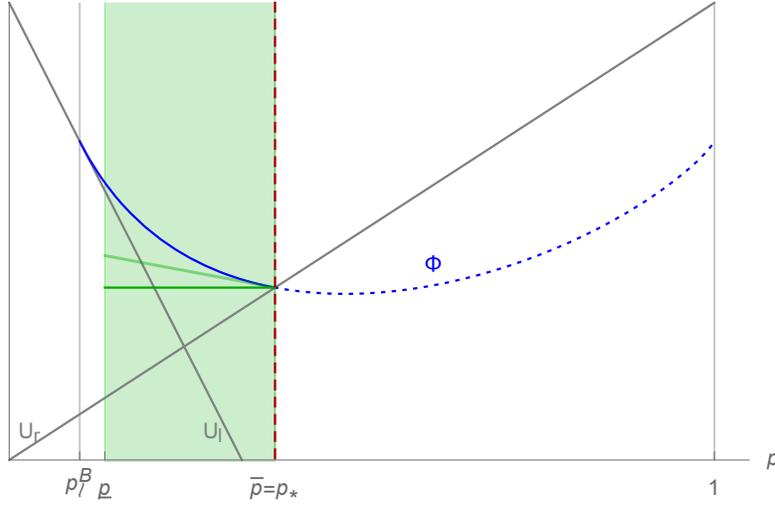


Figure 7: Case $\Phi'(p_*) < 0$. The value segment for $\bar{p} \nearrow p_*$ and $\bar{p} = p_*$.

sloping for states belonging to an interval $(p_*, p_* + \varepsilon)$. This property requires the DM to take action r with strictly positive probability at state $\bar{p} = p_*$. Graphically, the DM's randomization flips the left end of the value segment downwards, so that it is flat at $\bar{p} = p_*$, as illustrated in Figure 7.¹⁵

Uniqueness. In Appendix B.9, we prove that the intrapersonal equilibrium is unique when either $c < \underline{c}$ or $c \geq \bar{c}$. When costs are intermediate, however, there is a class of intrapersonal equilibria. These differ from the one described in Theorem 1 only in the position of their stopping boundary on the right. More specifically, recalling that \hat{p} is the belief at which the DM is indifferent between actions ℓ and r , fix any q such that $\underline{p}(q) < \hat{p} \leq q$. We then construct a

¹⁵The knife-edge case $\bar{p}_0 = p_*$ (in Case 2) is the only instance where the requirement that the DM's strategy constitutes an equilibrium regardless of the initial state actually restricts the set of equilibria. In this instance, requiring equilibrium conditions *only* for subsequent states $\bar{p} \leq \bar{p}_0$ would leave some flexibility in picking the randomization between experimentation and action r at state $\bar{p}_0 = p_*$. Indeed, as long as the value segment remains (weakly) downward sloping, all equilibrium conditions would be satisfied for the states to the left of $\bar{p}_0 = p_*$. It is the equilibrium conditions for states $\bar{p} > p_*$ that require the value segment to become flat at state p_* .

strategy profile $(\tilde{\sigma}, \tilde{\pi})$ indexed by $q \in [\hat{p}, \underline{p}^{-1}(\hat{p})]$ as follows:

$$\tilde{\sigma}(\bar{p}) = (\tilde{m}(\bar{p}), \tilde{\nu}(\bar{p}), \tilde{\rho}(\bar{p})) = \begin{cases} \sigma(\bar{p}) \\ (1, 0, \hat{\rho}) \\ (1, 0, 1) \end{cases} \quad \tilde{\pi}(\bar{p}) = \begin{cases} \pi(\bar{p}) & \text{if } \bar{p} \leq q \\ \hat{p} & \text{if } q < \bar{p} < \underline{p}^{-1}(\hat{p}) \\ \underline{p} & \text{if } \underline{p}^{-1}(\hat{p}) \leq \bar{p}, \end{cases}$$

where (σ, π) is the intrapersonal equilibrium defined in [Theorem 1](#). In words, the DM follows the same strategy as in [Theorem 1](#) for states $\bar{p} \leq q$ but stops immediately in states $\bar{p} > q$, followed by the mixed action $\hat{\rho}$ if $\hat{p} \in (\underline{p}(\bar{p}), \bar{p})$ and action r if $\hat{p} \leq \underline{p}(\bar{p})$. Since $c < \bar{c}$, this new strategy yields a strictly lower payoff for the DM in the states just to the right of q . In fact, the value function jumps down at the “switching” state $\bar{p} = q$ to the point where the two stopping payoffs cross, so the DM is worse off from this strategy relative to the one in [Theorem 1](#). Nevertheless, the strategy profile—there is a continuum of such indexed by $q \in [\hat{p}, \underline{p}^{-1}(\hat{p})]$ —satisfies all requirements of our solution concept, including the HJB conditions. The intuition is that, even though the HJB conditions capture control over vanishingly proximate future actions, these conditions are not strong enough. Given $c \geq \underline{c}$, there always exists some $\epsilon < \bar{p} - q$ such that deviating only until the state falls to $\bar{p} - \epsilon$ will not generate a strictly higher payoff.

Nevertheless, we view these equilibria as implausible. First, we see that $\tilde{\sigma}$, regardless of q , is Pareto worse than σ for all selves of the DM evaluated at the worst-case beliefs. Hence, $\tilde{\sigma}$ will be eliminated if the multiple selves of the DM can coordinate on the selection of equilibria, as is often invoked in the name of “consistent planning;” see [Siniscalchi \(2011\)](#).¹⁶ Second, we could strengthen the notion of vanishing commitment to eliminate the equilibria. Namely, instead of requiring the equilibrium strategy profile to be a limit of the ϵ -commitment solution *pointwise* for each \bar{p} , one could require the limit convergence to be *uniform* across all \bar{p} 's. The profile $(\tilde{\sigma}, \tilde{\pi})$ cannot survive this stronger test since for any $\epsilon > 0$, there will always exist a state $\bar{p} \in (q, q + \epsilon)$, such that the ϵ -commitment solution for the DM at that state will entail a deviation to σ to enjoy a discontinuous payoff jump.¹⁷

¹⁶More precisely, [Siniscalchi \(2011\)](#) requires the current self to break a tie in favor of her earlier self. See [Ebert and Strack \(2018\)](#) for a related approach.

¹⁷One could adopt this stronger requirement as a solution concept. We do not take this approach for two reasons. First, the uniform convergence requirement is not easy to check since it requires one to verify the equicontinuity property of the value function. Second, the requirement jeopardizes the existence in the more general Poisson model we discuss in [Appendix C](#). The unique intrapersonal equilibrium presented in [Theorem 4](#) involves a discontinuous value function when the cost is intermediate, which does not survive the stronger refinement.

Knightian Uncertainty. The extreme case in which $\underline{p} = 0$ and $\bar{p} = 1$, known as Knightian uncertainty, constitutes an important theoretical benchmark. Naturally, the equilibrium characterization for this case is obtained as the limit of the solution we characterized above for $(\underline{p}_0, \bar{p}_0) \rightarrow (0, 1)$.

Proposition 1. *Assume $\mathcal{P}_0 = [0, 1]$. If $c \geq \underline{c}$, the DM mixes between immediate actions ℓ and r with probability \hat{p} . Otherwise the DM mixes between experimentation and stopping followed by ℓ at a stationary rate*

$$\tilde{\nu} = \frac{\lambda}{\delta_\ell} (u_r^R - \tilde{u}), \text{ where } \tilde{u} = \frac{u_r^R + u_\ell^L}{2} - \sqrt{\left(\frac{u_r^R - u_\ell^L}{2}\right)^2 + \frac{c}{\lambda} \delta_\ell}.$$

This solution is the limit of the equilibrium strategy characterized in [Theorem 1](#) ([Theorem 3](#) for Case 2) for the case $0 < \underline{p}_0 < \bar{p}_0 < 1$ as $(\underline{p}_0, \bar{p}_0) \rightarrow (0, 1)$.

As $(\underline{p}_0, \bar{p}_0) \rightarrow (0, 1)$, Region 3 featuring randomized stopping takes over the entire region to the right of p_* . Starting from \bar{p}_0 close to one, the DM will then spend an arbitrarily long time in Region 3. In the extreme case, when $[\underline{p}_0, \bar{p}_0] = [0, 1]$, the DM's set of priors remains unchanged in the absence of breakthrough news. Hence, under any Markovian strategy, the DM either stops immediately or stops at a constant rate. The latter is optimal if and only if $c \leq \underline{c}$.

We can interpret the Knightian solution as a form of hedging by the DM to ensure her payoff equalizes between the states. Indeed, it is easy to verify that the stationary stopping rate $\tilde{\nu}$ satisfies the following equality:

$$u_\ell^L - \frac{c}{\tilde{\nu}} = \frac{\lambda}{\tilde{\nu} + \lambda} u_r^R + \frac{\tilde{\nu}}{\tilde{\nu} + \lambda} u_\ell^R - \frac{c}{\tilde{\nu} + \lambda}. \quad (17)$$

The LHS shows the DM's payoff from stopping at a constant rate $\tilde{\nu}$ when the state is L . In state L the DM never observes a breakthrough and thus always ends up taking the correct action. The downside is a relatively high expected cost of experimentation equal to $c/\tilde{\nu}$. The RHS shows the DM's payoff when the state is R . Now, the DM receives R -evidence before she stops experimenting, with positive probability $\lambda/(\tilde{\nu} + \lambda)$. With the complimentary probability, however, the DM takes the wrong action and receives the lower payoff u_ℓ^R . The payoff gross of experimentation costs is then lower in state R than in state L , which is compensated by the DM experimenting for shorter periods (in expectation) and therefore incurring a lower expected

cost of experimentation equal to $c/(\tilde{\nu} + \lambda)$.

The role of randomization. Whether randomization helps to hedge against ambiguity depends on the DM's internal view on how the uncertainty unfolds (see Saito (2015) and Ke and Zhang (2020)). We adopt the assumption that the DM views the game against nature as one with simultaneous moves. Hence, in the DM's mind, nature cannot condition her choice on the outcome of the DM's randomization. If the DM takes on an even more pessimistic view, believing that nature can adjust her choice *after* the DM's randomization outcome, any realization of a mixed strategy will be evaluated according to the worst-case belief for this realization, so the DM has no incentives to randomize.

Modifying the setting in this way clearly changes the characterization of our solution for the case when ambiguity is large. It does, however, not erase the DM's desire to stop preemptively in anticipation of a long experimentation phase when ambiguity is large. As we will show, stopping occurs again at intermediate sets of beliefs where the worst-case belief is about to switch. But instead of an extended period of randomized stopping for an interval of states, now stopping occurs at an isolated state \bar{p}^* . Anticipating this, earlier selves of the DM are happy to experiment until that point. Even though the game ends once the state reaches \bar{p}^* , if the DM would find herself at state $\bar{p} < \bar{p}^*$, she would again be experimenting. In this sense, the non-monotonic pattern of stopping persists even without randomization. To keep the proof short, we restrict attention to the symmetric case where $u_r^R = u_\ell^L =: \delta$ and $u_\ell^R = u_r^L = 0$.

Proposition 2. *Suppose the DM is restricted to pure strategies. Assume $\Phi'(p_*) = 0$, $c < \bar{c}$ and payoffs are symmetric. If Δ is sufficiently large, then there is a state $\bar{p}^* \in (0, 1)$ such that the DM takes action ℓ at \bar{p}^* and experiments on a left and right neighborhood of \bar{p}^* .*

To understand this result, let us meet the DM on the journey backward in time, at the exact moment when the worst-case belief is about to switch, which is captured by the right panel of Figure 4.¹⁸ Moving further back in time from that moment, the value segment associated with experimentation becomes upward-sloping, so the DM, as before, finds herself in a dilemma where further experimentation is met with the adversarial belief $\underline{p}(\bar{p})$ whereas stopping with ℓ is met with \bar{p} . With nature now wielding the second-mover advantage, the DM can no longer protect herself from these adversarial beliefs via randomization. Instead, the DM must choose

¹⁸Note the DM's equilibrium behavior in Regions 1 and 2 is unchanged.

the lesser of two bad choices. As long as $\bar{p} > p_*$ is relatively close to p_* , the value segment is sufficiently flat so that the value of experimentation at $\underline{p}(\bar{p})$ is higher than the stopping payoff U_ℓ at belief \bar{p} . Experimentation thus remains the better option for an interval of time before p_* is reached. However, going further back in time, and assuming that ambiguity is sufficiently large, there comes a state \bar{p}^* at which the value of experimentation at its worst belief $\underline{p}(\bar{p}^*)$ is as low as $U_\ell(\bar{p}^*)$. At that state, stopping becomes optimal for the DM. As noted above, given that the DM stops at \bar{p}^* , earlier selves are willing to experiment until \bar{p}^* is reached. The equilibrium strategy thus prescribes experimentation to the left and right of \bar{p}^* .

4 Behavioral Implications of Ambiguity

In this section, we perform comparative statics to further explore the behavioral implications of ambiguity.

Prolonged Learning and Indecision. We first investigate the DM's incentive to prolong her experimentation late in the game and how it varies with ambiguity. For the latter purpose, we denote by θ the *true* probability of R and consider the case where the DM's set of priors \mathcal{P} is centered around θ . Specifically, for each $\theta \in (0, 1)$ and each $\Delta \geq 0$, we define

$$\mathcal{P}_\Delta(\theta) := \left\{ p \in [0, 1] : \left| \ln \left(\frac{p}{1-p} \right) - \ln \left(\frac{\theta}{1-\theta} \right) \right| \leq \frac{\Delta}{2} \right\}$$

as the set of priors with midpoint θ in terms of the log-likelihood ratio and ambiguity level Δ . We then denote by $T_\Delta(\theta)$ the expected length of experimentation when the DM's set of priors is $\mathcal{P}_\Delta(\theta)$ and the true probability of state R is θ . We have the following result.

Proposition 3 (Length of experimentation). *Suppose $c < \bar{c}$ so that experimentation occurs under some θ . Let $0 \leq \Delta < \Delta'$ and define $\bar{\theta} := \sup\{\theta : T_\Delta(\theta) > 0\}$. There exists some $\hat{\theta} \in (0, \bar{\theta}]$ such that $T_{\Delta'}(\theta) \geq T_\Delta(\theta)$ if $\theta < \hat{\theta}$ and $T_{\Delta'}(\theta) \leq T_\Delta(\theta)$ if $\theta \in [\hat{\theta}, \bar{\theta})$. Further, $T'_{\Delta'}(\theta) \leq T'_\Delta(\theta)$ for all $\theta < \hat{\theta}$.*

In words, an increased ambiguity leads the DM to prolong her experimentation when the likelihood θ of state R is sufficiently small. For a benchmark, consider the Bayesian DM with accurate prior $p_0 = \theta$ and no ambiguity ($\Delta = 0$). The length of experimentation for the Bayesian

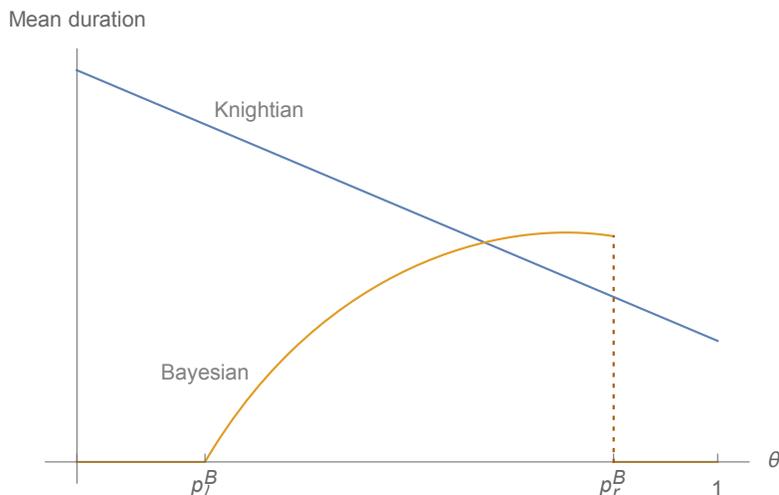


Figure 8: Mean sampling time for the Bayesian case and the case of Knightian uncertainty.

DM is concave and strictly increasing in θ for θ not too large,¹⁹ as illustrated in Figure 8. There are two competing forces. On the one hand, as θ falls, experimentation is less likely to produce conclusive R -evidence; this tends to lengthen the experimentation. On the other hand, the DM adjusts her learning strategy to shorten its duration as her (accurate) prior θ falls. For θ sufficiently small, the latter effect dominates the former effect, so the overall learning time decreases as R becomes less likely.

As stated in Proposition 3, ambiguity makes the DM’s strategy less responsive to the change in θ . The reduced responsiveness to θ corresponds to the behavioral trait of “indecisiveness.” In the extreme case of $\Delta \rightarrow \infty$ (i.e., Knightian ambiguity), the learning strategy becomes completely unresponsive to the true likelihood θ , so that, as θ falls, her learning time actually increases (see Figure 8). The reason is that even after an arbitrarily long duration of unsuccessful experimentation, the DM’s extreme ambiguity prevents her from dismissing R as sufficiently unlikely and stopping experimentation.

Premature Decisiveness. While the Bayesian optimal stopping time is deterministic conditional on not receiving R -evidence, a sufficiently large ambiguity leads to randomized stopping at earlier points in time. Therefore, the ambiguity-averse DM exhibits early “decisiveness” not shown by her Bayesian counterpart. At the same time, an increased ambiguity also leads to

¹⁹See Proposition 1 of Che and Mierendorff (2019).

prolonged experimentation, as seen in [Proposition 3](#), which increases the likelihood of delayed stopping. The combined effect is the increased dispersion of the stopping times. To formalize this result, consider two prior sets \mathcal{P} and \mathcal{Q} with $\mathcal{P} \subset \mathcal{Q}$. Let $F_{\mathcal{P}} : [0, +\infty) \rightarrow [0, 1]$ denote the cumulative distribution function over stopping times when the initial set of priors is \mathcal{P} . Hence, given prior set \mathcal{P} , $F_{\mathcal{P}}(t)$ is the probability that the DM stops before time t .

Proposition 4 (Dispersion of stopping times.). *Assume $c \leq \underline{c}$ and $\Phi'(p_*) \geq 0$. Let \mathcal{P} and \mathcal{Q} be such that $\mathcal{P} \subseteq \mathcal{Q}$. There exists a time \hat{t} such that $F_{\mathcal{Q}}(t) \geq F_{\mathcal{P}}(t)$ for all $t < \hat{t}$ and $F_{\mathcal{Q}}(t) \leq F_{\mathcal{P}}(t)$ for all $t > \hat{t}$.*

[Proposition 4](#) shows that the stopping time distribution under the two nested sets \mathcal{P} and \mathcal{Q} exhibits a single crossing property. This property indicates that the distribution over stopping times under the larger set \mathcal{Q} is more dispersed than the one under \mathcal{P} . Indeed, whenever the mean stopping time under \mathcal{P} is weakly greater than under \mathcal{Q} , then $F_{\mathcal{P}}$ second-order stochastically dominates $F_{\mathcal{Q}}$ (when the means are equal, $F_{\mathcal{Q}}$ constitutes a mean-preserving spread of $F_{\mathcal{P}}$). The single crossing property implies that the range of potential stopping points spreads as ambiguity becomes larger: an expansion of the initial set of priors leads to an earlier first potential stopping point and a later last potential stopping point. The property is illustrated in [Figure 9](#), where we plot the cumulative probability of stopping conditional on state L for the Bayesian DM (orange graph) and the ambiguity-averse DM (blue graph).

Can the two traits coexist? According to our theory, ambiguity-aversion leaves potentially testable behavioral markers—prolonged indecision and impulsive decisiveness. While seemingly contradictory, there is a coherent logic that ties these two traits together. For an illustration, consider again the theorist seeking to prove a new theorem with unknown validity. According to our theory, ambiguity—with regard to the unknown validity of the theorem—makes it difficult for the theorist to dismiss the validity of the theorem even after many unsuccessful attempts to prove it. Hence, ambiguity might lead the DM to a costly process of continued trial that could be seen by the outsider as a “wild goose chase.” Recognizing that such an obsessive pursuit is in the offing—perhaps also reminded by how such endeavors ended in the past—the DM may simply refrain from getting into the situation by stopping early; the theorist may prematurely give up after making a few unsuccessful attempts at proving the theorem.²⁰

²⁰[Barkley-Levenson and Fox \(2016\)](#) find indecision and impulsivity to be positively correlated among the

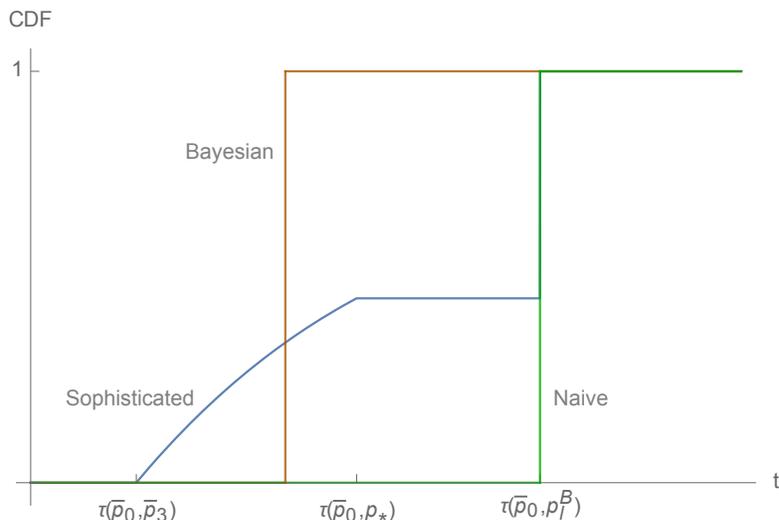


Figure 9: Cumulative distribution of stopping times (conditional on $\omega = L$) for the Bayesian DM and the ambiguity-averse DM when sophisticated (in blue) and naive (in green). In the graph, $\tau(p, q)$ denotes the time it takes for belief p to be updated to q in the absence of R -evidence.

The role of sophistication. While the feature of prolonged learning as a response to ambiguity does not hinge on the DM’s sophistication, premature decisiveness is driven by the DM’s ability to correctly forecast her future behavior and respond accordingly. Indeed, a fully naive DM—one who acts according to the best plan given by the current worst-case scenario but fails to carry out the plans of earlier selves—only exhibits the first behavioral trait and will thus always experiment longer than the corresponding Bayesian DM. The sophisticated DM, on the other hand, will engage in preemptive stopping when ambiguity is sufficiently large, so she will tend to stop earlier than her naive counterpart. Note, however, that both DMs share the same final stopping points, so with positive probability both stop at the same time. We illustrate this feature in **Figure 9**: the CDF for the naive DM first-order stochastically dominates that of the sophisticated DM.

It is also interesting to contrast the sophisticated and naive DM with a DM who has full commitment power. As we saw in **Section 2**, the commitment solution is given by the Bayesian

respondents of their behavioral survey, and reach an interpretation uncannily close to our theory: “Our findings suggest ... indecisiveness and impulsivity can be viewed as two sides of the same coin. Both indecisiveness and impulsivity are maladaptive behavioral responses to difficulty engaging with a decision. We surmise that these responses arise from a common desire to avoid negative affect that some individuals experience when making choices. ... they may behave in an impulsive manner, hastily choosing in order to avoid unpleasant deliberation over opportunity costs.”

optimal stopping rule for the prior in $[p_0, \bar{p}_0]$ that minimizes the Bayesian value function Φ^* . Assuming $\bar{p}_0 > p_*$, the graph illustrating the commitment solution is thus qualitatively similar to that of the Bayesian DM in [Figure 9](#). Specifically, a DM with full commitment will stop before the naive DM does. The comparison with a sophisticated DM is ambiguous, except that the latter's stopping time is more dispersed in the sense established in [Proposition 2](#).

5 Incremental Learning

So far, we have focused on Poisson learning, where news is one-sided and conclusive. Another canonical case is the incremental learning model in which the signal follows a diffusion process. In this section, we modify the baseline model and consider this alternative information technology. For simplicity, we restrict attention to the case where payoffs are symmetric across states: $u_r^R = u_\ell^L$, $u_r^L = u_\ell^R$.

Suppose the DM receives a signal $X \in \mathbb{R}$ about the state $\omega \in \{L, R\}$ at a flow cost of c per unit time, which follows the process:

$$dX_t = \mu_\omega dt + \sigma dB_t,$$

where $\mu_R > \mu_L$. As usual, it is more convenient to convert X into a belief, particularly into a log-likelihood ratio Z , where

$$dZ_t = \frac{\psi}{\sigma} dX_t - \frac{\psi}{\sigma} \left(\frac{\mu_R + \mu_L}{2} \right) dt = \pm \frac{\psi^2}{2} dt + \psi dB_t,$$

and $\psi := (\mu_R - \mu_L)/\sigma$ is the signal-to-noise ratio (see [Daley and Green \(2012, 2020\)](#) for example).

As before, the Bayesian optimal behavior is characterized by two stopping boundaries (in terms of log-likelihood ratio), z_ℓ and z_r , such that the DM stops and chooses ℓ if $z \leq z_\ell$, stops and chooses r if $z \geq z_r$, and experiments if $z \in (z_\ell, z_r)$. Let $\phi(z|z_\ell, z_r)$ be the value of this strategy when the DM's prior is z , where $z = 0$ corresponds to the belief of $p = 1/2$. The Bayesian optimal stopping boundaries are denoted by (z_ℓ^B, z_r^B) , where $z_\ell^B = -z_r^B < 0$ due to the symmetry of payoffs. As before, the stopping boundaries are graphically depicted by tangent points of the Bayesian value function at the stopping payoffs (which are nonlinear due to the

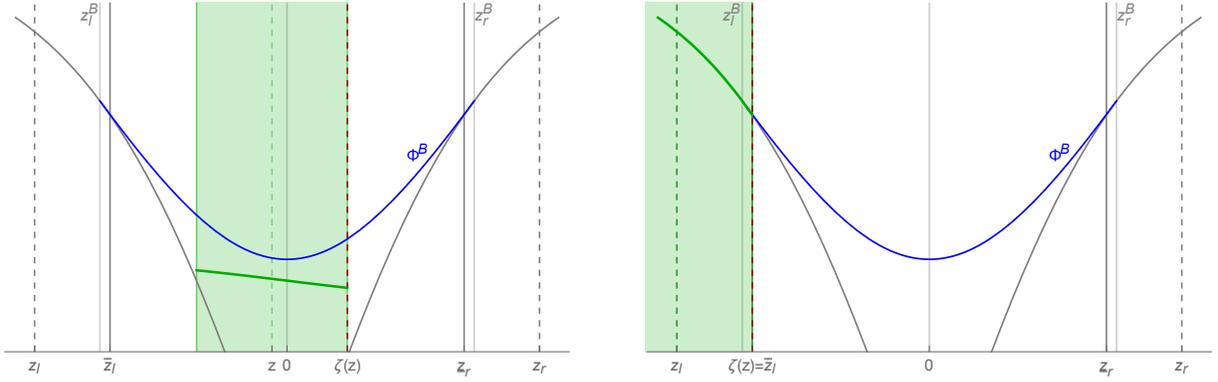


Figure 10: Agent’s value when experimenting (left) and when stopping (right), with worst-case belief $\zeta(z)$ and boundaries $\bar{z}_\ell := z_\ell + \Delta/2$ and $\bar{z}_r := z_r - \Delta/2$. The Bayesian value function is $\Phi^B(z) := \phi(z|z_\ell^B, z_r^B)$.

change of variable); see [Figure 10](#).²¹

Consider now an ambiguity-averse DM whose prior set is represented by an interval $\mathcal{Y}(z) := [z - \Delta/2, z + \Delta/2]$. In contrast to before, let us index the “state” by the midpoint z and the deviation from the midpoint in terms of the log-likelihood ratio by $\gamma \in [-\Delta/2, \Delta/2]$. Upon suitable modification, the intra-personal equilibrium is characterized much like in the baseline model.

Small Δ . If ambiguity is sufficiently small, there are again two stopping boundaries $z_\ell < 0 < z_r$ (where $z_r = -z_\ell$) such that the DM experiments if and only if the state z belongs to the interval (z_ℓ, z_r) . Suppose the DM follows this strategy. Then, at state z , the DM with belief $(z + \gamma) \in \mathcal{Y}(z)$ perceives an expected payoff of

$$V(z + \gamma, z) = \phi(z + \gamma | z_\ell + \gamma, z_r + \gamma).$$

Close to the stopping boundary, the DM’s worst-case belief is given by the “inner-most” belief of the set, just as in the Poisson model. The two stopping boundaries z_ℓ and z_r are thus determined by the smooth-pasting conditions with respect to these beliefs:

$$\begin{aligned} \phi' \left(z_\ell + \frac{\Delta}{2} \mid z_\ell + \frac{\Delta}{2}, z_r + \frac{\Delta}{2} \right) &= U'_\ell \left(z_\ell + \frac{\Delta}{2} \right) \\ \phi' \left(z_r - \frac{\Delta}{2} \mid z_\ell - \frac{\Delta}{2}, z_r - \frac{\Delta}{2} \right) &= U'_r \left(z_r - \frac{\Delta}{2} \right), \end{aligned}$$

²¹This value function and the Bayesian optimal stopping boundaries, z_ℓ^B and z_r^B , can be derived using standard methods.

where, with slight abuse of notation, U_a , $a = \ell, r$ denotes the DM's stopping payoff as a function of the log-likelihood ratio. Stopping for action ℓ occurs at state z_ℓ where the DM is indifferent between stopping and experimenting *at her right-most belief* $z_\ell + \Delta/2$, while stopping for action r occurs at state z_r where the DM is indifferent between stopping and experimenting *at her left-most belief* $z_r - \Delta/2$.

A notable difference relative to the Poisson model is that the stopping beliefs, $z_\ell + \Delta/2$ and $z_r - \Delta/2$, no longer coincide with the corresponding Bayesian stopping boundaries, z_ℓ^B and z_r^B . Instead, we have $z_\ell^B < z_\ell + \Delta/2 < 0 < z_r - \Delta/2 < z_r^B$, so the DM stops experimenting before stopping becomes Bayesian optimal for all beliefs in the set.²² To see this, consider the case where the DM's right-most belief is close to z_ℓ^B . If the DM continues experimentation, there is a small chance that the state drifts all the way to the other side of the right stopping boundary. The DM understands that if this happens, she will stop only when the left-most belief reaches the stopping boundary. From the viewpoint of the right-most belief, there is hence excessive experimentation close to the right stopping boundary, which reduces the value of experimenting. The same applies to the other side. The argument further implies that in the experimentation region, there is no belief in $[z - \Delta/2, z + \Delta/2]$ for which the equilibrium strategy is Bayesian optimal. This means that the value segment for the ambiguity-averse DM is strictly below, and thus never touches, the Bayesian value function, as illustrated in [Figure 10](#).

Large Δ . The above solution constitutes an equilibrium only when ambiguity is sufficiently small. Graphically, the value segment has to be sufficiently short so that its lower end does not fall below the stopping payoff. Otherwise stopping becomes optimal under the worst-case belief, as we illustrate in [Figure 11](#) (left panel). In such situations, the DM again randomizes between experimentation and stopping. Mixed stopping works here as follows: starting from $z = 0$, if z moves slightly to the left, the DM mixes between experimentation and stopping followed by action r ; if instead z moves slightly to the right, the DM mixes between experimentation and stopping followed by ℓ . This pattern is analogous to how randomized stopping occurs in the baseline model. Intuitively, mixing with action r when $z < 0$ (resp. ℓ when $z > 0$) hedges the DM against the possibility that the state is in fact R (resp. L) and, hence, that experimentation will take a very long time.

The DM's value and equilibrium stopping rate in the mixed stopping region of the drift-

²²Epstein and Ji (2022) obtain a qualitatively similar result for the case of rectangular sets of priors.

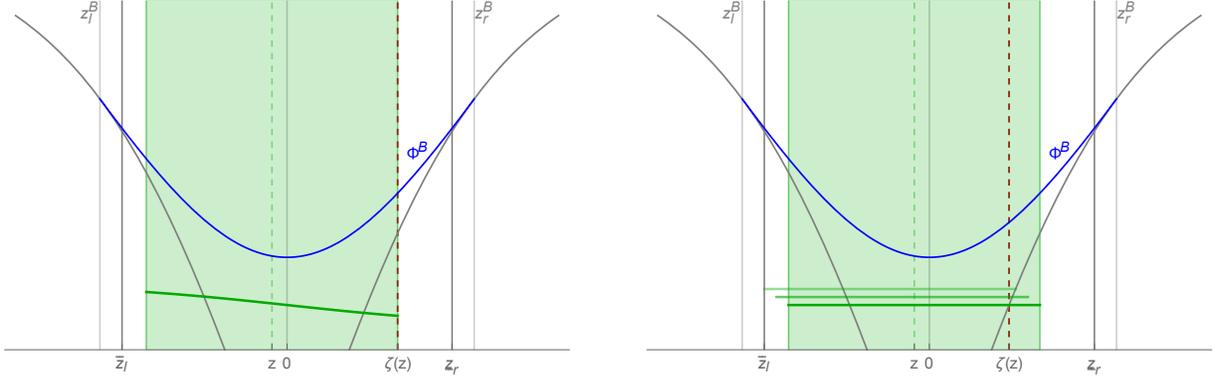


Figure 11: The case of large Δ : on the left, profitable deviation if the DM would experiment; on the right, intrapersonal equilibrium with mixed stopping (followed by r).

diffusion model can be derived via the same approach we developed for the model with Poisson learning. The saddle-point HJB condition is now

$$c = \max_{\nu} \min_{y \in \mathcal{Y}(z)} \nu(U(y) - V(y, z)) + \frac{\psi^2}{2}(2p(y) - 1)V_y(y, z) + \frac{\psi^2}{2}(2p(y) - 1)V_z(y, z) + \frac{\psi^2}{2}(V_{yy}(y, z) + V_{yz}(y, z) + V_{zz}(y, z)),$$

where $p(y) := e^y/(1 + e^y)$ is the probability of R corresponding to the log-likelihood ratio y . Letting $\zeta = \zeta(z)$ denote nature's choice in state z , we can derive $\zeta(z)$ and the DM's stopping rate $\nu(z)$ by following the same steps as in [Section 3.2](#). Let us consider the case $z < 0$ where stopping is followed by r (the case $z > 0$ is analogous). Randomization between experimentation and r requires that the coefficient of ν must vanish:

$$U_r(\zeta(z)) = V(\zeta(z), z). \tag{18}$$

To satisfy this condition, nature must choose an interior $\zeta(z)$, which requires that the value segment is flat. This means that V_y, V_{yz}, V_{yy} all vanish and we can write $V(y, z) = \hat{V}(z)$ for all y . We then use [\(18\)](#) to pin down nature's choice:

$$\zeta(z) = p^{-1} \left(\frac{\hat{V}(z) - u_r^L}{u_r^R - u_r^L} \right). \tag{19}$$

For nature to choose an interior likelihood ratio, the coefficient of y in the saddle-point condition must vanish, too:

$$\nu(z)(u_r^R - u_r^L) + \psi^2 V_z(\zeta(z), z) = 0. \quad (20)$$

Using again (18), this requirement ties down the DM's stopping rate $\nu(z)$. Finally, we need to determine the value function $\hat{V}(z)$. To this end, we substitute for $\zeta(z)$ and $\nu(z)$ to simplify the saddle-point condition and obtain the analog of (ODE).

$$c = \psi^2 \left(\frac{\hat{V}(z) - u_r^L}{u_r^R - u_r^L} - \frac{1}{2} \right) \hat{V}'(z) + \frac{\psi^2}{2} \hat{V}''(z).$$

The main complication in the drift-diffusion model is that the boundary condition for this ODE is no longer pinned down by the minimum of the Bayesian value function, as it was before, but has to be found as a fixed point.

Apart from this technical difference, the solution method and qualitative features of the solution extend to the drift-diffusion model. In particular: 1) the DM has incentives to engage in excessive experimentation close to the stopping boundaries; 2) if ambiguity is large and the flow cost c is small, the DM will exert self-control and preempt herself from experimenting excessively by randomly stopping at intermediate sets of beliefs.

6 Concluding Remarks

This paper studies the canonical Wald problem under the assumption that the DM faces uncertainty about the underlying probabilities of the different states. We consider a DM who seeks to maximize her payoff guarantee and updates beliefs prior by prior. Since the DM's worst-case belief may change over time, the DM's preference (at the worst-case belief) may violate dynamic consistency. Rather than ruling out such violations, we seek to understand their behavioral implications on information acquisition. The main technical challenge is then to solve the DM's optimization problem in the presence of dynamic preference reversals. We develop a new method, based on a version of HJB conditions, and provide a micro-foundation for the approach. The method proves highly tractable and is applicable to dynamic optimization problems beyond information acquisition, such as real options theory, dynamic portfolio choice, or search models.

We show that ambiguity gives rise to two behavioral traits—prolonged indecision and pre-

mature decisiveness. The inability by our DM to stop searching for R -evidence after many unsuccessful attempts is reminiscent of the *sunk cost fallacy*, i.e., the tendency not to give up on projects or activities in which one has invested. In our case, the DM struggles to dismiss the possibility of the state being R even after many failed bids to find evidence for it. The DM in that situation is guided by the right-most belief—namely, the belief that not only rationalizes the incurred experimentation costs so far but also calls for further continuation of learning. Ambiguity aversion could thus be viewed as an alternative mechanism behind the sunk cost fallacy, complementing other existing behavioral explanations.

Our model can be extended to allow for a choice in the type of information the DM acquires. In [Appendix C](#) we consider the general Poisson model introduced by [Che and Mierendorff \(2019\)](#), where the DM has access to two news sources, one emitting R -evidence (as in the baseline model), the other emitting L -evidence. Conditional on acquiring information, the DM must then allocate her attention between the two news sources. 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 show that excessive learning and randomized stopping continue to arise in this setting with sufficiently large ambiguity and sufficiently small information costs. In addition, the DM may “split” her attention between the two news sources in order to hedge, even when a Bayesian DM would not do so under any prior. Whenever the DM employs the split-attention strategy, she continues learning until she fully resolves the uncertainty. This behavior thus provides another pattern of excessive learning, when it is seen from the perspective of a Bayesian observer with comparable beliefs.

An open question is how our solution changes under different updating rules. Prior-by-prior updating entails that the DM updates on the state but not on the probability model. On the other extreme of the spectrum is maximum-likelihood updating ([Gilboa and Schmeidler, 1993](#)), according to which the DM only retains those priors which explain the data best. In our setting, this would not be very interesting, as the DM, in the absence of a breakthrough, would discard all priors but the left-most after the first instant of time. There are, however, less extreme versions that combine features of both updating rules (e.g. [Cheng \(2021\)](#)). It will be interesting to explore the implications of such rules on the incentives to acquire information in dynamic settings.

We focus on the case where the DM is fully sophisticated and thus correctly forecasts the behavior of her future selves. An open question is how the results extend to the case where the

DM is partially naive. For instance, the DM may be paying attention only at certain points in time. This could be modeled by having a Poisson clock run in parallel, which then determines the times at which the DM contemplates the actions of her future selves and optimizes accordingly (otherwise she follows the naively optimal plan). A higher Poisson rate would thus capture a more sophisticated DM. Given such a model, one could ask how the DM's welfare (say with respect to initial preferences) would depend on the level of sophistication. We leave these interesting questions for future research.

References

- Arrow, K. J., Blackwell, D., and Girshick, M. A. (1949). Bayes and minimax solutions of sequential decision problems. *Econometrica*, pages 213–244. [2](#)
- Auster, S. and Kellner, C. (2022). Robust bidding and revenue in descending price auctions. *Journal of Economic Theory*, page 105072. [5](#)
- Barberis, N. (2012). A model of casino gambling. *Management Science*, 58(1):35–51. [5](#)
- Bardi, M. and Capuzzo-Dolcetta, I. (1997). *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhäuser. [13](#)
- Barkley-Levenson, E. E. and Fox, C. R. (2016). The surprising relationship between indecisiveness and impulsivity. *Personality and Individual Differences*, pages 1–6. [29](#)
- Beauchêne, D., Li, J., and Li, M. (2019). Ambiguous persuasion. *Journal of Economic Theory*, 179:312–365. [5](#)
- Bose, S. and Daripa, A. (2009). A dynamic mechanism and surplus extraction under ambiguity. *Journal of Economic theory*, 144(5):2084–2114. [5](#)
- Bose, S. and Renou, L. (2014). Mechanism design with ambiguous communication devices. *Econometrica*, 82(5):1853–1872. [5](#)
- Che, Y.-K. and Mierendorff, K. (2019). Optimal dynamic allocation of attention. *American Economic Review*, 109:2993–3029. [2](#), [4](#), [28](#), [36](#), [82](#), [84](#)

- Cheng, X. (2021). Relative maximum likelihood updating of ambiguous beliefs. *Journal of Mathematical Economics*, page 102587. 36
- Cheng, X. and Riedel, F. (2013). Optimal stopping under ambiguity in continuous time. *Mathematics and Financial Economics*, 7(1):29–68. 4
- Christensen, S. and Lindensjö, K. (2018). On finding equilibrium stopping times for time-inconsistent markovian problems. *SIAM Journal on Control and Optimization*, 56(6):4228–4255. 5
- Christensen, S. and Lindensjö, K. (2020). On time-inconsistent stopping problems and mixed strategy stopping times. *Stochastic Processes and their Applications*, 130(5):2886–2917. 5
- Daley, B. and Green, B. (2012). Waiting for news in the market for lemons. *Econometrica*, 80(4):1433–1504. 31
- Daley, B. and Green, B. (2020). Bargaining and news. *American Economic Review*, 110(2):428–74. 31
- Ebert, S. and Strack, P. (2015). Until the bitter end: On prospect theory in a dynamic context. *American Economic Review*, 105(4):1618–33. 5
- Ebert, S. and Strack, P. (2018). Never, ever getting started: On prospect theory without commitment. *Available at SSRN 2765550*. 5, 24
- Epstein, L. G. and Ji, S. (2022). Optimal learning under robustness and time-consistency. *Operations Research*, 70(3):1317–1329. 4, 33
- Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31. 2
- Epstein, L. G. and Schneider, M. (2007). Learning under ambiguity. *The Review of Economic Studies*, 74(4):1275–1303. 4
- Fudenberg, D., Strack, P., and Strzalecki, T. (2018). Speed, accuracy, and the optimal timing of choices. *American Economic Review*, 108(12):3651–84. 2, 4

- Ghosh, G. and Liu, H. (2021). Sequential auctions with ambiguity. *Journal of Economic Theory*, 197:105324. 5
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153. 2, 8
- Gilboa, I. and Schmeidler, D. (1993). Updating ambiguous beliefs. *Journal of Economic Theory*, 59(1):33–49. 36
- Gul, F. and Pesendorfer, W. (2021). Evaluating ambiguous random variables from choquet to maxmin expected utility. *Journal of Economic Theory*, 192:105129. 8
- Henderson, V., Hobson, D., and Tse, A. (2017). Randomized strategies and prospect theory in a dynamic context. *Journal of Economic Theory*, 168:287–300. 5
- Ke, S. and Zhang, Q. (2020). Randomization and ambiguity aversion. *Econometrica*, 88(3):1159–1195. 26
- Kellner, C. and Le Quement, M. T. (2018). Endogenous ambiguity in cheap talk. *Journal of Economic Theory*, 173:1–17. 5
- Laraki, R. and Solan, E. (2005). The value of zero-sum stopping games in continuous time. *SIAM Journal on Control and Optimization*, 43(5):1913–1922. 13
- Miao, J. and Wang, N. (2011). Risk, uncertainty, and option exercise. *Journal of Economic Dynamics and Control*, 35(4):442–461. 4
- Moscarini, G. and Smith, L. (2001). The optimal level of experimentation. *Econometrica*, 69(6):1629–1644. 2, 4
- Nikandrova, A. and Pans, R. (2018). Dynamic project selection. *Theoretical Economics*, 13(1):115–143. 4
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press. 9
- Peskir, G. and Shiryaev, A. (2006). *Optimal stopping and free-boundary problems*. Springer. 4
- Riedel, F. (2009). Optimal stopping with multiple priors. *Econometrica*, 77(3):857–908. 4

- Saito, K. (2015). Preferences for flexibility and randomization under uncertainty. *American Economic Review*, 105(3):1246–71. 26
- Shishkin, D. and Ortoleva, P. (2023). Ambiguous information and dilation: An experiment. *Journal of Economic Theory*, 208:105610. 8
- Siniscalchi, M. (2011). Dynamic choice under ambiguity. *Theoretical Economics*, 6(3):379–421. 24
- Wald, A. (1947). Foundations of a general theory of sequential decision functions. *Econometrica*, pages 279–313. 2, 4
- Xu, Z. Q. and Zhou, X. Y. (2013). Optimal stopping under probability distortion. *The Annals of Applied Probability*, 23(1):251–282. 5
- Zhong, W. (2019). Optimal dynamic information acquisition. *Available at SSRN 3024275*. 4

A Mathematical Appendix

A.1 Markov strategies and their values

For any fixed prior $\mathcal{P}_0 = [\underline{p}_0, \bar{p}_0]$, we first express the DM's strategy as specifying her action after each time t of experimenting without receiving any news. (Recall that the latter are the only nontrivial subgames we focus on, since the DM chooses r immediately following the arrival of R -evidence.)

Specifically, a *time-indexed strategy* is a function $\tilde{\sigma} : [0, \infty) \rightarrow [0, \infty) \times [0, 1]^2$, $\tilde{\sigma}_t = (\tilde{\nu}_t, \tilde{m}_t, \tilde{\rho}_t)$, where $\tilde{\nu}_t$ is the stopping rate and \tilde{m}_t the instantaneous stopping probability used at time t if no R -signal has been received up until time t . $\tilde{\rho}_t$ is the probability of action r conditional on stopping at time t without receiving an R -signal.²³ To ensure that a strategy admits a well-defined stopping time, we require that (a) any connected set of t 's in which $\tilde{m}_t = 1$ contains its infimum, and (b) the set $\{t \geq 0 | \tilde{m}_t \in (0, 1)\}$ is countable. Time-indexed strategies satisfying these restrictions are called *admissible*, and $\tilde{\Sigma}$ denotes the set of all admissible time-indexed strategies. Fix any $\tilde{\sigma} \in \tilde{\Sigma}$. Then, starting at any time s , the probability that the strategy $\tilde{\sigma}$ prescribes the DM to stop by time $t > s$ is written as:

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

Using this function, we can write the expected payoff, or the *value*, of strategy $\tilde{\sigma}$ given belief p at time t as follows:

$$\begin{aligned} & \tilde{V}^{\tilde{\sigma}}(p, t) \\ &= p \left(\int_t^\infty (1 - F_t(s)) e^{-\lambda s} [\tilde{\nu}_s U^R(\tilde{\rho}_s) + \lambda u_r^R - c] ds + \sum_{s \in [t, \infty)} (1 - F_t(s)) \tilde{m}_s e^{-\lambda s} U^R(\tilde{\rho}_s) \right) \\ & \quad + (1 - p) \left(\int_t^\infty (1 - F_t(s)) e^{-\lambda s} [\tilde{\nu}_s U^L(\tilde{\rho}_s) - c] ds + \sum_{s \in [t, \infty)} (1 - F_t(s)) \tilde{m}_s U^L(\tilde{\rho}_s) \right), \quad (21) \end{aligned}$$

where $U^\omega(\rho) = \rho u_r^\omega + (1 - \rho) u_\ell^\omega$ denotes the utility from the mixed action ρ if the state is ω .

²³Appendix B.3 shows that the current formulation of a strategy is induced by a well-behaved stopping time distribution that the DM may choose.

The law of motion for the belief in the absence of stopping is exogenous and deterministic, and the belief always drifts downwards. Hence, there is a one-to-one mapping between time t and the belief set $\mathcal{P}_t := [\underline{p}_t, \bar{p}_t]$ for any initial prior belief set $\mathcal{P}_0 := [\underline{p}_0, \bar{p}_0]$ in the absence of stopping, *independently of the DM's strategy*. That is, both \underline{p}_t and \bar{p}_t drift down according to the law of motion (1) given the initial values, \underline{p}_0 and \bar{p}_0 . As in the text, we index the set \mathcal{P}_t by the highest possible belief \bar{p}_t , and call it the *state*. Noting that given \bar{p} , \underline{p} is uniquely pinned down, we denote $\underline{p}_t := \underline{p}(\bar{p}_t)$. Any $\bar{p} < \bar{p}_0$ can then be uniquely mapped to the time $t \geq 0$ it takes for the state to drift down to \bar{p} :

$$t = \tau(\bar{p}) := \frac{1}{\lambda} \ln \left(\frac{\bar{p}}{1 - \bar{p}} \frac{1 - \bar{p}_0}{\bar{p}_0} \right). \quad (22)$$

Consequently, any time-indexed strategy $\tilde{\sigma} = (\tilde{\sigma}_t)$ induces a strategy $\sigma = (\sigma(\bar{p})) := (\tilde{\sigma}_{\tau(\bar{p})})$, and conversely, any strategy induces a time-indexed strategy via $\tilde{\sigma}_t := \sigma(\tau^{-1}(t))$. Importantly, note that the Markov structure of the strategy—the dependence of the strategy only on the state— involves no restriction on the strategy space.²⁴ Clearly, the admissibility on $\tilde{\sigma}$ corresponds to the admissibility imposed on σ . The expected payoff from any $\sigma \in \Sigma$ can then be computed from (21) via a change of variable using (22):

$$V^\sigma(p, \bar{p}) := \tilde{V}^{\tilde{\sigma}_{\tau(\cdot)}}(p, \tau(\bar{p})). \quad (23)$$

A.2 Microfoundation of the saddle-point HJB equations

We now justify our equilibrium concept or the use of saddle-point HJB equations. The idea is to imagine that at each point in time, the DM can control, or commit to, her actions for a brief period of time, say $\epsilon > 0$. She thus acts dynamically consistently for a brief period of time, while taking as given the strategies of her future selves beyond that time interval. We then require the strategy to be a limit of the ϵ -dynamically consistent strategies as $\epsilon \rightarrow 0$.

To make this formal, fix a candidate strategy σ^* . Suppose that, at any state \bar{p} and any time t , the DM can commit to any strategy σ (possibly different from σ^*) during a time interval $[t, t + \epsilon)$, for some $\epsilon > 0$, and the strategy switches to σ^* thereafter.

²⁴Recall that this one-to-one mapping is a result of our environment. Beyond our setting, it may be restrictive to focus on Markov strategies. That is, if the mapping $t \mapsto \bar{p}_t$ is not injective, then there can be strategies that depend on calendar time that cannot be encoded as a Markov strategy.

Let $(\sigma_\epsilon(\cdot), \pi_\epsilon(\cdot))$ be the saddle point strategy for such an ϵ horizon problem (followed by σ^*), or an ϵ -commitment solution. We say that a pseudo equilibrium (σ^*, π^*) is *dynamically credible* if for each \bar{p} there is a sequence of ϵ -commitment solutions $(\sigma_\epsilon(\bar{p}), \pi_\epsilon(\bar{p}))$ that converges to (σ^*, π^*) as $\epsilon \rightarrow 0$. We further say that a pseudo equilibrium (σ^*, π^*) is *well-behaved* if, for each $\bar{p} \in (0, 1)$, $V^{\sigma^*}(\pi^*(\bar{p}), \bar{p})$, $V_p^{\sigma^*}(\pi^*(\bar{p}), \bar{p})$ and $V_p^{\sigma^*}(\pi^*(\bar{p}), \bar{p}_-)$ are well-defined and continuous at $(\pi^*(\bar{p}), \bar{p})$. We now provide the justification based on dynamic consistency.

Write $W_\epsilon^\sigma(p, \bar{p})$ for the DM's value when the DM follows strategy σ for ϵ -period, followed by the candidate strategy σ^* . Then,

$$\begin{aligned} & W_\epsilon^\sigma(p, \bar{p}) \\ &= m(\bar{p})U_{\rho(\bar{p})}(p) + (1 - m(\bar{p})) \left[-p \int_{\bar{p}^\epsilon}^{\bar{p}} e^{-\int_0^\tau (\bar{p}, \bar{p}')(\lambda + \nu(\bar{p}^\tau)) d\tau} \left(-c + \lambda u_r^R + \nu(\bar{p}') u_{\rho(\bar{p}')}^R \right) \frac{1}{\eta(\bar{p}')} d\bar{p}' \right. \\ & \quad - (1 - p) \int_{\bar{p}^\epsilon}^{\bar{p}} e^{-\int_0^\tau (\bar{p}, \bar{p}') \nu(\bar{p}^\tau) d\tau} \left(-c + \nu(\bar{p}') u_{\rho(\bar{p}')}^L \right) \frac{1}{\eta(\bar{p}')} d\bar{p}' \\ & \quad \left. + \left(p e^{-\int_0^\tau (\bar{p}, \bar{p}^\epsilon)(\lambda + \nu(\bar{p}^\tau)) d\tau} + (1 - p) e^{-\int_0^\tau (\bar{p}, \bar{p}^\epsilon) \nu(\bar{p}^\tau) d\tau} \right) V^{\sigma^*}(p^\epsilon, \bar{p}^\epsilon) \right], \end{aligned}$$

where for any $p \in [0, 1]$, p^τ denotes the τ -ahead update of belief p , and

$$\tau(p, p') := \frac{1}{\lambda} \ln \left(\frac{p}{1-p} \frac{1-p'}{p'} \right)$$

is the time it takes for the Bayesian update of p to reach p' in the absence of R -evidence. (Note that the earlier definition of $\tau(\bar{p})$ in [Appendix A.1](#) suppresses the first argument since it was fixed to \bar{p}_0 .) Let $(\sigma_\epsilon(\cdot), \pi_\epsilon(\cdot))$ denote an ϵ -commitment solution. Then, it must satisfy the following saddle point condition: for each \bar{p}

$$\sigma_\epsilon(\bar{p}) \in \arg \max_\sigma W_\epsilon^\sigma(\pi_\epsilon(\bar{p}), \bar{p}) \tag{24}$$

$$\pi_\epsilon(\bar{p}) \in \arg \min_{p \in \mathcal{P}(\bar{p})} W_\epsilon^{\sigma_\epsilon}(p, \bar{p}). \tag{25}$$

Let $W_\epsilon(\pi_\epsilon, \bar{p}) := W_\epsilon^{\sigma_\epsilon}(\pi_\epsilon, \bar{p})$ be the saddle point value of the ϵ -commitment solution at \bar{p} .

The ϵ -commitment saddle point value $W_\epsilon(\pi_\epsilon, \bar{p})$ is characterized as follows: For any $\delta \in (0, \epsilon)$,

$$W_\epsilon(\pi_\epsilon, \bar{p})$$

$$\begin{aligned}
&= \max_{m, \nu, \rho} m(\bar{p}) U_{\rho(\bar{p})}(\pi_\epsilon) + (1 - m(\bar{p})) \left[-\pi_\epsilon \int_{\bar{p}^\delta}^{\bar{p}} e^{-\int_0^\tau (\bar{p}, \bar{p}') (\lambda + \nu(p^\tau)) d\tau} \left(-c + \lambda u_r^R + \nu(\bar{p}') u_{\rho(\bar{p}')}^R \right) \frac{1}{\eta(\bar{p}')} d\bar{p}' \right. \\
&\quad - (1 - \pi_\epsilon) \int_{\bar{p}^\delta}^{\bar{p}} e^{-\int_0^\tau (\bar{p}, \bar{p}') \nu(p^\tau) d\tau} \left(-c + \nu(\bar{p}') u_{\rho(\bar{p}')}^L \right) \frac{1}{\eta(\bar{p}')} d\bar{p}' \\
&\quad \left. + \left(\pi_\epsilon e^{-\int_0^\tau (\bar{p}, \bar{p}^\delta) (\lambda + \nu(p^\tau)) d\tau} + (1 - \pi_\epsilon) e^{-\int_0^\tau (\bar{p}, \bar{p}^\delta) \nu(p^\tau) d\tau} \right) W_{\epsilon-\delta}(\pi_\epsilon^\delta, \bar{p}^\delta) \right].
\end{aligned}$$

Following the dynamic programming principle, the RHS does not depend on any $\delta \in (0, \epsilon)$. Hence, the derivative of the RHS with respect to δ evaluated at $\delta = 0$ must equal zero. So, we obtain the standard HJB equations:²⁵

$$0 = \max_{m, \nu, \rho} G(m, \nu, \rho, \pi_\epsilon, \bar{p}, W_\epsilon, dW_\epsilon), \quad (26)$$

$$\sigma_\epsilon(\bar{p}) = (m_\epsilon(\bar{p}), \nu_\epsilon(\bar{p}), \rho_\epsilon(\bar{p})) \in \arg \max_{m, \nu, \rho} G(m, \nu, \rho, \pi_\epsilon, \bar{p}, W_\epsilon, dW_\epsilon). \quad (27)$$

Likewise, by the saddle point condition, $W_\epsilon(\pi_\epsilon, \bar{p})$ is also equivalently characterized as: for any $\delta \in (0, \epsilon)$,

$$\begin{aligned}
&W_\epsilon(\pi_\epsilon, \bar{p}) \\
&= \min_p m_\epsilon(\bar{p}) U_{\rho_\epsilon(\bar{p})}(p) + (1 - m_\epsilon(\bar{p})) \left[-p \int_{\bar{p}^\delta}^{\bar{p}} e^{-\int_0^\tau (\bar{p}, \bar{p}') (\lambda + \nu_\epsilon(p^\tau)) d\tau} \left(-c + \lambda u_r^R + \nu_\epsilon(\bar{p}') u_{\rho_\epsilon(\bar{p}')}^R \right) \frac{1}{\eta(\bar{p}')} d\bar{p}' \right. \\
&\quad - (1 - p) \int_{\bar{p}^\delta}^{\bar{p}} e^{-\int_0^\tau (\bar{p}, \bar{p}') \nu_\epsilon(p^\tau) d\tau} \left(-c + \nu_\epsilon(\bar{p}') u_{\rho_\epsilon(\bar{p}')}^L \right) \frac{1}{\eta(\bar{p}')} d\bar{p}' \\
&\quad \left. + \left(p e^{-\int_0^\tau (\bar{p}, \bar{p}^\delta) (\lambda + \nu_\epsilon(p^\tau)) d\tau} + (1 - p) e^{-\int_0^\tau (\bar{p}, \bar{p}^\delta) \nu_\epsilon(p^\tau) d\tau} \right) W_{\epsilon-\delta}(p^\delta, \bar{p}^\delta) \right].
\end{aligned}$$

Following the dynamic programming principle, the RHS does not depend on any $\delta \in (0, \epsilon)$.

²⁵Since

$$W_{\epsilon-\delta}(\pi_\epsilon^\delta, \bar{p}^\delta) = W_\epsilon(\pi_\epsilon, \bar{p}_-) + \frac{\partial W_\epsilon(\pi_\epsilon, \bar{p}_-)}{\partial p} \eta(\pi_\epsilon) \delta + \frac{\partial W_\epsilon(\pi_\epsilon, \bar{p}_-)}{\partial \bar{p}} \eta(\bar{p}) \delta + o(\delta),$$

the derivative of $W_{\epsilon-\delta}(\pi_\epsilon^\delta, \bar{p}^\delta)$ with respect to δ is:

$$\lim_{\delta \downarrow 0} \frac{W_{\epsilon-\delta}(\pi_\epsilon^\delta, \bar{p}^\delta) - W_\epsilon(\pi_\epsilon, \bar{p}_-)}{\delta} = \frac{\partial W_\epsilon(\pi_\epsilon, \bar{p}_-)}{\partial p} \eta(\pi_\epsilon) + \frac{\partial W_\epsilon(\pi_\epsilon, \bar{p}_-)}{\partial \bar{p}} \eta(\bar{p}).$$

Hence, the derivative of the RHS with respect to δ evaluated at $\delta = 0$ must equal zero. So, we obtain the standard HJB equations:

$$\pi_\epsilon \in \arg \min_{p \in \mathcal{P}(\bar{p})} G(m_\epsilon, \nu_\epsilon, \rho_\epsilon, p, \bar{p}, W_\epsilon, dW_\epsilon). \quad (28)$$

The proof is Appendix [Appendix B.4](#). We begin with a useful observation:

Lemma A.1. *Suppose a well-behaved pseudo equilibrium (σ^*, π^*) is dynamically credible. Then, for each $\bar{p} \in (0, 1)$, for any sequence such that $(\sigma_\epsilon(\bar{p}), \pi_\epsilon(\bar{p})) \rightarrow (\sigma^*(\bar{p}), \pi^*(\bar{p}))$ as $\epsilon \rightarrow 0$, $W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p}_-) \rightarrow V^{\sigma^*}(\pi^*(\bar{p}), \bar{p}_-)$, $\partial W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p}_-)/\partial p \rightarrow V_p^{\sigma^*}(\pi^*(\bar{p}), \bar{p}_-)$, and $\partial W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p}_-)/\partial \bar{p} \rightarrow V_{\bar{p}}^{\sigma^*}(\pi^*(\bar{p}), \bar{p}_-)$.*

The result now follows.

Theorem 2. *If a well-behaved pseudo equilibrium is dynamically credible, then it is an intrapersonal equilibrium. Conversely, if there is a unique intrapersonal equilibrium, then it is dynamically credible.*

Proof. Fix any \bar{p} . By the dynamic credibility of (σ^*, π^*) , for each $\bar{p} \in (0, 1)$, there exists a sequence of ϵ such that an associated ϵ -commitment solution $(\sigma_\epsilon(\bar{p}), \pi_\epsilon(\bar{p}))$ converges to $(\sigma^*(\bar{p}), \pi^*(\bar{p}))$ as $\epsilon \rightarrow 0$. Given the continuity assumption, by [Lemma A.1](#), $W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p}_-) \rightarrow V^{\sigma^*}(\pi^*(\bar{p}), \bar{p}_-)$, $\partial W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p}_-)/\partial p \rightarrow V_p^{\sigma^*}(\pi^*(\bar{p}), \bar{p}_-)$, and $\partial W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p}_-)/\partial \bar{p} \rightarrow V_{\bar{p}}^{\sigma^*}(\pi^*(\bar{p}), \bar{p}_-)$.

Next, observe that G is continuous. Hence, by Berge's maximum theorem, the maximized value is continuous and the set of maximizers must be upper hemicontinuous, which in this case should mean a closed graph. Hence, [\(11\)](#) and [\(12\)](#) must follow from [\(26\)](#) and [\(27\)](#) in the limit as $\epsilon \rightarrow 0$.

Applying the same argument, [\(13\)](#) follows from [\(28\)](#) in the limit as $\epsilon \rightarrow 0$. Again, by Berge's maximum theorem, the minimized value is continuous and the set of minimizers must be upper hemicontinuous, which implies a closed graph. Hence, the conditions must hold in the limit as $\epsilon \rightarrow 0$.

Conversely, suppose (σ^*, π^*) is intrapersonal equilibrium. Then, it satisfies [\(11\)](#), [\(12\)](#), and [\(13\)](#), for each \bar{p} . Hence, for each \bar{p} , $V^{\sigma^*}(\pi^*(\bar{p}), \bar{p})$, $V_p^{\sigma^*}(\pi^*(\bar{p}), \bar{p})$ and $V_{\bar{p}}^{\sigma^*}(\pi^*(\bar{p}), \bar{p}_-)$ are well-defined and continuous at $(\pi^*(\bar{p}), \bar{p})$. Take a sequence (subsequence if necessary) of $\epsilon > 0$ such that $(\sigma_\epsilon, \pi_\epsilon, W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p}), \partial W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p})/\partial p, \partial W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p})/\partial \bar{p})$ all converge to some limit, say $(\sigma_0, \pi_0, W_0(\pi_0(\bar{p}), \bar{p}), \tilde{W}_{0,p}(\pi_0(\bar{p}), \bar{p}), \tilde{W}_{0,\bar{p}}(\pi_0(\bar{p}), \bar{p}))$. The reason that such a sequence can

be found is as follows. First, $(W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p}), \partial W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p})/\partial p)$ is bounded within a compact set $[\min\{u_\ell^R, u_r^L\}, \max\{u_r^R, u_\ell^L\}] \times [u_\ell^R - u_\ell^L, u_r^R - u_r^L]$. This together with (27) means that $\nu_\epsilon(\bar{p})$ can be bounded within $[0, \bar{\nu}]$ for some large $\bar{\nu}$. The boundedness of $\nu_\epsilon(\bar{p})$ in turn means that $\partial W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p})/\partial \bar{p}$ can be bounded within an compact set. This proves that the converging sequence can be found for each \bar{p} . Next, since $\partial^2 W_\epsilon(\pi_\epsilon(\bar{p}), \bar{p})/\partial p \partial \bar{p} = 0$ for all ϵ (which follows from the fact that $W_\epsilon(\cdot, \bar{p})$ is linear), we have $\tilde{W}_{0,p}(\pi_0(\bar{p}), \bar{p}) = \partial W_0(\pi_0(\bar{p}), \bar{p})/\partial p$ and $\tilde{W}_{0,\bar{p}}(\pi_0(\bar{p}), \bar{p}) = \partial W_0(\pi_0(\bar{p}), \bar{p})/\partial \bar{p}$. It is immediate that the limiting strategy profile (σ_0, π_0) is a pseudo equilibrium. Also using the same argument as above, we conclude that the limit strategy profile (σ_0, π_0) , together with the value function W_0 , satisfies (11), (12), and (13) for each \bar{p} . Therefore, (σ_0, π_0) is also an intrapersonal equilibrium. The uniqueness of the equilibrium then means that $(\sigma_0, \pi_0) = (\sigma^*, \pi^*)$, so for each \bar{p} , $(\sigma^*(\bar{p}), \pi_0^*(\bar{p}))$ is a limit of some ϵ -commitment solution $(\sigma_\epsilon(\bar{p}), \pi_\epsilon(\bar{p}))$ as $\epsilon \rightarrow 0$. Hence, (σ^*, π^*) is dynamically credible. \square

A.3 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 A.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

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

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. \square

For the proof, for each region of \bar{p} 's, we will specify the strategy profile (σ, π) and the associated value $V^\sigma(p, \bar{p})$ 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 $c \in (\underline{c}, \bar{c}]$ but $\Delta < \Delta_c$, where Δ_c will be specified later (in Region 3 in Appendix A.3.1).

We call this case *main case*. We will first treat the main case in [Appendix A.3.1](#). The remaining cases will be treated in [Appendix A.3.2](#) and [Appendix B.8](#).

A.3.1 Main Case

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

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

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

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

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

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 $\bar{p} \leq \bar{p}_1$. Hence $\sigma = (1, 0, 0)$ is a maximizer and both (11) and (12) are satisfied for $V^\sigma(p, \bar{p}) = U_\ell(p)$. Setting $m(\bar{p}) = 1$ and $\rho(\bar{p}) = 0$, the objective in (13) is independent of p and $\pi(\bar{p}) = \bar{p}$ is a minimizer. Hence we have shown that for the posited strategy profile (σ, π) satisfy the equilibrium conditions, (7)–(13) for all $\bar{p} \in [0, \bar{p}_1]$.

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

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

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

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

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

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

$$V^\sigma(p, \bar{p}) := \Phi(\bar{p}) + (p - \bar{p})\Phi'(\bar{p}), \text{ for all } \bar{p} \in (\bar{p}_1, \bar{p}_2]. \quad (29)$$

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

$$V^\sigma(\pi(\bar{p}), \bar{p}) = \Phi(\bar{p}), \quad V_p^\sigma(\pi(\bar{p}), \bar{p}) = \Phi'(\bar{p}), \quad V_{\bar{p}}^\sigma(\pi(\bar{p}), \bar{p}) = 0.$$

Therefore (11) simplifies to

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

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

$$-c + \bar{p}\lambda(u_r^R - \Phi(\bar{p})) + \Phi'(\bar{p})\eta(\bar{p}) = 0$$

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

$$\begin{aligned} & -c + p\lambda(u_r^R - (\Phi(\bar{p}) + (p - \bar{p})\Phi'(\bar{p}))) - \lambda p(1 - p)\Phi'(\bar{p}) - \lambda \bar{p}(1 - \bar{p})(p - \bar{p})\Phi''(\bar{p}) \\ & = -c + p\lambda(u_r^R - (\Phi(\bar{p}) + (1 - \bar{p})\Phi'(\bar{p}))) - \lambda \bar{p}(1 - \bar{p})(p - \bar{p})\Phi''(\bar{p}). \end{aligned}$$

Differentiating this with respect to p yields

$$\lambda(u_r^R - (\Phi(\bar{p}) + (1 - \bar{p})\Phi'(\bar{p}))) - \lambda \bar{p}(1 - \bar{p})\Phi''(\bar{p}).$$

Differentiating the ODE for the Bayesian Value function (see (2)) with respect to \bar{p} we obtain that this expression is equal to zero. Hence $\pi(\bar{p})$ 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 $\bar{p} \in (\bar{p}_1, \bar{p}_2]$.

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

This region exists only when the ambiguity is sufficiently large so that $V^\sigma(\underline{p}(\bar{p}_2), \bar{p}_2) < U_\ell(\underline{p}(\bar{p}_2))$. If this inequality is reversed, then we set $\bar{p}_3 = \bar{p}_2$, and Region 3 is empty. Assuming the inequality, we will specify the upper bound \bar{p}_3 , the critical cost \underline{c} , and ambiguity level Δ_c referred to in the statement of the theorem. Finally, we will specify $\nu(\bar{p})$ and $\pi(\bar{p})$ fully for \bar{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 $\nu(\bar{p})$. Meanwhile, nature chooses belief $\pi(\bar{p}) \in \mathcal{P}(\bar{p})$. Here, we specify (ν, π) precisely, together with the value $V^\sigma(p, \bar{p})$ associated with the strategy. We first construct a value function $V(p, \bar{p})$ that satisfies the equilibrium conditions, (7)–(13), given the candidate strategy σ . We will then establish that the constructed function $V(p, \bar{p})$ indeed coincides with the value of σ .

To begin, fix any $\bar{p} > \bar{p}_2$. 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 $\pi(\bar{p})$ to be nature’s choice satisfying (7), it is sufficient, and will be seen also necessary, to have $V_p(p, \bar{p}) = 0$. Namely,

$$V(p, \bar{p}) = \hat{V}(\bar{p}), \forall p. \quad (30)$$

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 (30) yields (16).

It now remains to specify the function $\hat{V}(\bar{p})$. To this end, we use (15), (16) and (30) to simplify (11) and obtain the differential equation ($\widehat{\text{ODE}}$). Together with the boundary condition that $\hat{V}(\bar{p}_2) = \Phi(\bar{p}_2)$, ($\widehat{\text{ODE}}$) admits a unique solution:²⁶

$$\hat{V}(\bar{p}) := \frac{C}{\left(\frac{1-\bar{p}}{\bar{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}} + C} u_1 + \frac{\left(\frac{1-\bar{p}}{\bar{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}}}{\left(\frac{1-\bar{p}}{\bar{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}} + C} u_2, \quad (31)$$

²⁶The derivation of the solution appears in [Appendix B.6](#).

where

$$C := \frac{u_2 - \Phi(\bar{p}_2)}{\Phi(\bar{p}_2) - u_1} \left(\frac{1 - \bar{p}_2}{\bar{p}_2} \right)^{\frac{u_2 - u_1}{\delta_\ell}}, \quad (32)$$

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

$$u_{1,2} := \frac{u_r^R + u_\ell^L}{2} \pm \sqrt{\left(\frac{u_r^R - u_\ell^L}{2} \right)^2 + \frac{c}{\lambda} \delta_\ell}. \quad (33)$$

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

$$\hat{V}(\bar{p}_3) = U_\ell(\underline{p}(\bar{p}_3)). \quad (34)$$

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

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

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

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

Proof. The proof can be found in [Appendix B.5](#) below. □

[Lemma A.3](#)-(i) implies that the value segment shifts up (as well as shifts left) as \bar{p} drifts down.²⁷ Next, [Lemma A.3](#)-(ii) implies that \bar{p}_3 is well defined and above \bar{p}_2 whenever the region exists. Last, [Lemma A.3](#)-(iii) ensures that the value segment remains above \hat{u} for all $\bar{p} \in (\bar{p}_2, \bar{p}_3]$,

²⁷One can also see that V^σ is continuous, including at the boundary. This follows from the fact that $\hat{V}(\bar{p}_2) = \Phi(\bar{p}_2)$, and that $\bar{p}_2 = p_*$ and $\Phi'(p_*) = 0$.

if either $c \leq \underline{c}$ or if $c \in [\underline{c}, \bar{c})$ 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 $V(p, \bar{p})$ is constructed from HJB equations along with (7). We now claim the value function indeed represents the value of the candidate strategy σ :

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

Proof. See Appendix B.7. □

Verification of equilibrium conditions. To verify (11)–(12), note that the definition of $\pi(\bar{p})$ in (15) implies that $U_\ell(\pi(\bar{p})) - V^\sigma(\pi(\bar{p}), \bar{p}) = 0$. We have already observed from Lemma A.3-(iii) that $V^\sigma(\cdot, \bar{p}) \geq \hat{u}$ for all $\bar{p} \in (\bar{p}_2, \bar{p}_3]$. This further implies that $\pi(\bar{p}) \leq \hat{p}$ for $\bar{p} \in (\bar{p}_2, \bar{p}_3]$. Therefore $\rho(\bar{p}) = 0$ is a maximizer in (11)–(12). By (14), $\nu(\bar{p})$ as defined in (16) is a maximizer as well. It remains to show that $m(\bar{p}) = 0$ 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 (30) to simplify the coefficient of $(1 - m)$ to:

$$-c + \pi(\bar{p})\lambda(u_r^R - \hat{V}(\bar{p})) - \lambda\bar{p}(1 - \bar{p})\hat{V}(\bar{p}), \quad (35)$$

which vanishes precisely because \hat{V} solves (ODE). We thus conclude that $m(\bar{p}) = 0$ is a minimizer in (12). We have thus verified (31) and (12).

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

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

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

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

Here, \bar{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 $\bar{p} > \bar{p}_3$ until state \bar{p}_3 is reached. To compute the associated value $V^\sigma(p, \bar{p})$,

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 $v_{**} := \hat{V}(\bar{p}_3)$:²⁸

Auxiliary Problem: Imagine a hypothetical Bayesian DM with any belief $p \geq \underline{p}(\bar{p}_3)$ who at each instant may experiment or stops. She may experiment until either breakthrough occurs or her belief reaches $\underline{p}(\bar{p}_3)$. If she stops at any point, she collects the payoff $\hat{V}(\bar{p}_3)$ (independent of p). The optimal value of this stopping problem is:

$$\Psi(p) = \max_{p' \in [\underline{p}(\bar{p}_3), p]} \phi(p; p', v_{**}). \quad (36)$$

Let p_{**} denote the optimal stopping belief. We note that $\Psi(p_{**}) = v_{**}$ and $\Psi'(p_{**}) \geq 0$, with equality holding whenever $p_{**} > \underline{p}(\bar{p}_3)$ (a consequence of smooth pasting). We can easily see that $p_{**} \in [\underline{p}(\bar{p}_3), \bar{p}_3]$.²⁹

Fix any state $\bar{p} > \bar{p}_3$. We ask: *for what belief p is the continuation strategy $\sigma_{\bar{p}}$ optimal over Region 3?* Recall $\sigma_{\bar{p}}$ prescribes: “experiment for the duration of $\tau(\bar{p}, \bar{p}_3)$ and, absent breakthrough by the end of the experimentation, stop and collect $\hat{V}(\bar{p}_3)$,” 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, $q(\bar{p})$, such that

$$\tau(\bar{p}, \bar{p}_3) = \tau(q(\bar{p}), p_{**}). \quad (37)$$

In words, it is the belief that would be updated to p_{**} after the prescribed duration $\tau(\bar{p}, \bar{p}_3)$ of experimentation.³⁰ Since p_{**} is the optimal stopping belief in the above Auxiliary Problem, a DM with belief $q(\bar{p})$ finds the prescribed strategy optimal at state \bar{p} and thus will realize the

²⁸Recall that if Region 3 does not exist, then $\bar{p}_3 = \bar{p}_2 = p_*$, so $\hat{V}(\bar{p}_3) = \hat{V}(\bar{p}_2) = \Phi(\bar{p}_2)$, as defined before. All subsequent results hold since $V_p(p, \bar{p}_2) = \Phi'(\bar{p}_2) = \Phi'(p_*) = 0$.

²⁹To see that $p_{**} < \bar{p}_3$, it suffices to show that $p_{**} < p_*$ since $p_* \leq \bar{p}_3$. 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 $\bar{p}_3 \geq p_*$.

³⁰Obviously, $q(\bar{p}_3) = p_{**}$.

value of $\Psi(q(\bar{p}))$. Consequently,

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

What about a DM with belief $p \neq q(\bar{p})$? 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 \bar{p} :

$$V^\sigma(p, \bar{p}) := \Psi(q(\bar{p})) + (p - q(\bar{p}))\Psi'(q(\bar{p})). \quad (38)$$

We define \bar{p}_4 to be such that

$$V^\sigma(\underline{p}(\bar{p}_4), \bar{p}_4) = U_r(\underline{p}(\bar{p}_4)). \quad (39)$$

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

$$V_p^\sigma(p, \bar{p}) = \Psi'(q(p)) \geq 0,$$

where the equality follows from (38) and the inequality from the convexity of $\Psi(\cdot)$ and $\Psi'(p_{**}) \geq 0$.

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

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

Proof. Note $V^\sigma(\underline{p}(\bar{p}_3), \bar{p}_3) = \hat{V}(\bar{p}_3) \geq U_\ell(\underline{p}(\bar{p}_3))$.³¹ Next, note that $\underline{p}(\bar{p}_3) > \underline{p}_r^B$ and $V^\sigma(\underline{p}(\bar{p}_3), \bar{p}_3) = \phi(\underline{p}(\bar{p}_3); \underline{p}(\bar{p}_3), U_\ell(\underline{p}(\bar{p}_3)))$. Since, by Lemma A.2, $\phi(p; p, U_\ell(p))$ can only cross $U_\ell(p)$ from below, we have

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

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

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

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

³¹If $V^\sigma(\underline{p}(\bar{p}_2), \bar{p}_2) < U_\ell(\underline{p}(\bar{p}_2))$, then the inequality holds with equality by definition. If $V^\sigma(\underline{p}(\bar{p}_2), \bar{p}_2) \geq U_\ell(\underline{p}(\bar{p}_2))$, then the inequality follows from the fact that $\bar{p}_3 = \bar{p}_2$.

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

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

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

where we suppress the argument of $\underline{p}(\bar{p})$ 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

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

where we have used [\(38\)](#) for the first equality and $q'(\bar{p}) = \eta(q(\bar{p}))/\eta(\bar{p})$ for the second equality.³² 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 $m(\bar{p}) = \nu(\bar{p}) = 0$, 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:

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

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

³²This holds since $\tau(q(\bar{p}), q(\bar{p}_3)) = \tau(\bar{p}, \bar{p}_3)$ implies that $\ln(q(\bar{p})/(1 - q(\bar{p}))) - \ln(p_*/(1 - p_*)) = \ln(\bar{p}/(1 - \bar{p})) - \ln(\bar{p}_3/(1 - \bar{p}_3))$. Hence the difference in log-likelihood ratios of $q(\bar{p})$ and \bar{p} is constant

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

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

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 $p = \underline{p}(\bar{p})$ is a minimizer in (13) is satisfied.

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

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

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

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

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

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

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

A.3.2 Case 1 with $c \geq \underline{c}$ and $\Delta > \Delta_c$

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

$$(m(\bar{p}), \nu(\bar{p}), \rho(\bar{p})) = \left(1, 0, \frac{\delta_\ell}{\delta_r + \delta_\ell}\right), \quad (40)$$

and $\pi(\bar{p}) = \hat{p}$, for $\bar{p} \in [\bar{p}_3, \bar{p}_4)$, where \bar{p}_3 is now set at \bar{p}'_3 , which satisfies $\hat{V}(\bar{p}'_3) = \hat{u}$. We note that this new \bar{p}'_3 is smaller than the original \bar{p}_3 and is still larger than \bar{p}_2 .³³ \bar{p}_4 is now set at \bar{p}'_4 , which

³³Recall $\hat{V}(\cdot)$ is strictly decreasing. Since $\hat{V}(\bar{p}_3) < \hat{u}$ when $c \in (\underline{c}, \bar{c})$ and $\Delta > \Delta_c$, we have $\bar{p}'_3 < \bar{p}_3$. We also have $\bar{p}'_3 > \bar{p}_2$ since $\hat{V}(\bar{p}_2) = \Phi(\bar{p}_2) > \max\{U_\ell(\bar{p}_2), U_r(\bar{p}_2)\} \geq \hat{u}$.

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

It is straightforward to verify that $V^\sigma(p, \bar{p}) = \hat{u}$ together with $\pi(\bar{p}) = \hat{p}$ and (40) satisfy (7)–(13) for $\bar{p} \in (\bar{p}_3, \bar{p}_4)$. Indeed, inserting $\pi(\bar{p}) = \hat{p}$ and $V^\sigma(p, \bar{p}) = \hat{u}$ as well as $V_p^\sigma(p, \bar{p}) = 0$ and $V_{\bar{p}}^\sigma(p, \bar{p}) = 0$, 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 $\bar{p} = \bar{p}_3$, we have $V^\sigma(p, \bar{p}_-) = \hat{V}^\sigma(\bar{p}_3) = \hat{u}$ and $V_p^\sigma(p, \bar{p}_-) = 0$, and $V_{\bar{p}}^\sigma(p, \bar{p}_-) < 0$, 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 (40), and $V^\sigma(p, \bar{p}_-) = V^\sigma(p, \bar{p}) = \hat{u}$ in the objective in (13), we see that the objective is equal to zero and hence $\pi(\bar{p}) = \hat{p}$ is a minimizer. Finally, $V_p^\sigma(p, \bar{p}) = 0$ implies that (7) holds.

B Online Appendix

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 2.*

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:

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

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

$$c = \lambda p (u_r^R - \phi(p)) + \phi'(p) \eta(p). \quad (41)$$

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

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

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

$$\xi U_r' + (1 - \xi) \Phi'(p_r^B) = 0; \tag{42}$$

- $c \geq \bar{c}$: the DM randomizes between actions r and ℓ with probabilities \hat{p} and $1 - \hat{p}$, respectively, where \hat{p} 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.³⁴

Since the DM's strategy is Bayesian optimal against nature's choice p_{min} ,³⁵ 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 as a function of nature's choice $p \in [0, 1]$:

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

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,

³⁴ 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

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

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.

³⁵Note 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 $c < \bar{c}$.

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, $V_0'(\cdot) = \Phi^{*'}(p_{min})$, 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 $F_t : [0, \infty] \rightarrow [0, 1]$, for each $t \geq 0$,³⁶ 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 stopping plan from then on.

We require a couple of technical assumptions:

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

³⁶We let the stopping time be ∞ if the DM never stops.

(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 $\tilde{\rho} : [0, \infty) \rightarrow [0, 1]$ 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 $\tilde{\sigma} : [0, \infty) \rightarrow [0, \infty) \times [0, 1] \times [0, 1]$, $\tilde{\sigma}_t = (\tilde{\nu}_t, \tilde{m}_t, \tilde{\rho}_t)$, 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.

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 $\tilde{\Sigma}$ 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 $\tilde{\sigma} \in \tilde{\Sigma}$. Conversely, any admissible strategy $\tilde{\sigma} \in \tilde{\Sigma}$ 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 $s < t$,

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

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 $(\tilde{\nu}, \tilde{m}, \tilde{\rho}) \in \tilde{\Sigma}$.

Conversely, fix any $\tilde{\sigma} \in \tilde{\Sigma}$. 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, $\{F_t\} \in \mathcal{F}$. □

B.4 Proof of Lemma A.1

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

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

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

Next, consider

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

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

Finally, consider

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

$$= \left(\pi_\epsilon e^{-\int_0^{\tau(\bar{p}, \bar{p}^\epsilon)} (\lambda + \nu_\epsilon(p^\tau)) d\tau} + (1 - \pi_\epsilon) e^{-\int_0^{\tau(\bar{p}, \bar{p}^\epsilon)} \nu_\epsilon(p^\tau) d\tau} \right) V_{\bar{p}}^{\sigma^*}(\pi^\epsilon, \bar{p}^\epsilon) \frac{\eta(\bar{p}^\epsilon)}{\eta(\bar{p})}$$

$$\rightarrow V_{\bar{p}}^{\sigma^*}(\pi^*, \bar{p}_-) \text{ as } \epsilon \rightarrow 0,$$

where the second line holds because of the HJB for ϵ -commitment solution, i.e., by (27).

B.5 Proof of Lemma A.3

We first prove (i). To show $\hat{V}(\bar{p}) \in (u_1, u_2)$, it suffices to show that $C > 0$ since in this case $\hat{V}^\sigma(\bar{p})$ is a convex combination of u_1 and u_2 . We show that $\Phi(\bar{p}_2) \in (u_1, u_2)$ which implies that $C > 0$. To see that $\Phi(\bar{p}_2) > u_1$, 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 $\Phi(\bar{p}_2) < u_2$, note that $u_2 \geq \frac{1}{2}(u_r^R + u_\ell^L) + \frac{1}{2}|u_r^R - u_\ell^L| = u_r^R \vee u_\ell^L > \Phi(\bar{p}_2)$. Since $u_1 < u_2$, the term $\left(\frac{1-\bar{p}}{\bar{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}}$ is decreasing in \bar{p} and therefore $\hat{V}^{\sigma'}(\bar{p}) < 0$.

We next prove (ii). By the assumption, we have $V^\sigma(\underline{p}(\bar{p}_2), \bar{p}_2) < U_\ell(\underline{p}(\bar{p}_2))$. As $\bar{p} \rightarrow 1$, $\hat{V}(\bar{p}) \rightarrow u_1$, whereas $U_\ell(\underline{p}(\bar{p})) \rightarrow u_\ell^R$, which by assumption is less than u_1 .³⁷ Hence, by the intermediate value theorem, there exists $\bar{p}_3 \in (\bar{p}_2, 1)$ such that $V^\sigma(\underline{p}(\bar{p}_3), \bar{p}_3) = \hat{V}(\bar{p}_3) = U_\ell(\underline{p}(\bar{p}_3))$. The fact that such a \bar{p}_3 is unique follows from the fact that $\hat{V}(\cdot)$ can only cross $U_\ell(\underline{p}(\cdot))$ from below:

$$\left. \frac{d(\hat{V}^\sigma(\bar{p}) - U_\ell(\underline{p}(\bar{p})))}{d\bar{p}} \right|_{U_\ell(\underline{p}(\bar{p})) = \hat{V}(\bar{p})} = \frac{(u_r^R - \underline{p}u_\ell^R - (1-\underline{p})u_\ell^L)(u_\ell^L - \underline{p}u_\ell^R - (1-\underline{p})u_\ell^L)}{u_\ell^L - u_\ell^R} - \frac{c}{\lambda} - (u_\ell^R - u_\ell^L) \underline{p}(1-\underline{p}) > 0$$

$$\iff (u_r^R - \underline{p}u_\ell^R - (1-\underline{p})u_\ell^L) \underline{p} - \frac{c}{\lambda} - (u_\ell^R - u_\ell^L) \underline{p}(1-\underline{p}) > 0$$

$$\iff (u_r^R - u_\ell^R) \underline{p} - \frac{c}{\lambda} > 0$$

$$\iff \underline{p} > p_\ell^B,$$

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

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

$$u_1 > \hat{u} \iff \frac{u_r^R + u_\ell^L}{2} - \sqrt{\left(\frac{u_r^R - u_\ell^L}{2}\right)^2 + \frac{c}{\lambda} \delta_\ell} > \frac{u_r^R u_\ell^L - u_r^L u_\ell^R}{\delta_r + \delta_\ell}$$

³⁷This follows from the fact that $\lim_{\bar{p} \rightarrow 1} \left(\frac{1-\bar{p}}{\bar{p}}\right)^{\frac{u_2-u_1}{\delta_\ell}} = 0$. Note that $\underline{p}(\bar{p}) \rightarrow 1$ as $\bar{p} \rightarrow 1$.

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

For the last statement, fix any $c \in (\underline{c}, \bar{c})$. Then, $u_1 < \hat{u}$, as shown above. Rewrite (34)

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

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

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

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

B.6 Solution to $\widehat{\text{ODE}}$

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

$$\begin{aligned}
&\bar{p}(1 - \bar{p})\hat{V}^{\sigma'}(\bar{p}) - \frac{(u_r^R - \hat{V}^\sigma(\bar{p}))(u_\ell^L - \hat{V}^\sigma(\bar{p}))}{\delta_\ell} + \frac{c}{\lambda} = 0 \\
&\bar{p}(1 - \bar{p})\hat{V}^{\sigma'}(\bar{p}) - \frac{u_r^R u_\ell^L - (u_r^R + u_\ell^L)\hat{V}^\sigma(\bar{p}) + (\hat{V}^\sigma(\bar{p}))^2}{\delta_\ell} + \frac{c}{\lambda} = 0 \\
&\bar{p}(1 - \bar{p})\hat{V}^{\sigma'}(\bar{p}) + \frac{u_r^R + u_\ell^L}{\delta_\ell} \hat{V}^\sigma(\bar{p}) - \frac{1}{\delta_\ell} (\hat{V}^\sigma(\bar{p}))^2 + \frac{c}{\lambda} - \frac{u_r^R u_\ell^L}{\delta_\ell} = 0.
\end{aligned}$$

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

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

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

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

. 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

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

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

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 $\hat{V}^\sigma(\bar{p})$:

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

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

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

where

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

We note that

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

Boundary condition:

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

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

B.7 Proof of Lemma A.4

The value of strategy σ is given by:

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

Using the change of variables $\tau = \tau(\bar{p}, \bar{p}'')$ and hence $\bar{p}_\tau = \bar{p}''$ and $d\tau = -\frac{1}{\lambda \bar{p}'' (1 - \bar{p}'')} d\bar{p}''$ we can rewrite this as follows.³⁸

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

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

³⁸Remember that $\bar{p}_t = \frac{\bar{p}_0}{\bar{p}_0 + (1 - \bar{p}_0) e^{\lambda t}}$. Setting $\bar{p}_0 = \bar{p}$ and $\bar{p}_t = \bar{p}'$ we get

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

Hence,

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

defined $V(p, \bar{p})$ in Region 3 so that $V_p(p, \bar{p}) = 0$ and that

$$c = \lambda p (u_r^R - V(p, \bar{p})) + \nu(\bar{p}) (U_\ell(\bar{p}) - V(p, \bar{p})) + \eta(p) V_p(p, \bar{p}) + \eta(\bar{p}) V_{\bar{p}}(p, \bar{p}) \quad (43)$$

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

To show that $V(p, \bar{p}) = V^\sigma(p, \bar{p})$, substitute $p = 0$ and $p = 1$ in (43) to obtain two ODEs:

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

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

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

and

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

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

B.8 Case 2: $\Phi'(p_*) < 0$

Theorem 3. *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 $c \leq \underline{c}$ such that the intrapersonal equilibrium is described by:*

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

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

with $m(\bar{p}_2) \in (0, 1)$, $\nu^*(\bar{p}) > 0$ and $\pi^*(\bar{p}) \in (\underline{p}(\bar{p}), \bar{p})$, for $\bar{p} \in (\bar{p}_2, \bar{p}_3)$, and $\bar{p}_2 < \bar{p}_3$ 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 $\bar{p} \in [\bar{p}_3, \bar{p}_4)$, where

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

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

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

$$\begin{aligned} V^\sigma(p, \bar{p}_2) &= U_r(\bar{p}_2) + (p - \bar{p}_2) [m(\bar{p}_2)U'_r(\bar{p}_2) + (1 - m(\bar{p}_2))V_p^\sigma(p, \bar{p}_2^-)] \\ &= U_r(\bar{p}_2) + (p - \bar{p}_2) [m(\bar{p}_2)U'_r(\bar{p}_2) + (1 - m(\bar{p}_2))\Phi'(\bar{p}_2)]. \end{aligned}$$

To understand the first line, note that conditional on stopping, the DM gets a payoff of $U_r(p) = U_r(\bar{p}_2) + (p - \bar{p}_2)U'_r(\bar{p}_2)$, and conditional on not stopping, she gets $V^\sigma(p, \bar{p}_2^-) = V^\sigma(\bar{p}_2^-, \bar{p}_2^-) + (p - \bar{p}_2)V_p^\sigma(p, \bar{p}_2^-) = U_r(\bar{p}_2) + (p - \bar{p}_2)\Phi'(\bar{p}_2)$.³⁹

³⁹Here we have to distinguish carefully between $V^\sigma(p, \bar{p}_2^-)$, $V_p^\sigma(p, \bar{p}_2^-)$ and $V^\sigma(p, \bar{p}_2)$, $V_p^\sigma(p, \bar{p}_2)$.

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

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

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

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

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

The terms inside the square brackets are equal to zero (from (41)), so (11) holds, and $\sigma(\bar{p}_2) = (m(\bar{p}_2), 0, 1)$ is a maximizer of (12). Finally, note that the objective in (13) is

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

where we have used (29) to obtain the second line. Differentiating this with respect to p yields

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

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

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

B.9 Uniqueness

Here we prove the intrapersonal equilibrium, $\sigma = (\nu, \mu, \rho)$, together with nature's choice π , described in [Theorem 1](#) and [Theorem 3](#), is unique for the case $c < \underline{c}$ and $c \geq \bar{c}$. 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

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

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

1. $\tilde{p} \in [0, \bar{p}_1)$: Recall $m(\bar{p}) = 1$ and $\rho(\bar{p}) = 0$ for all $\bar{p} \in [0, \bar{p}_1)$. Assuming $\tilde{p} < \bar{p}_1$, the admissibility of $\tilde{\sigma}$ means that there must exist an open interval $(\tilde{p}, \tilde{p} + \varepsilon) \subset [0, \bar{p}_1]$ such that $\tilde{m}(\bar{p}) = 0$ for all $\bar{p} \in (\tilde{p}, \tilde{p} + \varepsilon)$.⁴⁰ Consider $\bar{p} \in (\tilde{p}, \tilde{p} + \varepsilon)$. Since $V^{\tilde{\sigma}}(p, \bar{p}) < \Phi^*(p)$ for all $p \in [\underline{p}(\bar{p}), \bar{p}]$ and since $\Phi^*(p) \leq U_\ell(p)$ for all $p \leq \bar{p}_1 = p_r^B$, we have $V^{\tilde{\sigma}}(p, \bar{p}) < U_\ell(p), \forall p \in [\underline{p}(\bar{p}), \bar{p}]$ and hence $V^{\tilde{\sigma}}(\tilde{\pi}(\bar{p}), \bar{p}) < U_\ell(\tilde{\pi}(\bar{p}))$, which violates condition (8) of [Definition 1](#).
2. $\tilde{p} \in [\bar{p}_1, \bar{p}_2)$: Recall $m(\bar{p}) = \nu(\bar{p}) = 0$ for all $\bar{p} \in (\bar{p}_1, \bar{p}_2)$.
 - (a) Suppose there is some $\varepsilon > 0$ such that $\tilde{m}(\bar{p}) = 0$ for all $\bar{p} \in [\tilde{p}, \tilde{p} + \varepsilon)$. We then have $V^{\tilde{\sigma}}(p, \tilde{p}) = V^\sigma(p, \tilde{p})$ for all p , with the value $V^{\tilde{\sigma}}(p, \bar{p})$ being continuous in \bar{p} at $\bar{p} = \tilde{p}$. Since $V_p^\sigma(p, \bar{p}) < 0$, it follows $V_p^{\tilde{\sigma}}(p, \bar{p}) < 0$ and hence $\tilde{\pi}(\bar{p}) = \bar{p}$ for $\bar{p} \geq \tilde{p}$ sufficiently close to \tilde{p} . Next, since $V^\sigma(\bar{p}, \bar{p}) > U_\rho(\bar{p})$ for all $\bar{p} \in (\bar{p}_1, \bar{p}_2)$ and all $\rho \in [0, 1]$, the same property holds under strategy $\tilde{\sigma}$ for $\bar{p} \geq \tilde{p}$ sufficiently close to \tilde{p} , that is, $V^{\tilde{\sigma}}(\bar{p}, \bar{p}) > U_\rho(\bar{p})$ for all $\rho \in [0, 1]$. Condition (9) of [Definition 1](#) then implies $\tilde{\nu}(\bar{p}) = \nu(\bar{p}) = 0$ for all \bar{p} belonging to a right neighborhood of \tilde{p} , which is a contradiction.
 - (b) Suppose next that for every $\varepsilon > 0$ there is some $\bar{p} \in [\tilde{p}, \tilde{p} + \varepsilon)$ such that $\tilde{m}(\bar{p}) > 0$. Since $\tilde{p} < \bar{p}_2 = p_* < \hat{p}$, we have $\bar{p} < \hat{p}$ for all $\bar{p} \geq \tilde{p}$ sufficiently close to \tilde{p} . Hence, by condition (10), if $\tilde{m}(\bar{p}) > 0$ for such values of \bar{p} , then $\tilde{\rho}(\bar{p}) = 0$. We start by considering the DM's strategy at state \tilde{p} . With probability $\tilde{m}(\tilde{p})$ the DM stops and

⁴⁰Our admissibility restrictions rule out the possibilities of $m(\bar{p}) < 1$ on an isolated point as well as $m(\bar{p}) \in (0, 1)$ for all \bar{p} belonging to an interval of states.

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

$$V^{\tilde{\sigma}}(p, \tilde{p}) = \tilde{m}(\tilde{p})U_{\tilde{\rho}(\tilde{p})}(p) + (1 - \tilde{m}(\tilde{p}))V^\sigma(p, \tilde{p}), \quad (44)$$

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 $\bar{p} \in (\tilde{p}, \tilde{p} + \varepsilon)$ such that $\tilde{m}(\bar{p}) > 0$ 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 $\tilde{m}(\bar{p}) = 1$ (and $\tilde{\rho}(\bar{p}) = 0$). For all states \bar{p} belonging to this interval, we thus have $V_p^{\tilde{\sigma}}(p, \bar{p}) = U'_\ell(p) < 0$ and hence $\tilde{\pi}(\bar{p}) = \bar{p}$. The HJB functional then simplifies to

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

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

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

- (a) Suppose again there is some $\varepsilon > 0$ such that $\tilde{m}(\bar{p}) = 0$ for all $\bar{p} \in [\tilde{p}, \tilde{p} + \varepsilon)$. For every $\varepsilon' \in (0, \varepsilon)$, we can then find some $\bar{p} \in [\tilde{p}, \tilde{p} + \varepsilon')$ such that $\tilde{\nu}(\bar{p}) \neq \nu(\bar{p})$.⁴¹ As before, we have $V^{\tilde{\sigma}}(p, \tilde{p}) = V^\sigma(p, \tilde{p})$, with $V^{\tilde{\sigma}}(p, \bar{p})$ being continuous in \bar{p} at $\bar{p} = \tilde{p}$. If $\tilde{\nu}(\bar{p}) > 0$ on a right neighborhood of \tilde{p} , the same HJB conditions pinning down $\nu(\bar{p})$ must hold. Uniqueness of the solution of these conditions then implies $\tilde{\nu}(\bar{p}) = \nu(\bar{p})$ for all \bar{p} belonging to this neighborhood. We are left with the following possibility: for every $\varepsilon' \in (0, \varepsilon)$, we can find some $\bar{p} \in [\tilde{p}, \tilde{p} + \varepsilon')$ such that $\tilde{\nu}(\bar{p}) = 0$. As $\varepsilon' \rightarrow 0$, the probability that R -evidence arrives within time $\tau(\bar{p}, \tilde{p})$ vanishes, so both $V_p^{\tilde{\sigma}}(p, \bar{p})$ and $V_{p\bar{p}}^{\tilde{\sigma}}(p, \bar{p})$ vanish (recall that at $\bar{p} = \tilde{p}$, 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 \bar{p} such that $\tilde{\nu}(\bar{p}) = 0$. Nature thus chooses $\tilde{\pi}(\bar{p}) = \underline{p}(\bar{p})$. Since however $U_\ell(\underline{p}(\bar{p})) > V^{\tilde{\sigma}}(\underline{p}(\bar{p}), \bar{p})$ for \bar{p} sufficiently close to \tilde{p} , this violates condition

⁴¹The possibility where there is an $\varepsilon' \in (0, \varepsilon)$ such that $\tilde{\nu}(\bar{p}) = \nu(\bar{p})$ for all $\bar{p} \in [\tilde{p}, \tilde{p} + \varepsilon')$, but then $\tilde{\rho}(\bar{p}) \neq \rho(\bar{p}) = 0$ for \bar{p} close or equal to \tilde{p} would clearly violate the HJB conditions.

(9).

(b) Consider next the case where for every $\varepsilon > 0$ there is some $\bar{p} \in [\tilde{p}, \tilde{p} + \varepsilon)$ such that $\tilde{m}(\bar{p}) > 0$. As before, we start by showing $\tilde{m}(\tilde{p}) = 0$. At state \tilde{p} the DM's value is described by (44).

- 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 $\tilde{\rho}(\tilde{p}) = 0$.⁴² 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, $V_p^{\tilde{\sigma}}(p, \tilde{p}) \geq 0$. However, $V_p^\sigma(p, \tilde{p}) = 0$ and $U'_\ell(p) < 0$, so $V_p^{\tilde{\sigma}}(p, \tilde{p}) \geq 0$ requires $\tilde{m}(\tilde{p}) = 0$, as can be seen from (44).

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 $\tilde{m}(\bar{p}) = 1$ for all $\bar{p} \in (\tilde{p}, \tilde{p} + \varepsilon]$. If $\bar{p} \leq \hat{p}$ for all \bar{p} in that interval, then (10) implies $\tilde{\rho}(\bar{p}) = 0$, in which case we have $V^{\tilde{\sigma}}(p, \bar{p}) = U_\ell(p)$ and hence $\tilde{\pi}(\bar{p}) = \bar{p}$. We can then follow the same argument as in point 2(b) to show that $\tilde{m}(\bar{p}) = 1$ violates the HJB condition (12). Suppose instead $\tilde{p} \geq \hat{p}$. Given $p(\tilde{p}) < p(\bar{p}_3) < \hat{p}$, the saddle point conditions (7) and (10) of Definition 1 imply $\tilde{\rho}(\bar{p}) = \hat{p}$ and $\tilde{\pi}(\bar{p}) = \hat{p}$ for all $\bar{p} > \tilde{p}$ sufficiently close to \tilde{p} such that $\tilde{m}(\bar{p}) = 1$, with $V^{\tilde{\sigma}}(p, \bar{p}) = \hat{u}$ for all $p \in \mathcal{P}(\bar{p})$ as the DM's value. Substituting for nature's choice and $V^{\tilde{\sigma}}$,⁴³ the HJB functional simplifies to

$$G(m, \nu, 0, \tilde{\pi}(\bar{p}), \bar{p}, V^{\tilde{\sigma}}, dV^{\tilde{\sigma}}) = (1 - m) (-c + \hat{p}\lambda(u_r^R - \hat{u})).$$

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

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

⁴²Unless $\tilde{m}(\tilde{p}) = \tilde{v}(\tilde{p}) = 0$, in which case $\tilde{\rho}(\tilde{p})$ can take an arbitrary value.

⁴³Note that, given $V^{\tilde{\sigma}}(p, \bar{p}) = \hat{u}$, we have $V_p^{\tilde{\sigma}}(p, \bar{p}) = V_{\bar{p}}^{\tilde{\sigma}}(p, \bar{p}) = 0$.

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

4. $\tilde{p} \in [\bar{p}_3, \bar{p}_4)$: Recall $m(\bar{p}) = \nu(\bar{p}) = 0$ for all $\bar{p} \in [\bar{p}_3, \bar{p}_4)$. The argument is analogous to that in Point 2. In the current case, $U_{\tilde{\rho}(\bar{p})}(p) < V^\sigma(p)$ 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 $U_{\tilde{\rho}(\bar{p})}(\underline{p}(\tilde{p})) < V^\sigma(\underline{p}(\tilde{p}), \tilde{p})$.
5. $\tilde{p} \in [\bar{p}_4, 1]$: Recall $m(\bar{p}) = 1$ and $\rho(\bar{p}) = 1$ for all $\bar{p} \in [\bar{p}_4, 1]$. For every $\bar{p} \geq \bar{p}_4$, we have $\underline{p}(\bar{p}) > \hat{p}$ and hence $U_\rho(p) < U_r(p)$ for all $\rho < 1$ and all $p \in \mathcal{P}(\bar{p})$. Condition (12) thus requires $\tilde{\rho}(\bar{p}) = 1$ for all $\bar{p} \geq \tilde{p}$. At $\bar{p} = \bar{p}_4$, $U_r(\underline{p}(\bar{p})) = V^\sigma(\underline{p}(\bar{p}), \bar{p})$, so any $\tilde{m}(\bar{p}_4) \in [0, 1]$ solves the DM's optimization problem. Leaving aside this point of indifference, we want to show that for any $\bar{p} > \bar{p}_4$, $\tilde{m}(\bar{p}) = m(\bar{p}) = 1$ must hold. Assuming this property fails, we must have an open interval $(\tilde{p}, \tilde{p} + \varepsilon)$, $\varepsilon > 0$ such that $\tilde{m}(\bar{p}) = 0$ for all $\bar{p} \in (\tilde{p}, \tilde{p} + \varepsilon)$.⁴⁴ If $\tilde{p} > \bar{p}_4$, 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 $\tilde{p} = \bar{p}_4$ and $\tilde{m}(\bar{p}) = 0$ for all $\bar{p} \in (\bar{p}_4, \bar{p}_4 + \varepsilon)$ and some $\varepsilon > 0$. If $\tilde{\nu}(\bar{p}) = 0$ on a right neighborhood of \tilde{p} , the DM's value on that neighborhood is described by

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

with Ψ and q as defined in (36) and (37). But, by definition of \bar{p}_4 in (39), $V^{\tilde{\sigma}}(p, \bar{p}) < U_r(\bar{p})$ holds for all $\bar{p} > \bar{p}_4 = \tilde{p}$ and $p \in \mathcal{P}(\bar{p})$, so the strategy violates condition (9). For interior stopping rates $\tilde{\nu}(\bar{p}) > 0$ on a right neighborhood of \tilde{p} , (12) requires $V^{\tilde{\sigma}}(\underline{p}(\bar{p}), \bar{p}) = U_r(\underline{p}(\bar{p}))$ for all \bar{p} belonging to that neighborhood, which in turn requires $\tilde{m}(\bar{p}) = 1$ 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 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_* .

⁴⁴See Footnote 40.

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 $\tilde{m}(\bar{p}) = 0$ for all \bar{p} belonging to the neighborhood. Under this assumption the value $V^{\tilde{\sigma}}(p, \bar{p})$ is right-continuous in \bar{p} at $\bar{p} = p_*$.⁴⁵ For $\bar{p} > p_*$ sufficiently close to p_* we thus have $\tilde{\pi}(\bar{p}) = \bar{p}$. Since stopping and taking action r is strictly Bayesian optimal for all $\bar{p} > p_*$, we then have $U_r(\tilde{\pi}(\bar{p})) > V^{\tilde{\sigma}}(\tilde{\pi}(\bar{p}), \bar{p})$ for all $\bar{p} > p_*$ sufficiently close to p_* , thus violating (9) of Definition 1. Suppose next $\tilde{m}(\bar{p}) = 1$ for all \bar{p} belonging to an interval $(p_*, p_* + \varepsilon)$, $\varepsilon > 0$. We distinguish two cases. If $\underline{p}(p_*) \geq \hat{p}$, then (10) implies $\tilde{\rho}(\bar{p}) = 1$ for all $\bar{p} \in (p_*, p_* + \varepsilon)$ 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 $\tilde{\rho}(\bar{p}) = \hat{\rho}$ and $\tilde{\pi}(\bar{p}) = \hat{p}$ for all $\bar{p} \in (p_*, p_* + \varepsilon)$. 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

$$V_p^{\tilde{\sigma}}(p, p_*) = \tilde{m}(p_*)U'_{\tilde{\rho}(p_*)}(p) + (1 - \tilde{m}(p_*))\Phi'(p) = 0. \quad (45)$$

Since $\Phi'(p) < 0$ for all $p \in \mathcal{P}(p_*)$, this equality requires $\tilde{\rho}(p_*) \in [\hat{\rho}, 1]$ and $\tilde{m}(p_*) \in [m(p_*), 1]$.⁴⁶ For the DM to optimally choose $\tilde{m}(p_*) \geq m(p_*)$, it must further hold $U_{\tilde{\rho}(p_*)}(\tilde{\pi}(p_*)) = \Phi(\tilde{\pi}(p_*))$.

⁴⁵As we approach p_* from the right, the probability of stopping before p_* vanishes.

⁴⁶Recall that $m(p_*)$ is determined by

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

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 (45) 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 [Appendix B.2](#)) implies that under any alternative strategy there must be a state at which the DM can profitably deviate to the commitment solution.

B.10 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 .⁴⁷ 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

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

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

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

with respect to p vanishes, so

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

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

⁴⁷That this is the correct value of $\tilde{\sigma}$ can be seen by confirming that $\tilde{u} = \frac{\tilde{\nu}U_\ell(\tilde{p}) + \tilde{p}\lambda u_r^R}{\tilde{\nu} + \tilde{p}\lambda}$.

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 [Appendix A.3.2](#), which applies here since $\Delta = \infty$ if $[\underline{p}_0, \bar{p}_0] = [0, 1]$.

To see that the equilibrium is the limit of the equilibrium strategies as $(\underline{p}_0, \bar{p}_0) \rightarrow (0, 1)$, note that for $c < \underline{c}$ and $(\underline{p}_0, \bar{p}_0)$ sufficiently close to $(0, 1)$, \bar{p}_0 falls into Region 3. Indeed, as $(\underline{p}_0, \bar{p}_0) \rightarrow (0, 1)$, \bar{p}_3 rises to 1 faster than \bar{p}_0 does. To see this, suppose $(\underline{p}_0, \bar{p}_0)$ is sufficiently close to $(0, 1)$ so that $\bar{p}_0 > p_* = \bar{p}_2$ and $\underline{p}_0 < p_\ell^B$. Recall \bar{p}_3 is given by $V^\sigma(\underline{p}(\bar{p}_3), \bar{p}) = U_\ell(\underline{p}(\bar{p}_3))$, that is, as the value of \bar{p} for which the left end of the value segment touches U_ℓ . Since $V^\sigma(\underline{p}(\bar{p}_3), \bar{p}_3)$ lies below the Bayesian value function, we clearly have $\underline{p}(\bar{p}_3) \geq p_\ell^B$. Since $\underline{p}_0 = \underline{p}(\bar{p}_0) < p_\ell^B$, \bar{p}_0 is strictly smaller than \bar{p}_3 (and of course strictly greater than p_*). Given that \bar{p}_0 falls into Region 3, it is then routine to check that $(\nu(\bar{p}_0), \pi(\bar{p}_0)) \rightarrow (\tilde{\nu}, \tilde{\pi})$ as $\bar{p}_0 \rightarrow 1$. Similarly, when $c \geq \underline{c}$, the same conclusion follows since the condition for [Appendix A.3.2](#) is satisfied for \bar{p}_0 sufficiently close to 1.

B.11 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 $\lim_{p \uparrow 1} \Phi'(\bar{p}) = +\infty$. For all states $\bar{p} \in (p_\ell^B, p_*]$ experimentation until the Bayesian update of \bar{p} reaches the Bayesian stopping boundary p_ℓ^B (or R -evidence arrives) is clearly optimal: by the assumption $c < \bar{c}$, the worst-case payoff associated to this strategy, $\Phi(\bar{p})$, is strictly greater than the worst-case payoff for action ℓ , given by $U_\ell(\underline{p}(\bar{p}))$ and strictly greater than the worst-case payoff for action r , given by $U_r(\bar{p})$.

Fixing this part of the strategy and moving backwards in time from $\bar{p} = p_*$, 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 \bar{p} , given by

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

Convexity of Φ implies that the value $l(0, \bar{p})$ is strictly decreasing in \bar{p} with $\lim_{\bar{p} \uparrow 1} l(0; \bar{p}) =$

$-\infty < U_\ell(p)$ for all p . Given $l(0, p_*) = \Phi(p_*) > U_\ell(p_*)$, the difference $U_\ell(\bar{p}) - l(0; \bar{p})$ as a function of \bar{p} must then have an intersection with zero. Notice next that for any given value of \bar{p} , we can find a Δ sufficiently large such that $\underline{p}(\bar{p})$ is arbitrarily close to zero. Together with the previous property, this implies that for Δ sufficiently large, there is a state $\bar{p} > p_*$ such that $l(\underline{p}(\bar{p}); \bar{p}) = U_\ell(\bar{p})$. Let \bar{p}_s denote the smallest of such states. For all $\bar{p} \in (\pi_\ell^B, \bar{p}_s)$, the DM will then experiment, whereas at $\bar{p} = \bar{p}_s$ she stops to take action ℓ .⁴⁸

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 $\bar{p} \in (\bar{p}_s, \bar{p}_s + \varepsilon)$. The payoff from experimentation at $\bar{p} > \bar{p}_s$ evaluated at belief p is⁴⁹

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

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

$$\bar{p}(1 - e^{-\lambda\tau(\bar{p}, \bar{p}_s)})\delta - c\tau(\bar{p}, \bar{p}_s) > 0, \tag{46}$$

where

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

It can be easily verified that the left-hand side of (46) approaches zero from above as $\bar{p} \downarrow \bar{p}_s$ if and only if $\bar{p}_s > p_\ell^B$, which is satisfied since $\bar{p}_s > p_* > p_\ell^B$.

⁴⁸For $\underline{p}(\bar{p}_s)$ sufficiently close to zero, we must have $U_\ell(\bar{p}_s) \geq U_r(\underline{p}(\bar{p}_s))$, so action ℓ is indeed optimal.

⁴⁹Recall that $\tau(\bar{p}, \bar{p}_s)$ denotes the time it takes for belief $\bar{p} > \bar{p}_s$ to be updated to \bar{p}_s in the absence of a breakthrough.

B.12 Proof of Proposition 3

It is useful to characterize the learning time through the following recursion equation.⁵⁰

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

where $\nu_\Delta(\theta)$ is the stopping rate for ambiguity level Δ when the state \bar{p} is such that $\ln\left(\frac{\bar{p}}{1-\bar{p}}\right) = \ln\left(\frac{\theta}{1-\theta}\right) + \frac{\Delta}{2}$. 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.13 Proof of Proposition 4

We indicate by \bar{p}_n and \bar{q}_n for $n = 1, \dots, 4$ 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 $\text{Supp}(F_\mathcal{P})$, 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 $\mathcal{P} = [\underline{p}_0, \bar{p}_0]$, which requires $\bar{p}_0 \in (p_*, \bar{p}_4)$ as well as $\Phi(p_*) < U_\ell(\hat{p}(\tau(\bar{p}_0, p_*), \underline{p}_0))$.

⁵⁰This 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 $dt \rightarrow 0$ while using $\dot{\theta} = -\lambda\theta(1-\theta)$, we obtain the equation.

We show first $\tau(\bar{p}_0, \bar{p}_3) \geq \tau(\bar{q}_0, \bar{q}_3)$. Towards a contradiction, suppose $\tau(\bar{p}_0, \bar{p}_3) < \tau(\bar{q}_0, \bar{q}_3)$, so the intrapersonal equilibrium under \mathcal{P} prescribes randomization at time $\tau(\bar{q}_0, \bar{q}_3)$. By [Lemma A.3](#), $\hat{V}(\bar{p})$ is strictly decreasing in \bar{p} in Region 3. Recalling that p^δ denotes the δ update of p ,⁵¹ we have

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

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

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 $\hat{\nu}(t; \mathcal{P}) := \nu(\bar{p}^t; \mathcal{P})$ the DM's stopping rate as a function of time. We want to show for all $t \leq \bar{t}_\mathcal{P}$,

$$\hat{\nu}(t; \mathcal{P}) \leq \hat{\nu}(t; \mathcal{Q}). \tag{47}$$

When $t < \tau(\bar{p}_0, \bar{p}_3)$ or $t \in (\tau(\bar{p}_0, \bar{p}_*), \bar{t}_\mathcal{P}]$, $\hat{\nu}(t; \mathcal{P}) = 0$, so (47) trivially holds. Considering $t \in [\tau(\bar{p}_0, \bar{p}_3), \tau(\bar{p}_0, p_*)]$, we have

$$\hat{\nu}'(t; \mathcal{Q}) = \nu'(\hat{p}(t, \mathcal{Q})) \frac{d\hat{p}(t, \bar{q}_0)}{dt} < 0, \tag{48}$$

since for all \bar{p} belonging to Region 3, $\nu'(\bar{p}) = \lambda \hat{V}'(\bar{p}) / \delta^\ell < 0$ (see (16)). (48) directly implies

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

and thus (47). Since (47) 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.

⁵¹That 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 $c < \bar{c} = \frac{\lambda\delta}{2}$ so that experimentation is optimal for some beliefs.⁵² The optimal strategy, characterized in [Che and Mierendorff \(2019\)](#), is succinctly described in [Figure 12](#) for two different ranges of learning costs. Suppose first the cost is *intermediate*; i.e., $c \in [\underline{c}^*, \bar{c})$, 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 12](#).

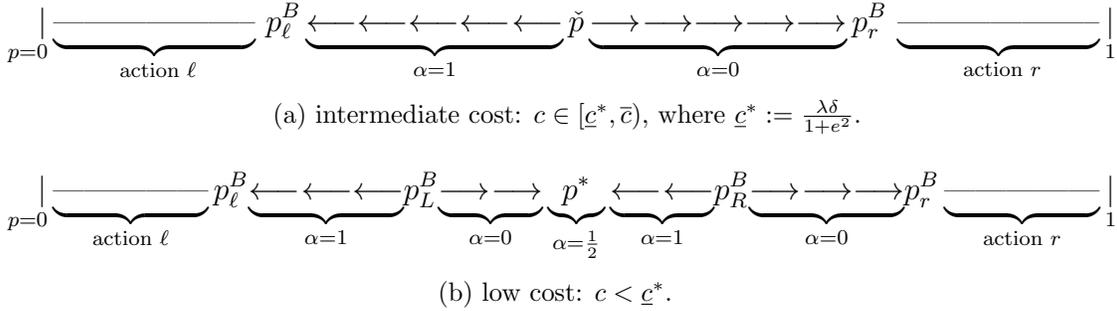


Figure 12: Bayesian optimal strategy.

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 12](#). When such evidence does not arrive, the DM’s belief

⁵²Otherwise, 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 information 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 $\underline{p}(\bar{p})$ denotes the lowest possible belief when the state is \bar{p} . Below, we suppress the dependence of \underline{p} on \bar{p} when no confusion arises. The equilibrium for this case is characterized as follows.

Theorem 4. *Suppose the cost is low or the ambiguity is small. The intrapersonal equilibrium of the ambiguity-averse DM is as follows.*

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

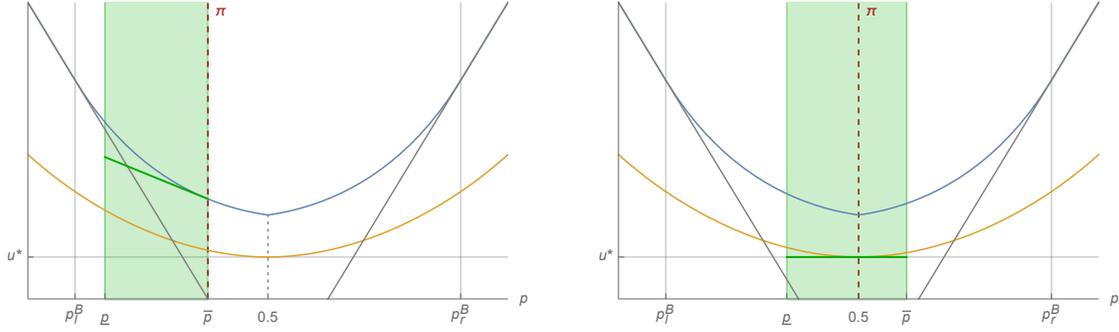


Figure 13: The case of intermediate cost and small ambiguity.

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

Proof. The proof is completely analogous to the baseline case for either $\bar{p} < p_* = 1/2$ or $\underline{p} > p_*$. Specifically, consider the former case. With $\pi(\bar{p}) = \bar{p}$, the DM's HJB coincides with that of the Bayesian DM for belief \bar{p} , so the result follows from [Che and Mierendorff \(2019\)](#). For nature's optimality, since the Bayesian value function is decreasing, clearly $\pi(\bar{p}) = \bar{p}$ 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 $\underline{p} < 1/2 < \bar{p}$. Nature chooses $\pi(\bar{p}) = 1/2$. 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 $V(p, \bar{p}) = \hat{u} = \frac{1}{2}\delta$ for all p , so nature does not wish to deviate from $\pi(\bar{p}) = 1/2$. 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 $V(p, \bar{p}) = u^*$ for all p , nature does not wish to deviate from $\pi(\bar{p}) = 1/2$. \square

To explain, suppose first the highest belief \bar{p} is less than $1/2$. Then, as before, the worst-case belief is the right-most belief \bar{p} , and the DM adopts the Bayesian optimal strategy for that belief. See the left panel of [Figure 13](#). In case $c \in [\underline{c}^*, \bar{c})$, the worst-case belief \bar{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 \bar{p} , actually drifts rightward when $\bar{p} > p_L^B$. 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 $\underline{p} < 1/2 < \bar{p}$. 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 \bar{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 13](#), the value segment touches the Bayesian 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 $c \in [\underline{c}^*, \bar{c})$, 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 12](#)-(a)). This explains the discontinuity in the right panel of [Figure 13](#).

Intermediate cost and large ambiguity. Suppose next that the cost is *intermediate*, namely $c \in [\underline{c}^*, \bar{c})$, and the ambiguity is *large* in the sense that $U_\ell(\underline{p}(\frac{1}{2})) > \Phi^*(\frac{1}{2})$. Recall the value $\hat{V}(\bar{p})$ defined in [Theorem 1](#) for the randomized stopping region (Region 3). This value solves [\(ODE\)](#) with the boundary condition $(\bar{p}_2, \Phi^*(\bar{p}_2)) = (\frac{1}{2}, \Phi^*(\frac{1}{2}))$ and is explicitly expressed in [\(31\)](#) (with $\bar{p}_2 = \frac{1}{2}$). 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 5. *Suppose the cost is intermediate, and the ambiguity is large. Then, the following is an equilibrium.*

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

- For any $\underline{p} \geq 1/2$, the DM employs the Bayesian optimal strategy for belief $\pi(\bar{p}) = \bar{p}$.
- Suppose $[\underline{p}, \bar{p}] \supset [p_-, p_+]$. Then, the DM hedges either by employing the split-attention learning or by mixing between r and ℓ , with the worst-case belief $\pi(\bar{p}) = 1/2$.
- Suppose $\underline{p} < 1/2$ and $\bar{p} \in (1/2, p_+]$. 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 [Appendix A.3](#).
- Suppose $\underline{p} \in [p_-, 1/2)$ and $\bar{p} > 1/2$. 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 [Appendix A.3](#).⁵³

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 $\bar{p} \in (1/2, p_+]$, we simply focus on that case. This case in turn corresponds precisely to Region 3 of [Theorem 1](#). The construction of $(\nu(\bar{p}), \pi(\bar{p}), \hat{V}(\bar{p}))$ follows exactly without any modification. The verification of HJB is also the same except that $\alpha(\bar{p}) = 1$ 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 :

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

It suffices to show that this derivative is nonnegative when evaluated at $(\nu(\bar{p}), \pi(\bar{p}), \hat{V}(\bar{p}))$. Indeed, when substituting $(\nu(\bar{p}), \pi(\bar{p}), \hat{V}(\bar{p}))$ from [\(16\)](#), [\(15\)](#), and [\(ODE\)](#), this expression is non-negative if and only if

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

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

If $\bar{p} \leq 1/2$ 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.

⁵³Given 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.

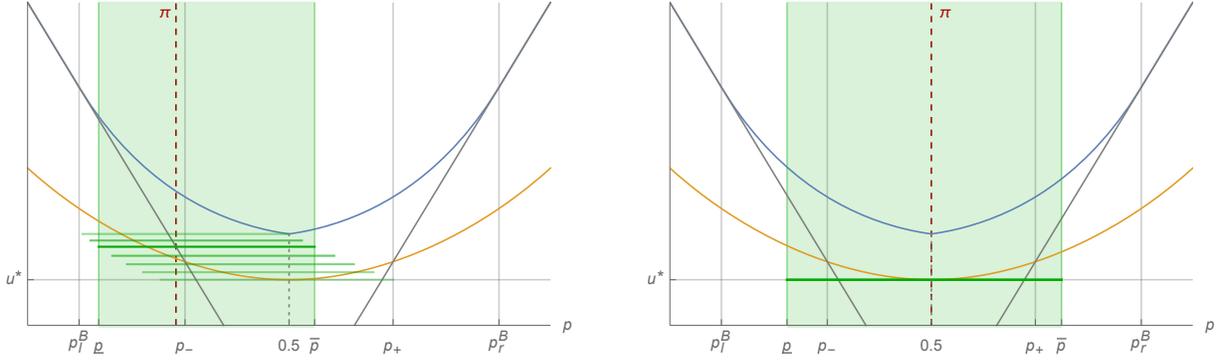


Figure 14: The case of intermediate c and large Δ .

If $\underline{p} < 1/2 < \bar{p}$, the DM employs a hedging strategy. If $[\underline{p}, \bar{p}] \supset [p_-, p_+]$ 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 14. If instead $\bar{p} \in (1/2, p_+]$, 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 14. 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.