

Optimal Sequential Decision with Limited Attention*

Yeon-Koo Che Konrad Mierendorff

This draft: July 2, 2017

First draft: April 18, 2016

Abstract

We consider a dynamic model of information acquisition. Before taking an action, a decision maker may direct her limited attention to collecting different types of evidence that support alternative actions. The optimal policy combines three strategies: (i) *immediate action*, (ii) a *contradictory strategy* seeking to challenge the prior belief, and (iii) a *confirmatory strategy* seeking to confirm the prior. The model produces a rich dynamic stochastic choice pattern as well as predictions in settings such as jury deliberation and political media choice.

KEYWORDS: Wald sequential decision problem, choice of information, contradictory and confirmatory learning strategies, limited attention.

1 Introduction

In many situations, decision makers (DMs) must choose actions whose payoffs are initially unknown. For instance, a firm or government agency may not know the merit of a potential investment project, or a candidate for a position. A similar uncertainty may be present when a prosecutor decides whether to indict a suspect, when a judge or a jury deliberates on the guilt or innocence of the accused, when voters decide which candidate to vote for, or when a researcher seeks to ascertain a hypothesis. In these situations, before making a decision, the DM can acquire information or spend time processing information already available to her. Acquiring more information allows her to make better decisions

*We thank Martino Bardi, Guy Barles, Hector Chade, Martin Cripps, Mark Dean, Mira Frick, Johannes Hörner, Han Huynh, Philippe Jehiel, Ian Krajbich, Botond Kőszegi, Suehyun Kwon, George Mailath, Filip Matejka, Pietro Ortoleva, Andrea Prat, Ed Schlee, Philipp Strack, Tomasz Strzalecki, Glen Weyl, Michael Woodford, Xingye Wu, Weijie Zhong, and audiences at numerous conferences and seminars for helpful comments. We also thank Timur Abbyasov for excellent research assistance.

Che: Department of Economics, Columbia University, 420 W. 118th Street, 1029 IAB, New York, NY 10025, USA; email: yeonkooche@gmail.com; web: <http://www.columbia.edu/~yc2271/>.

Mierendorff: Department of Economics, University College London, 30 Gordon Street, London WC1H 0AX, UK; email: k.mierendorff@ucl.ac.uk; web: <http://www.homepages.ucl.ac.uk/~uctpkmi/>.

but can be costly and may delay the action. The DM must therefore decide how much information she gathers, and typically this is not a one time decision. During the process of information acquisition, she can decide sequentially to continue acquiring information or to stop and take an action, depending on what she has learned so far. Seminal papers by Wald (1947) and Arrow, Blackwell, and Girshick (1949) have analyzed this stopping problem.

In many problems, a critical aspect of sequential decision making is not just *how much* information the DM acquires, but often more importantly *what kind* of information she seeks. For instance, a firm may seek evidence in favor of or against investing in a given project. A judge or jury may scrutinize incriminating or exculpatory evidence. A researcher may set out to prove a hypothesis by way of a mathematical proof or empirical evidence, or attempt to disprove it via a counter-example or contradictory evidence. How decision makers direct their resources and attention to alternative types of evidence influences the information they receive and ultimately the quality of decisions they make.

The present paper focuses on this aspect of sequential decision problems. We consider a model with binary actions, a and b , which are optimal in states A and B , respectively. The state is initially unknown, and the DM has a prior belief $p_0 \in (0, 1)$ about the probability that the state is A . In continuous time, the DM may decide to acquire information, or to stop and take a final and irreversible action. Information acquisition has a direct flow cost, and/or payoffs are discounted exponentially.

The DM may seek different types of evidence that would reveal alternative states in varying degrees of accuracy.¹ In the baseline model evidence is fully conclusive. The DM may seek A -evidence, and if the state is A , she will find such evidence at Poisson rate λ . If the state is B , she will not discover any A -evidence. Alternatively, she may seek B -evidence, which arrives at Poisson rate λ if the state is B , and she will not discover B -evidence if the state is A . Since evidence never arrives for the wrong state, it is conclusive. It is thus optimal to stop learning and take the optimal action immediately after discovery.

To capture the idea that decision makers often have limited resources (cognitive, financial, time, equipment, manpower) that can be devoted to acquiring or processing information, we assume that the DM has a unit budget of “attention,” which she must allocate between seeking A -evidence and B -evidence. Allocating less attention to one type of evidence will proportionally reduce the corresponding arrival rate. Being Bayesian, the DM also updates her belief in the absence of discovery. For example, seeking A -evidence but failing to find it makes her more pessimistic that the state is A .

Given this model, we characterize the optimal strategy, and provide comparative statics with respect to the information cost, discount rate, and payoff parameters. Our model yields rich predictions which we explore in two real-world settings: deliberation by grand

¹Seeking evidence could also mean introspective deliberation, i.e., scrutinizing and revisiting information that is already available.

and trial juries, and media choice by voters. Going beyond the baseline model, we analyze a model with non-conclusive evidence. We extend our main characterization to this general model and explore the implications for stochastic choice and response time.

Before discussing these results in greater detail, we describe the main characterization and provide an intuition. While our model allows for general strategies, we show that for each prior belief, the DM optimally uses one of three simple heuristics: (i) *immediate action*; (ii) *contradictory learning*; and (iii) *confirmatory learning*; and never switches between these different modes of learning.

Immediate action is the simple strategy where the DM takes an optimal action without acquiring any information. **Contradictory learning** involves seeking evidence that would confirm the state the DM finds relatively *unlikely*. An example is to seek *B*-evidence when state *A* is relatively likely. Seeking *B*-evidence but not discovering any makes the DM even more certain that the state is *A*. Ultimately, if no *B*-evidence arrives, she becomes so certain about the state being *A* that she chooses action *a* without further learning. To the extent that this is a likely event, a DM is effectively trying to “disprove” or “rule out” the unlikely state by playing the “devil’s advocate” to her own mind.

Confirmatory learning involves seeking evidence that would confirm the state the DM finds relatively *more likely*. An example is to seek *A*-evidence when state *A* is relatively *likely*. When following this strategy, the DM becomes less confident about the likely state when no discovery occurs. Eventually, she becomes so uncertain that she switches to a second phase where she divides her attention seeking both *A*-evidence and *B*-evidence until she observes a signal that reveals the true state.

We characterize which strategy is optimal for each prior belief, and show how the structure of the optimal policy depends on the cost of information. Not surprisingly, if information is very costly, the DM takes an immediate action for all beliefs. For moderate information acquisition costs, we show that the DM optimally takes an immediate action when she is extremely certain, while she employs contradictory learning when she is more uncertain. Finally, if information acquisition costs are low, immediate action is again optimal for extreme priors, and contradictory learning is optimal for less extreme priors. For very uncertain priors, however, confirmatory learning becomes optimal.

The intuition behind the optimal policy is explained by a trade-off between *accuracy* and *delay*. With an extreme belief, a fairly accurate decision can be made, and further information acquisition has a smaller benefit than the cost of delaying the action. Conversely, with a less extreme belief, information acquisition is more valuable. This explains why the *experimentation region* contains moderate beliefs and the *stopping region* is located at the extreme ends of the belief space. The trade-off also explains which strategy is optimal inside the experimentation region. Confirmatory learning will lead to a fully accurate decision because the DM never takes an action before learning the state, but this could take a long time. By contrast, contradictory learning seeks evidence only for

a limited duration. The DM is more likely to make a quick decision, but one that is not as accurate when the DM mistakenly “rules out” the unlikely state. Again, the more certain the DM is, the less valuable the evidence is. This explains why the DM chooses contradictory learning when she is more certain and confirmatory learning when she is more uncertain. An implication is that a “*skeptic*” is more likely to make an accurate decision with a longer delay than a “*believer*.”

Our model yields rich implications which we explore in two settings. First, it predicts distinct ways in which grand juries and trial juries deliberate: Trial juries adopt a high evidentiary standard for conviction and focus attention on incriminating evidence, whereas grand juries adopt a lower evidentiary standard for an indictment and focus attention on exculpatory evidence. We also extend the model to analyze how the possibility of a “hung jury” affects jury decisions.

Second, we derive implications of our model for the choice of news media by a voter. We show that optimal media choice leads to an “echo-chamber effect,” where voters subscribe to media that are likely to push them in the direction of their prior belief. Interestingly, with sufficiently informative media, this effect is reversed for voters with moderate beliefs. They optimally seek opposite-biased outlets, creating an “anti-echo chamber effect”. We extend the model to allow for a trade-off between bias and informativeness in media choice and show that voters with more extreme beliefs value informativeness less than moderates.

Finally, we formulate a generalized model that allows for *non-conclusive* evidence. We focus on the case where the optimal strategy of the DM has the *single experimentation property* (SEP)—that is, she finds it optimal to take an action as soon as she observes either *A*- or *B*-evidence. We show that this is the case if the cost of information is sufficiently high. If SEP is satisfied, the optimal policy is composed of immediate action, contradictory learning, and confirmatory learning strategies as in the baseline model.² While the structure of the optimal policy is preserved, the possibility of non-conclusive evidence leads to richer implications due to the imperfect learning and permits an interesting comparison with other models of stochastic choice.

First, we show that a DM with a more uncertain prior—a “skeptic”—ends up making more accurate decisions but with a longer delay, compared with a DM with a more extreme prior—a “believer.” Generally, the stochastic choice function in our model depends on the prior belief, which is not the case in classic random utility models of stochastic choice such as logit. This prior dependence is a common feature of models of information acquisition such as the rational inattention (RI) model, or drift diffusion models (DDM), which we discuss below. Unlike our model, however, these models predict that the *accuracy* of the decision following information acquisition is *independent* of the prior belief.³

²In Section 6 we go beyond SEP and construct examples where the DM optimally engages in repeated experimentation. These examples show that the central conclusion of the baseline model are robust to the introduction of non-conclusive evidence.

³Fudenberg, Strack, and Strzalecki (2017) show that extending the DDM to a rich state-space leads

Second, our model generates an interesting prediction on how the accuracy of the decision relates to *the response time*—that is, the amount of time it takes to reach that decision. If the DM obtains a signal, say in favor of A , after a long duration of trying, then it is relatively weak evidence in favor of A , compared to obtaining the same signal only after a short duration of trying. This means that for a given prior belief, a longer deliberation is associated with a less accurate decision, thus capturing a sense of *speed-accuracy complementarity* documented in cognitive psychology experiments.⁴

Related Literature. Wald’s sequential decision problem has been formulated in continuous time as a drift-diffusion model (DDM) as well as with Poisson signals (see [Peskir and Shiryayev, 2006](#), ch. VI, for rigorous treatments and references). In DDMs, the decision maker observes a Brownian motion whose drift is correlated with the true state. Our Poisson model assumes a learning technology in which the DM can seek “lumpy” evidence, which is applicable in situations where signals arrive rarely but reveal a lot of information. In contrast, DDMs are more suitable for situations where small amounts of information are revealed frequently.⁵

The Poisson signal structure builds on the exponential bandit model of [Keller, Rady, and Cripps \(2005\)](#). Closest to our baseline model is the planner’s problem for negatively-correlated bandits in [Klein and Rady \(2011\)](#) and [Francetich \(2014\)](#). However, the two models are fundamentally different. In our model, exploiting a payoff requires the DM to take an irreversible action and stop learning. By contrast, in bandit problems exploiting the payoff of an arm always generates information. For this reason, a distinct characterization emerges; for instance, there is no analogue of our “contradictory strategy.”⁶

A recent incarnation of Wald’s problem is [Fudenberg, Strack, and Strzalecki \(2017\)](#), who introduce rich states in DDM with two actions and obtain speed-accuracy complementarity. [Moscarini and Smith \(2001\)](#) endogenize the signal precision in a DDM through costly effort decisions.⁷ By contrast we focus on endogenous *types* of information. Related to this [Nikandrova and Pansc \(2017\)](#) have considered the problem of selectively learning about different investment projects. They use Poisson signals as in our baseline model, but the payoffs of final actions are uncorrelated whereas our model assumes negatively

to prior dependence in the accuracy.

⁴Subjects of perceptual or choice experiments often exhibit speed-accuracy complementarity (see [Ratcliff and McKoon \(2008\)](#) for a survey). While our model produces this prediction, the extent to which our model, particularly the Poisson signal, describes adequately the neuro-physiology underpinning subjects’ behavior in these experiments is unclear (see also footnote 46 below).

⁵While different applications call for different assumptions on the signal structure, an interesting question is which type of signal a DM prefers if she can acquire information *flexibly*, and incurs a cost that depends on the informativeness of the signal structure. [Zhong \(2017\)](#) shows that for a large class of (flow) cost functions, in the continuous time limit, it is optimal to obtain Poisson signals.

⁶This is also the case in [Damiano, Li, and Suen \(2017\)](#) who add learning to a Poisson bandit model.

⁷[Chan, Lizzeri, Suen, and Yariv \(2016\)](#) analyze a stopping decision by a heterogeneous committee. [Henry and Ottaviani \(2017\)](#) turn Wald’s problem into a persuasion game by splitting authority over final decisions and stopping between two players.

correlated payoffs.⁸ In addition to the different payoff structure and applications, these papers do not consider a general model with noisy signals or explore the stochastic choice implications.⁹

Our model shares a common theme with the rational inattention model introduced by Sims (2003) and further developed by Matejka and McKay (2015). Like our paper, they explain individual choice as resulting from one’s optimal allocation of limited attention over diverse information. While RI abstracts from the dynamic process of deliberation, we attempt to unpack the “black box” and explicitly model a dynamic learning process that gives rise to the predicted choice outcome.¹⁰

The rest of the paper is organized as follows. Section 2 presents the baseline model and Section 3 characterizes the optimal policy. Section 4 applies the model to jury deliberation and media choice. Section 5 generalizes the model under SEP and explores implications on stochastic choice and response time. Section 6 treats the case of repeated experimentation. Section 7 concludes. Omitted proofs can be found in Appendix A and in the Supplemental Material (Che and Mierendorff, 2017).

2 Baseline Model

We consider a DM who must take an action with unknown payoff. In the exposition, we will refer to three canonical examples: (i) a voter subscribing to news media; (ii) a jury deliberating on a verdict; and (iii) a scientist performing experiments on a hypothesis.

States, Actions and Payoffs. The DM must choose from two *actions*, a or b , whose payoffs depend on the unknown *state* $\omega \in \{A, B\}$. The *payoff* of taking action x in state ω is denoted by $u_x^\omega \in \mathbb{R}$. We label states and actions such that it is optimal to “match the state,” and assume that the optimal action yields a positive payoff—that is, $u_a^A > \max\{0, u_b^A\}$ and $u_b^B > \max\{0, u_a^B\}$.¹¹ The DM may delay her action and acquire information. In this case, she incurs a flow cost of $c \geq 0$ per unit of time. In addition, her payoffs (and the flow cost) are discounted exponentially at rate $r \geq 0$.¹²

⁸See also the recent papers by Mayskaya (2016) and Ke and Villas-Boas (2016) which are concurrent to our paper.

⁹In a binary action model, Fudenberg, Strack, and Strzalecki (2017) also consider the allocation of (limited) attention between two Brownian motions, each indicating the (unknown) payoff of one action, and show that it is optimal to devote equal attention to each process. Woodford (2016) considers a DM optimizing on the types of evidence she seeks subject to entropy-based capacity constraint but in an exogenously specified DDM framework with exogenous boundaries.

¹⁰The dynamic RI model of Steiner, Stewart, and Matejka (2015) considers a sequence of multiple actions. Applied to a Wald setup with a single irreversible action, the predictions of the model are similar to the static case.

¹¹Note that we allow for $u_x^A = u_x^B$ so that one action $x \in \{a, b\}$ can be a *safe action*. We rule out the trivial case in which $u_x^\omega \geq u_y^\omega$ for $x \neq y$, in both states $\omega = A, B$.

¹²One of c and r may be zero but not both.

In the voter example, the actions a and b may correspond to voting for “right-leaning” and “left-leaning” candidates, respectively, and the state captures which candidate has a better platform. In the jury example, the actions correspond to “convict” (a) or “acquit” (b), and the state corresponds to the defendant’s guilt (A) or innocence (B).

The DM’s *belief* is denoted by the probability $p \in [0, 1]$ that the state is A . Her *prior belief* at time $t = 0$ is denoted by p_0 . If the DM chooses her action optimally without information acquisition, then given belief p , she will realize an expected payoff of $U(p) := \max\{U_a(p), U_b(p)\}$, where $U_x(p) := pu_x^A + (1 - p)u_x^B$ is the expected payoff of taking action x . $U(p)$ takes a piece-wise linear form as depicted in Figure 2 on page 14, and we denote the belief where $U(p)$ has a kink by \hat{p} . We denote the optimal action by $x^*(p) \in \arg \max_{x \in \{a, b\}} U_x(p)$, which is unique if $p \neq \hat{p}$.

Information Acquisition and Attention. We model information acquisition in continuous time. At each point in time, the DM may allocate one unit of *attention* to seeking one of two types of evidence: *A-evidence* which reveals state A conclusively, and *B-evidence* which reveals state B conclusively. If the DM allocates a fraction $\alpha \in [0, 1]$ of her attention to seeking *A-evidence*, and the remaining fraction $\beta = 1 - \alpha$ to seeking *B-evidence*, then she receives *A-evidence* at the Poisson arrival rate of $\alpha\lambda$ in state A , and *B-evidence* at the Poisson arrival rate of $\beta\lambda$ in state B . We denote the DM’s attention strategy by $(\alpha_t) = (\alpha_t)_{t \geq 0}$, and assume that α_t is a measurable function of t .¹³

A more concrete interpretation of the signal structure can be given depending on the context. In the voter example, α corresponds to a particular news medium (e.g., MSNBC or Fox), which publishes evidence in favor of either candidate or platform. In the absence of evidence each medium publishes partisan rhetoric corresponding to its bias. In jury deliberation, α corresponds to the attention the jury devotes to finding incriminating evidence (as opposed to exculpatory evidence). In the scientist example, α corresponds to the nature of the experiment the scientist designs—that is, proving or disproving a hypothesis.

Bayesian Updating. Suppose the DM uses the attention strategy (α_t) . Given her belief p_t , she observes signals confirming state A with Poisson rate $\alpha_t\lambda p_t$, and signals confirming state B with rate $\beta_t\lambda(1 - p_t)$. As long as she does not observe any signal, Bayes rule yields¹⁴

$$\dot{p}_t = -\lambda(\alpha_t - \beta_t)p_t(1 - p_t) = -\lambda(2\alpha_t - 1)p_t(1 - p_t). \quad (2.1)$$

¹³Given the linearity of the arrival rates, the DM cannot benefit from randomization. For this reason, we only consider deterministic strategies $(\alpha_t)_{t \in \mathbb{R}_+}$.

¹⁴By the martingale property we have $\lambda\alpha_t p_t dt + (1 - \lambda\alpha_t p_t dt - \lambda\beta_t(1 - p_t) dt)[p_t + \dot{p}_t dt] = p_t$. Dividing by dt and letting $dt \rightarrow 0$ yields (2.1).

For example, if the DM devotes all her attention to seeking A -evidence ($\alpha = 1$), she may not achieve discovery because the state is B , or because the state is A but no evidence has arrived. The longer she waits for a signal, the more convinced she becomes that the state is B . Finally, note that $\dot{p}_t = 0$ if $\alpha_t = 1/2$. That is, if the DM divides her attention equally between two types of evidence, she never updates her belief in case of no discovery.

The Decision Maker's Problem. The DM chooses an attention strategy (α_t) and a stopping time $T \geq 0$ at which a decision will be taken if no discovery has been made by then. Her problem is thus given by

$$V^*(p_0) = \max_{(\alpha_\tau), T} \left\{ \int_0^T e^{-rt} P_t(p_0, (\alpha_\tau)) [p_t \lambda \alpha_t u_a^A + (1 - p_t) \lambda \beta_t u_b^B - c] dt \right. \\ \left. + e^{-rT} P_T(p_0, (\alpha_\tau)) U(p_T) \right\}, \quad (2.2)$$

where $\beta_t = 1 - \alpha_t$, $P_t(p_0, (\alpha_\tau)) = p_0 e^{-\lambda \int_0^t \alpha_s ds} + (1 - p_0) e^{-\lambda \int_0^t \beta_s ds}$ is the probability that no signal is received by time t given strategy (α_τ) , and p_t satisfies (2.1). The first line of the objective function captures the flow cost which is incurred until the DM stops, and the payoffs from taking an action following discovery of evidence.¹⁵ The second line accounts for the payoff from the optimal decision in case of no discovery by T .¹⁶

The Hamilton-Jacobi-Bellman (HJB) equation for this problem is

$$c + rV(p) = \max_{\alpha \in [0,1]} \left\{ \lambda \alpha p (u_a^A - V(p)) + \lambda (1 - \alpha) (1 - p) (u_b^B - V(p)) \right. \\ \left. - \lambda (2\alpha - 1) p (1 - p) V'(p) \right\}, \quad (2.3)$$

for p such that $V(p) > U(p)$. If the LHS of (2.3) is larger than the RHS we must have $V(p) = U(p)$ —that is, $T(p) = 0$ and immediate action is optimal. Since the problem is autonomous, the optimal policy and the stopping decision at time t depend only on the current belief p_t . Note also that the objective in (2.3) is linear in α , which implies that the optimal policy is a bang-bang solution: the optimal attention strategy must satisfy $\alpha^*(p) \in \{0, 1\}$, except when the derivative of the objective vanishes.

3 Optimal Strategy in the Baseline Model

We begin with a description of several intuitive learning heuristics that the DM could employ. These learning heuristics form basic building blocks for the DM's optimal strategy. The details of the formal construction are presented in Appendix A.

¹⁵Specifically, at each time t , conditional on no discovery so far, A -evidence is discovered at the instantaneous rate of $p_t \lambda \alpha_t$, B -evidence is discovered at the instantaneous rate of $(1 - p_t) \lambda \beta_t$, and the flow cost c is always incurred.

¹⁶For a given $(\alpha_t)_{t \in \mathbb{R}_+}$, conditional on no discovery, posterior belief evolves according to a deterministic rule (2.1). Since stopping matters only when there is no discovery, it is without loss to focus on a *deterministic* stopping time T .

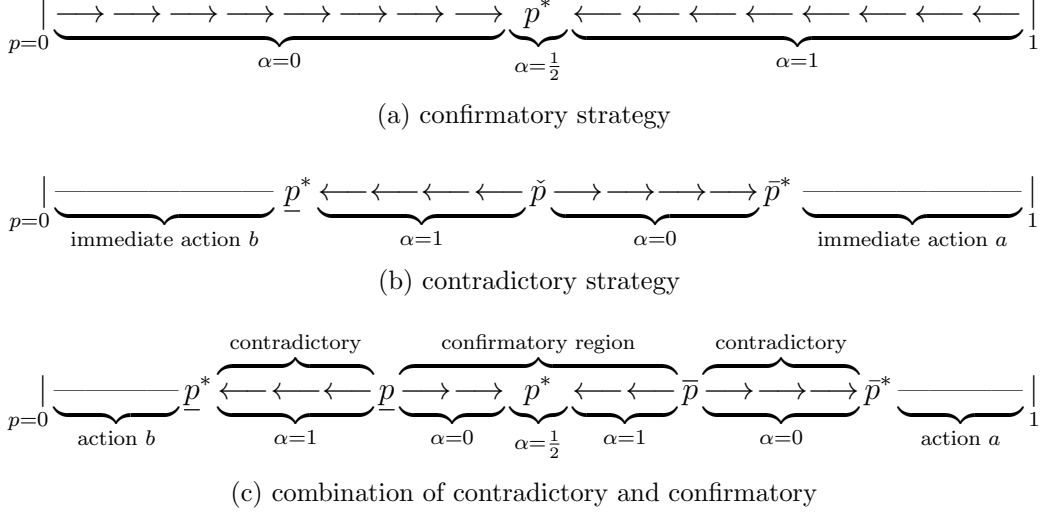


Figure 1: Structure of Heuristic Strategies and Optimal Solution

3.1 Learning Heuristics

Immediate action (without learning). A simple strategy the DM can choose is to take an immediate action and realize $U(p)$ without any information acquisition. Since information acquisition is costly, this can be optimal if the DM is sufficiently confident in her belief—that is, if p is either sufficiently high or sufficiently low.

Confirmatory learning. When the DM decides to experiment, one natural strategy is to seek evidence that “confirms” her current belief. Formally, a confirmatory learning strategy prescribes:

$$\alpha(p) = \begin{cases} 0 & \text{if } p < p^* \\ \frac{1}{2} & \text{if } p = p^* \\ 1 & \text{if } p > p^*, \end{cases}$$

for some the reference belief $p^* \in (0, 1)$, which will be chosen optimally. For beliefs above p^* , the DM seeks A -evidence. Receiving A -evidence confirms her belief that A is relatively likely and moves it to $p = 1$. For beliefs below p^* , she seeks B -evidence, which moves the belief to $p = 0$. In the absence of discovery, the DM becomes more uncertain. As is clear from equation (2.1), her belief drifts from both extremes towards the absorbing point p^* . This is illustrated in Panel (a) of Figure 1, with arrows indicating the direction of Bayesian updating. Once the reference belief p^* is reached, the DM divides her attention equally between seeking both types of evidence. In this case, no further updating occurs, and she repeats the same strategy until she obtains evidence that confirms the true state of the world.

One clear advantage of the confirmatory strategy is that it will never lead to a mistake in the action chosen by the DM. On the other hand, since the DM always waits for evidence before taking an action, the confirmatory strategy involves a potentially long delay, which

is costly.

Contradictory learning. Alternatively, the DM could try to obtain evidence that “contradicts” her current belief. Formally, contradictory learning prescribes

$$\alpha(p) = \begin{cases} 1 & \text{if } p \in (\underline{p}^*, \check{p}), \\ 0 & \text{if } p \in [\check{p}, \bar{p}^*), \end{cases} \quad (3.1)$$

for some reference belief \check{p} and boundaries of the *experimentation region* $(\underline{p}^*, \bar{p}^*)$, which will each be chosen optimally. Within the experimentation region, the DM seeks B -evidence if state A is relatively likely, and A -evidence if state B is relatively likely. In the absence of a signal, the DM’s belief drifts outward in the direction confirming her original belief. In effect, the DM is gradually ruling out the unlikely, if she seeks but fails to obtain contradictory evidence. Eventually, the DM’s belief will reach one of the boundary points \underline{p}^* or \bar{p}^* , at which she is sufficiently certain to take an immediate action without conclusive evidence. The belief updating is illustrated in Panel (b) of Figure 1.

The contradictory strategy may seem counter-intuitive, since the DM is looking for evidence for the unlikely state, which is less likely to arrive. Its value is that it may generate a “surprise” which changes the DM’s optimal action. In the absence of a surprise the DM can “rule out” the unlikely, and will reach a quick decision for the optimal action in the likely state.

3.2 Optimal Strategy

The structure of the optimal policy depends on the *cost of information* c . Intuitively the higher the flow cost, the lower is the value of experimenting. As will be seen, the experimentation region expands as the cost of information falls. More interestingly, the type of learning strategy employed also changes in a nontrivial way. The following characterization shows that there are three cases. If c is very high, immediate action is always optimal. For intermediate values of c , the optimal strategy involves only contradictory learning, and for low values of c , both contradictory and confirmatory learning occur.

Theorem 1. *For given utilities u_x^ω , $\lambda > 0$, and $r \geq 0$, there exist $\bar{c} = \bar{c}(r, u_x^\omega, \lambda)$ and $\underline{c} = \underline{c}(r, u_x^\omega, \lambda)$, $\bar{c} \geq \underline{c} \geq 0$, with strict inequalities for r sufficiently small, such that the optimal strategy is characterized as follows:¹⁷*

- (a) (**No learning**) *If $c \geq \bar{c}$, the DM takes action $x^*(p)$ without any information acquisition.*
- (b) (**Contradictory learning**) *If $c \in [\underline{c}, \bar{c})$, there exist $0 < \underline{p}^* < \check{p} < \bar{p}^* < 1$ such that for $p \in (\underline{p}^*, \bar{p}^*)$, the optimal α is given by (3.1). If $p \notin (\underline{p}^*, \bar{p}^*)$, the DM takes action $x^*(p)$ without any information acquisition.*

¹⁷See (A.20) and (A.19) in Appendix A for explicit expressions for \underline{c} and \bar{c} .

(c) (**Contradictory and Confirmatory learning**) If $c < \underline{c}$, then there exist $0 < \underline{p}^* < \underline{p} < p^* < \bar{p} < \bar{p}^* < 1$ such that for $p \in (\underline{p}^*, \bar{p}^*)$, the optimal α is given by

$$\alpha(p) = \begin{cases} 1, & \text{if } p \in (\underline{p}^*, \underline{p}), \\ 0, & \text{if } p \in (\underline{p}, p^*), \\ \frac{1}{2}, & \text{if } p = p^*, \\ 1, & \text{if } p \in (p^*, \bar{p}), \\ 0, & \text{if } p \in (\bar{p}, \bar{p}^*). \end{cases} \quad (3.2)$$

If $p \notin (\underline{p}^*, \bar{p}^*)$ the DM takes action $x^*(p)$ without any information acquisition.

We sketch the main steps of the proof in Section 3.4 below. The formal proof can be found in Appendix A.

The optimal policy in case (b) is depicted in Panel (b) of Figure 1. In case (c) the pattern is more complex (see Panel (c) of Figure 1). Both confirmatory and contradictory learning are optimal for some beliefs. Theorem 1 also states that the confirmatory region (\underline{p}, \bar{p}) is always sandwiched between two regions where the contradictory strategy is employed. That is, near the boundaries of the experimentation region, the contradictory strategy is always optimal.

The intuition for the optimal strategy is explained as follows. First, since learning is costly, there are levels of confidence, given by \bar{p}^* and \underline{p}^* , that the DM finds sufficient for making decisions without any evidence. These beliefs constitute the boundaries of the experimentation region; more extreme beliefs result in an immediate action.

Within the experimentation region $(\underline{p}^*, \bar{p}^*)$, the DM's choice depends on the trade-off between the confirmatory and contradictory strategies. The confirmatory strategy has the advantage that the DM will eventually discover evidence, and will thus never make a mistake. At the same time, full discovery of evidence could lead to a long delay. This is particularly the case when the DM is fairly certain in her belief. In that case, the contradictory strategy becomes relatively more appealing. Suppose for instance the DM's belief $p < \bar{p}^*$ is very close to \bar{p}^* . In that case, the contradictory strategy either yields a "surprise" (B -evidence), leading to a perfectly accurate decision b , or (more likely) allows the DM to "rule out" the unlikely state B and reach the desired level of confidence \bar{p}^* for action a with very little delay. Of course, the DM could instead look for A -evidence only briefly and take an action if no information arrives, but obtaining A -evidence has no value since it does not change her action. The only sensible alternative is to "go all the way" toward full discovery—i.e., the confirmatory strategy, but it takes a long time. Hence, the contradictory strategy is optimal near the stopping boundaries. As the DM becomes less certain, however, the trade-off tilts toward the confirmatory strategy, as the contradictory strategy involves a longer delay to reach the stopping region.

3.3 Comparative Statics

It is instructive to study how the optimal strategy varies with the parameters. Of particular interest are the experimentation region $(\underline{p}^*, \bar{p}^*)$, and the confirmatory region (\underline{p}, \bar{p}) . We say a region *expands* when a parameter change leads to a superset of the original region. (This includes the case that the region appears when it was empty before.) We say a region *shifts up* (*down*) when the boundaries of the region increase (decrease) strictly.

- Proposition 1** (Comparative statics). *(a) The experimentation region expands as r or c falls, and covers $(0, 1)$ in the limit as $(r, c) \rightarrow (0, 0)$.*
- (b) The experimentation region expands as u_a^B or u_b^A falls, and covers $(0, 1)$ in the limit as $(u_a^B, u_b^A) \rightarrow (-\infty, -\infty)$. Specifically, if $c < \bar{c}$, then $\underline{p}^* \searrow 0$ monotonically as $u_b^A \rightarrow -\infty$, and $\bar{p}^* \nearrow 1$ monotonically as $u_a^B \rightarrow -\infty$.*
- (c) If $c < \bar{c}$, then the experimentation region shifts down as u_a^A increases and up as u_b^B increases.*
- (d) If $\underline{c} < c < \bar{c}$, then \check{p} , the cutoff in the contradictory strategy, increases in u_b^A and decreases u_a^B .*
- (e) If $c < \bar{c}$, then the confirmatory region expands as (u_a^B, u_b^A) falls and covers $(0, 1)$ in the limit as $(u_a^B, u_b^A) \rightarrow (-\infty, -\infty)$. Specifically, if $c < \underline{c}$, then $\underline{p} \searrow 0$ monotonically as $u_b^A \rightarrow -\infty$, and $\bar{p} \nearrow 1$ monotonically as $u_a^B \rightarrow -\infty$.*

The proof can be found in Appendix B.2 in the Supplemental Material.

Parts (a) and (b) are quite intuitive. The DM acquires information for a wider range of beliefs if the cost of learning (r, c) falls, or if mistakes become more costly in the sense that (u_a^B, u_b^A) falls. A similar intuition holds for (c). The intuition for (d) is more subtle. For example, if $p < \check{p}$, the contradictory strategy may lead to taking action b in the wrong state (A). If this becomes more costly, \check{p} shifts down to avoid this mistake.

Part (e) shows that the cost of mistakes also matters for the relative appeal of the *alternative* learning strategies: the confirmatory strategy becomes more appealing when mistakes become more costly. In the limit where mistakes become completely unacceptable, the confirmatory strategy becomes optimal for all beliefs. One could imagine this limit behavior as that of a scientist who views collecting evidence of either kind—proving or disproving a hypothesis—as the only acceptable way of advancing science. Such a scientist will rely solely on the confirmatory strategy: she will initially strive to prove a hypothesis for instance if she conjectures it to be true (believes it to be more likely true than not); after a series of unsuccessful attempts to prove the hypothesis, however, she begins to doubt her initial conjecture, and when the doubt reaches a “boiling point” (i.e., p^*), she begins to put some effort to disprove it.¹⁸

¹⁸Contradictory learning could also describe some aspect of scientific inquiry if a scientist is willing to accept a small margin of error. For instance, even a careful theorist may not verify thoroughly her “proof” if she believes it to be correct. Rather, she may look for a mistake in her argument—a contradictory strategy, and without finding one, may declare it as a correct proof.

The Role of Discounting. Intuitively, one would interpret r as a cost of learning and would thus expect that c and r are substitutes in the sense that a higher discount rate requires a lower flow cost for the same structure to emerge. Formally, $\partial\bar{c}/\partial r < 0$ and $\partial\underline{c}/\partial r < 0$. The following proposition shows that this is indeed the case if at least one loss payoff (u_b^A or u_a^B) is not too small.

Proposition 2. (a) Suppose $\bar{c} > 0$. Then $\partial\bar{c}/\partial r < 0$ iff $U(\hat{p}) > 0$, which is equivalent to $u_a^A u_b^B - u_b^A u_a^B > 0$.
(b) Suppose $\underline{c} > 0$. Then $\partial\underline{c}/\partial r < 0$ if both $u_a^A > |u_b^A|$ and $u_b^B > |u_a^B|$; $\partial\underline{c}/\partial r > 0$ if $\min\{u_b^A, u_a^B\}$ is sufficiently small.

The proof can be found in Appendix B.3 in the Supplemental Material. If both loss payoffs u_b^A and u_a^B are negative and sufficiently large in absolute value, we have $\partial\bar{c}/\partial r > 0$ and $\partial\underline{c}/\partial r > 0$. A higher discount rate calls for more experimentation in this case. Intuitively, if losses are sufficiently large, the DM would prefer to delay their realization, which favors longer experimentation. This explains Part (a) of Proposition 2. Moreover, large losses in case of a mistake increase the need for accuracy, favoring the confirmatory strategy. This explains Part (b) of Proposition 2.

3.4 Sketch of the Proof

We now sketch the main arguments leading to Theorem 1 in several steps. A less technically interested reader may want to skip this section. Our method is first to “guess” the structure of the policy and then to verify its optimality. To this end, we first compute the value of alternative learning strategies. Taking an action immediately simply yields $U(p)$. To compute the value of the other strategies, we first obtain two ODEs by substituting $\alpha = 0$ and $\alpha = 1$ in the HJB equation (2.3). For given boundary conditions (\tilde{p}, W) , where $\tilde{p} \in (0, 1)$, the ODEs admit unique solutions $V_0(p; \tilde{p}, W)$ and $V_1(p; \tilde{p}, W)$, respectively.

To compute the value of the confirmatory strategy, recall that it prescribes, for some $p^* \in (0, 1)$, $\alpha = 1/2$ until evidence arrives, whereupon the DM takes an action according to the evidence. Let $\bar{U}(p)$ denote the value of this “stationary” strategy.¹⁹ We can then use $V(p^*) = \bar{U}(p^*)$ as a boundary condition for the value function and obtain

$$V_{cf}(p) := \begin{cases} V_0(p; p^*, \bar{U}(p^*)), & \text{for } p \leq p^*, \\ V_1(p; p^*, \bar{U}(p^*)), & \text{for } p \geq p^*. \end{cases}$$

We postulate smooth pasting at p^* and require $V_{cf}(p) \geq \bar{U}(p)$ for all p in a neighborhood of p^* . This uniquely pins down the reference belief as $p^* = (ru_b^B + c) / (ru_a^A + ru_b^B + 2c)$.

¹⁹We have $\bar{U}(p) = \frac{\lambda}{2r+\lambda} (pu_a^A + (1-p)u_b^B) - \frac{2c}{2r+\lambda}$. Intuitively, without mistakes, the DM achieves the “first-best” payoff $pu_a^A + (1-p)u_b^B$, but discounted and net off the expected discounted flow cost.

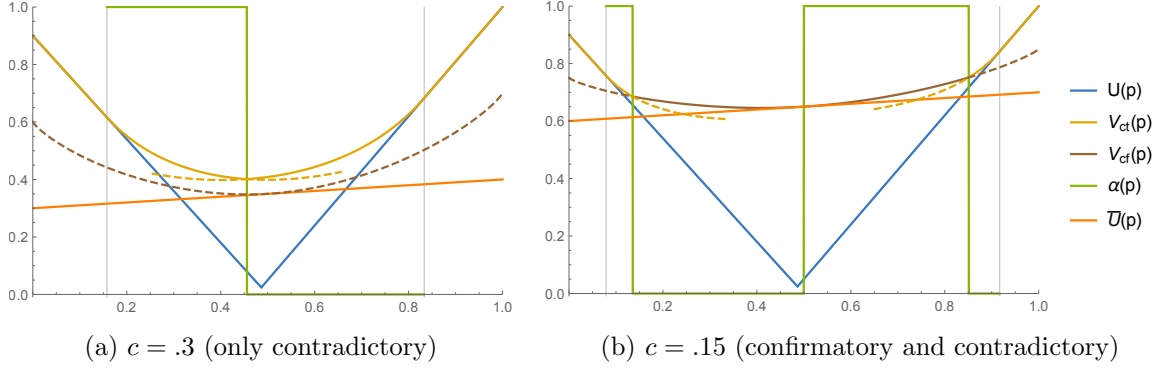


Figure 2: Value Function and Optimal Policy. The value function is the upper envelope of V_{ct} and V_{cf} (solid). ($\lambda = 1$, $r = 0$, $u_a^A = 1$, $u_b^B = .9$, $u_b^A = u_a^B = -.9$)

The value of contradictory learning is computed similarly. Intuitively, \bar{p}^* and \underline{p}^* are the beliefs at which the DM is indifferent between an optimal immediate action and contradictory learning for an instant followed by an immediate action in case of no discovery. This yields the boundary conditions $V(\underline{p}^*) = U_b(\underline{p}^*)$ and $V(\bar{p}^*) = U_a(\bar{p}^*)$. Next, we postulate that the slope of the value is equal to $U'(p)$ at these values (smooth pasting). Combining these two conditions pins down the critical beliefs \bar{p}^* and \underline{p}^* . We can thus construct the *left branch* and *right branch* of the value function for the contradictory strategy:

$$\underline{V}_{ct}(p) := \begin{cases} U_b(p), & \text{if } p \leq \underline{p}^*, \\ V_1(p; \underline{p}^*, U_b(\underline{p}^*)), & \text{if } p > \underline{p}^*; \end{cases} \quad \bar{V}_{ct}(p) := \begin{cases} U_a(p), & \text{if } p \geq \bar{p}^*, \\ V_0(p; \bar{p}^*, U_a(\bar{p}^*)), & \text{for } p < \bar{p}^*. \end{cases}$$

We combine these branches and construct the value of contradictory learning as $V_{ct}(p) := \max\{\underline{V}_{ct}(p), \bar{V}_{ct}(p)\}$.

Second, we consider a candidate solution $V_{Env}(p) := \max\{V_{ct}(p), V_{cf}(p)\}$.²⁰ In Proposition 6, we show that this yields the strategies stated in Theorem 1. Two observations are crucial. First, consider a hypothetical “full attention payoff”—the payoff the DM would attain if she could set $\alpha = \beta = 1$.²¹ We show that contradictory learning achieves this upper bound at the stopping boundaries \bar{p}^* and \underline{p}^* , while confirmatory learning achieves a strictly lower value (Lemmas 3 and 4). This implies that contradictory learning is always part of the optimal strategy since it dominates confirmatory learning at the stopping boundaries \bar{p}^* and \underline{p}^* . Second, we establish a *Crossing Property* (Lemma 5): a solution $V_1(p; x, W)$ can cross $V_0(p; x', W')$ only from above if they exceed $\bar{U}(p)$. This means that branches of the confirmatory and contradictory value function must intersect in the way illustrated in Figure 2. In particular, this implies that the choice between confirmatory and contradictory learning leads to one of the structures in Panels (b) and (c) of Figure

²⁰This value is obtained if the DM chooses optimally between confirmatory and contradictory learning based on the prior belief, and never switches.

²¹Choosing $\alpha = \beta = 1$ is not feasible for the DM in our model with limited attention. The payoff from $\alpha = \beta = 1$ is only used as a benchmark.

1.²² Case (b) of Theorem 1 arises if $V_{ct}(p^*) \geq \bar{U}(p^*)$ and case (c) arises if $V_{ct}(p^*) < \bar{U}(p^*)$.

Finally, we show that the candidate $V_{Env}(p)$ is the value function by verifying that it satisfies the HJB equation (2.3). Since $V_{Env}(p)$ is not everywhere differentiable, we show that it is a *viscosity solution* of the HJB equation. By Theorem III.4.11 in Bardi and Capuzzo-Dolcetta (1997), this implies that it is the value function of the DM’s problem.

4 Applications

The model of information acquisition with limited attention can be applied in various situations. We discuss two applications: jury deliberation and media choice by voters. For each we also discuss generalizations that go beyond the baseline model: In the case of jury deliberation we consider the inclusion of a third action (a “hung jury”) and the case of media choice, we consider news media with varying degrees of informativeness. The formal treatment of the extensions can be found in Appendices D and E in the Supplemental Material.

4.1 Jury Deliberation

In many states of the USA, when a prosecutor accuses an individual of a crime, a grand jury is impaneled to decide whether to indict the accused. Should that occur, a trial jury hears the case presented by the prosecutor and the defense attorney before returning its verdict. Both types of juries deliberate based on the evidence or lack thereof over a period of time. Our model can be used to consider their deliberation behavior and its implications on the outcome.

Suppose that the DM is a member of a jury deciding whether to “indict/convict” (action a) or “not indict/acquit” (action b) a defendant.²³ The state is either “guilty” ($\omega = A$) or “innocent” ($\omega = B$). Suppose the jury has already heard all the evidence presented in the courtroom,²⁴ and has formed a preliminary opinion summarized by the prior belief p_0 . The jury proceeds with deliberation which involves revisiting the evidence, testimonies, and arguments presented to them and scrutinizing some details jurors may have missed (i.e., juries often ask for transcripts of testimonies they wish to scrutinize). The attention decision α corresponds to the type of evidence/testimony that the jury may scrutinize in depth.

²²Note that as postulated in footnote 20, the DM never switches the mode of learning.

²³Here, we view the DM as a single juror, or a coalition of jurors acting like a single individual. The conclusions would also apply for a single judge. In practice, the collective action aspect of jury deliberation brings another layer to the problem, which is beyond the scope of this paper. See Persico (2004), Stephenson (2010), and Chan, Lizzeri, Suen, and Yariv (2016). Our application to jury deliberation is a first step toward exploring information choice in the jury context.

²⁴Only the prosecutor presents the case in grand jury proceedings, whereas in a trial, both prosecutor and defense attorney present the evidence.

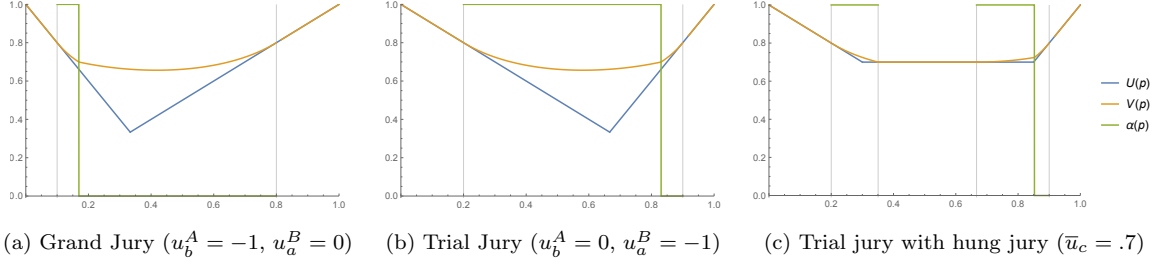


Figure 3: Deliberation of grand and trial juries. ($\lambda = 1$, $r = 0$, $c = .2$, $u_a^A = u_b^B = 1$.)

One salient difference between the two juries is the costs of making alternative errors: while a trial jury perceives a high cost of convicting an innocent (i.e., a low value of u_a^B), a grand jury must be more concerned about failing to indict a guilty defendant (i.e., a low value of u_b^A). The evidentiary standards as well as the decision rule employed by US courts reflect this difference.²⁵ Proposition 1 predicts how this difference causes grand juries and trial juries to deliberate differently, and shows that the de facto evidentiary standards—reflected by the boundaries \underline{p}^* and \bar{p}^* —differ for the two juries.

For ease of exposition we use examples to illustrate our findings which hold generally. Panels (a) and (b) of Figure 3 depict examples where the parameters are chosen so that Part (b) of Theorem 1 obtains.

For the grand jury in panel (a), probable cause corresponds to $p > \bar{p}^* = .8$. If the prior falls in this range, the jury returns an indictment right away. Conversely, the accused is acquitted immediately only if $p_0 < \underline{p}^* = .1$. This shows that the stronger concern about acquitting a guilty person leads the grand jury to require a higher standard for this decision. If the grand jury does not return a decision immediately, for most beliefs ($p \in (.17, .8)$ in the example), the jury will look for exculpatory evidence.²⁶ If no such evidence is found, the jury becomes more convinced that there is “probable cause,” and ultimately returns an indictment.

The trial jury deliberates differently. The evidentiary standard for conviction is higher ($\bar{p}^* = .9$ in the example), whereas the standard for acquittal is lower ($\underline{p}^* = .2$), reflecting the greater concern about convicting an innocent person. Moreover, if no immediate verdict is returned, for most beliefs ($p \in (.2, .83)$ in the example), the jury scrutinizes incriminating evidence. Not finding such evidence pushes the jury’s belief toward “reasonable doubt” and an acquittal.

In many cases, unanimity is required for a jury decision. If no consensus is reached, a *hung jury* arises. In the US this leads to a mistrial with the possibility to retry the case. Although a consequence of unanimity rule adopted by Common Law courts for criminal

²⁵In a criminal case, conviction in a criminal court requires proving guilt “beyond the reasonable doubt” whereas the standard for grand jury indictment is “probable cause.” The decision rule for trial jury is unanimity, whereas a grand jury indictment requires concurrence by 12 members out of 16-23 total members.

²⁶If $u_a^A = u_b^B$ and $u_a^B > u_b^A$, then $\underline{p}^* < 1 - \bar{p}^*$ by part (b), and $\check{p} < 1/2$ by part (d) of Proposition 1.

cases, in principle a mistrial need not be the only solution to a jury deadlock,²⁷ and its effect on jury deliberation is an important question. To analyze the effect, we introduce hung jury in “reduced form” as a third action c , which we assume to have a safe payoff of $u_c^A = u_c^B = \bar{u}_c$.²⁸ Panel (c) of Figure 3 depicts this case for a trial jury. The effects of introducing the option of a hung jury can be summarized as follows.

First, the possibility of a hung jury does not change the evidentiary standard for either verdict: they remain $\underline{p}^* = .2$ for acquittal and $\bar{p}^* = .9$ for conviction. The reason is that these standards are set by indifference between contradictory learning and immediate action, and the values of these strategies are not affected by the hung jury option.

Second, the hung jury option affects the deliberation strategy in favor of seeking incriminating evidence (B -evidence). The value of seeking incriminating evidence increases since the jury faces a more appealing option, namely to “settle” for mistrial instead of deliberating until they reach the reasonable doubt. This means that for some pre-deliberation beliefs ($p_0 \in [.83, .85]$), the jury switches attention from exculpatory evidence to incriminating evidence. The flip side of this behavior is that for a large range of pre-deliberation beliefs ($p_0 \in [.35, .67]$) a mistrial is declared.

Third, the change in deliberation behavior also affects verdicts. For $p_0 \in [.67, .85]$, the probability of guilty verdict falls.²⁹ Remarkably, the probability of “not-guilty” verdict also falls. In fact, with the mistrial option, the jury never returns a “not-guilty” verdict for any pre-deliberation belief $p_0 \in [.35, .85]$ above the lower bound of the mistrial region, whereas, without that option, the jury would have returned “not guilty” verdict with positive probability after a long deliberation (when it reaches \underline{p}^*).

These findings suggest that the possibility of hung jury has complex and nontrivial effects on jury deliberation and verdicts and also point to the richness of the prediction that can be brought out by use of a dynamic model such as the current one.³⁰

²⁷For instance, in England and New Zealand, the unanimity requirement is relaxed in the event of a jury deadlock, allowing a judge to require supermajority instead. Hung jury is not possible in a Scottish court where a simple majority is required for a verdict. See https://en.wikipedia.org/wiki/Hung_jury.

²⁸An interpretation is that this is a continuation value jurors assign to a mistrial. We formally analyze the model with a third action in Appendix D in the Supplemental Material.

²⁹Fix any belief $p_0 \in [0.83, 0.85]$ in the switching region. Without hung jury, the probability of a guilty verdict exceeds p_0 since the jury never finds the guilty innocent whereas it may find the innocent guilty (in case the jury reaches \bar{p}^* without finding exculpatory evidence). With hung jury, the probability of guilty verdict falls strictly below p_0 by an analogous argument. For $p_0 \in [0.67, 0.83]$, mistrial is declared in any event in which, without hung jury, the jury would have found incriminating evidence after p fell below 0.67. Hence, the probability of a guilty verdict is reduced commensurately.

³⁰The analysis so far focused on the situation in which, without hung jury, jury deliberation involves only contradictory learning. The case with both confirmatory and contradictory learning is analogous: introduction of mistrial increases the value of confirmatory learning, which reduces the probability of the guilty verdict for high p (around \bar{p}) and that of the not guilty verdict for a low p (around \underline{p}), at the expense of increased probability of mistrial.

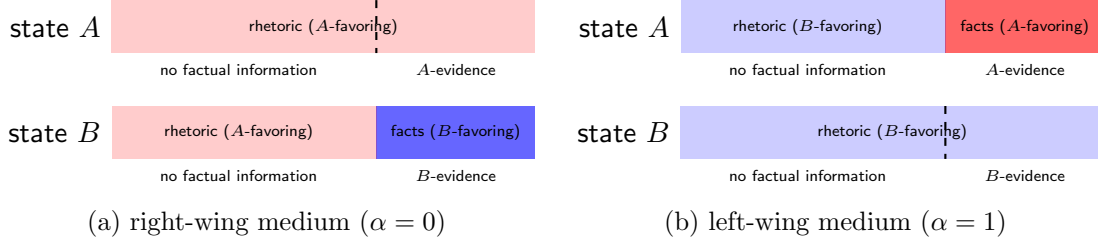


Figure 4: Media bias: Labels below the bars describe the information the medium has learned. Colors and labels inside the bars describe informational content (the message) published in each case.

4.2 Choice of News Media by Voters

Our second application is the choice of news media by a voter.³¹ Suppose a citizen votes for one of two candidates, a or b . Candidate a has a right-wing platform, and candidate b has a left-wing platform.³² The state $\omega \in \{A, B\}$ captures which platform serves the interest of the voter better, and u_x^ω represents the voter’s utility from voting for x in state ω . Her belief about the state is captured by p . We say the voter is more right-leaning the higher p is.³³

Before voting, the voter may obtain information from a range of “news media” or “news channels,” at a flow cost $c > 0$.³⁴ We now interpret each choice $\alpha \in [0, 1]$ as a particular news medium the voter may subscribe to.³⁵ Medium α publishes three possible types of content: (i) factual information in favor of A , (ii) factual information in favor of B , (iii) partisan rhetoric that corresponds to the bias of a medium. Rhetoric is a message that does not convince voters that there is factual information in favor of any candidate. Over a short time interval $dt > 0$, the voter receives factual messages supporting A or B with probabilities (i) $p\alpha\lambda dt$, or (ii) $(1-p)(1-\alpha)\lambda dt$, respectively, and (iii) in the absence of such messages, the voter perceives the pulished content as rhetoric, which happens with the remaining probability $1 - [p\alpha + (1-p)(1-\alpha)]\lambda dt$.

³¹This application shares a theme with Suen (2004), Burke (2008), Oliveros (2015), and Yoon (2016)—particularly the optimality of consuming a biased medium for a Bayesian agent. These papers are largely static, unlike the current model which is fully dynamic. Meyer (1991) makes a similar observation in a dynamic contest environment. Mullainathan and Shleifer (2005) assumes “behavioral” bias on the part of consumers to predict media slanting. Gentzkow and Shapiro (2008) and Perego and Yuksel (2016) study media competition, which we take as given.

³²In the context of a referendum, one could interpret $x = a$ as voting for the status quo and $x = b$ for an alternative, or vice versa.

³³Alternatively one could model a voter’s bias in terms of her payoffs. In this case u_x^ω could incorporate her “partisan” preference.

³⁴Since the time at which the payoff is realized (the election or implementation of a policy) is independent of when the voter makes up her mind, we assume $r = 0$. Our model does not capture the effect of a “deadline,” which is clearly relevant for the election example. We conjecture that the salient features of our characterization which are discussed in the context of this application carry over to a model with sufficiently long deadline, but leave the formal analysis of this case for future work.

³⁵The notation suggests that voters subscribe to a single medium. Given that the arrival rates of evidence are linear in α , this is without loss—a “multi-homing” voter cannot achieve a higher payoff.

The informational content of partisan rhetoric—and thus the partisan bias—depends on the medium α . To illustrate this, consider the most extreme media. (Media bias for the general case $0 \leq \alpha \leq 1$ is more precisely micro-founded in Appendix C of the Supplemental Material.) As depicted in Panel (a) of Figure 4, medium $\alpha = 0$ sends B messages in state B only when it learns facts supporting B , which happens with *some* probability. Otherwise it sends partisan rhetoric favoring the right-wing candidate a both in state B and in state A .³⁶ We therefore call medium $\alpha = 0$ *right-wing*. Intuitively, a right-wing medium will advocate the left-wing candidate b only if it possesses hard facts supporting state B . On the other hand, it is willing to advocate the right-wing candidate even in the absence of factual information. Therefore, even if it possesses factual information in favor of a , voters will not distinguish this from partisan rhetoric.

As depicted in Panel (b), medium $\alpha = 1$ has the opposite bias. It always publishes (left-wing) rhetoric in state B . In state A it also publishes left-wing rhetoric, except when it learns factual evidence in favor of candidate a . Corresponding to the bias of the rhetoric we call such a medium *left-wing*. More generally, a medium $\alpha \in (0, 1/2)$ is *moderately right-wing*, a medium $\alpha \in (1/2, 1)$ is *moderately left-wing*, and the medium $\alpha = 1/2$ is called *unbiased*.

We treat the behavior of media as exogenous to the model and focus on optimal subscription choices by voters for a given set of media. With this interpretation of the model, Theorem 1 has the following implications for the choice of news media.

Corollary 1. *Voters with extreme beliefs $p \notin (\underline{p}^*, \bar{p}^*)$ always vote for their favorite candidates without consulting media. Those subscribing to media exhibit the following behavior.*

(a) (**Moderately informative media**) Suppose $\underline{c} \leq c < \bar{c}$:

- All voters to the right of \check{p} subscribe to right-wing media and all voters to the left of \check{p} subscribe to left-wing media.
- Over time absent breakthrough news, all voters become progressively more extreme and polarized.

(b) (**Highly informative media**) Suppose $c < \underline{c}$:

- (i) Right-leaning voters ($p \in (\bar{p}, \bar{p}^*)$) and moderate left-leaning voters ($p \in (\underline{p}, \underline{p}^*)$) subscribe to right-wing media. (ii) Left-leaning voters ($p \in (\underline{p}^*, \underline{p})$) and moderate right-leaning voters ($p \in (\underline{p}^*, \bar{p})$) subscribe to left-wing media. (iii) Undecided voters ($p = p^*$) subscribe to unbiased media.
- Over time absent breakthrough news, moderate voters ($p \in (\underline{p}, \bar{p})$) become increasingly undecided and more extreme voters become increasingly more extreme and polarized.

Panel (a) of Figure 5 shows the choice of media α by voters with different beliefs.

³⁶Note that the partisan bias of the rhetoric is consistent with Bayesian updating. Since rhetoric is more frequent in state A , it must be in favor of the right-wing candidate.

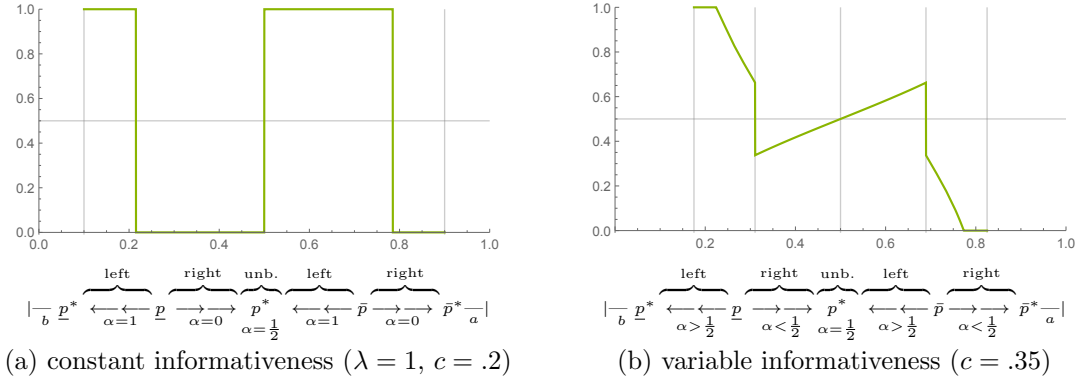


Figure 5: Optimal media choice. ($r = 0, u_a^A = u_b^B = 1, u_a^B = u_b^A = -1$)

Figure 6 shows how the distribution of beliefs evolves over time (measured at three different times), where colors represent the media choice by voters who are still subscribing to media. The initial distribution is $p_0 \sim U[0, 1]$ in this example.

Among those subscribing to news media, voters on the far right choose right-wing media. Such a medium is valuable to them since it mostly publishes a content reinforcing their belief, and publishes left-favoring information *only if* it is accurate enough (in fact, fully conclusive in the baseline model) to actually convince them to change their votes. Over time, in the absence of convincing contradictory news, the right-wing medium feeds such voters with what they believe, leading them to become more extreme in their beliefs. Hence, applied to dynamic media choice with fully Bayesian voters, our model generates self reinforcing beliefs—sometimes called an “echo-chamber” effect—that persists until strong contradictory evidence arrives.

The moderately right voter’s media choice is quite different. Their moderate beliefs cause them to seek accurate evidence (of either kind) for voting. Initially they look for right-favoring evidence which they expect more likely to arise given their beliefs. Interestingly, they expect to find such evidence in left-wing media, since these media scrutinize the right-favoring information more and apply a very high standard for reporting

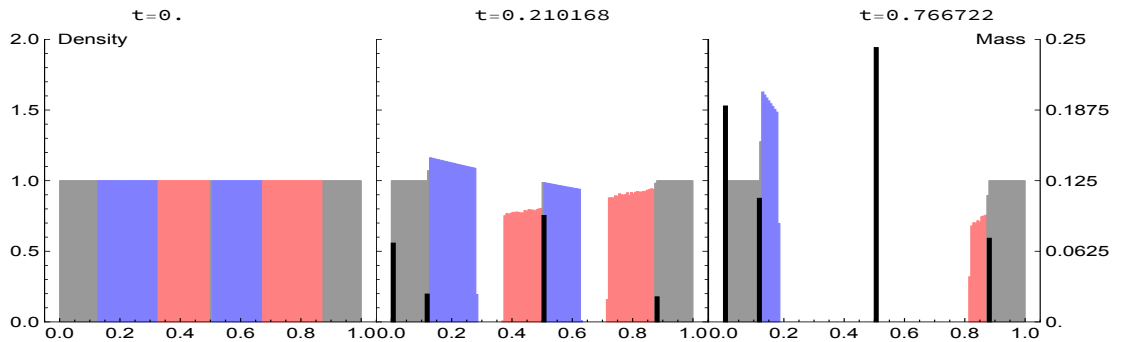


Figure 6: Evolution of media choice and beliefs when the true state is B . Shaded areas represent the density of beliefs (left axis). Bold bars represent mass points of beliefs (right axis).

such information. The moderate’s media choice thus differs from the extreme voter’s choice. The prediction of an “anti-echo chamber” effect is novel and has no analogue in previous literature. Over time, absent breakthrough news, their anti-echo chamber choice leads voters to be undecided. Ultimately, they switch to the unbiased medium.

In sum, our dynamic model of media choice predicts two different dynamics of belief evolution resulting from optimal media choice: the beliefs for those who are sufficiently extreme become more polarized, and the beliefs of those who are sufficiently moderate converge toward the middle and result in the subscription of unbiased media.

Trade-off between Skewness and Informativeness of Media. One prediction of the baseline model is that voters only choose from three media, right-wing, left-wing, and unbiased. This is a consequence of the assumption that all media have access to information with identical arrival rates, and differ only in the rhetoric they publish. Formally, each news medium is characterized by a pair of arrival rates (λ^A, λ^B) for A - and B -evidence, respectively. The sum of arrival rates $\lambda^A + \lambda^B$ can be viewed as a measure of the *informativeness* of a medium. In the baseline model, if we normalize $\lambda = 1$, we have $\lambda^A = \alpha$ and $\lambda^B = \beta = 1 - \alpha$, so that $\lambda^A + \lambda^B = 1$ for all media. The analysis of this case suggests that there is no demand for moderate media and the market will be dominated by extremely-biased outlets, as long as all media are identical in their informativeness.

We relax this assumption and introduce a trade-off between bias and informativeness by assuming less biased media to be more informative. This generates predictions about which voters have a stronger preference for informativeness versus bias.

Formally we assume that $\lambda^B = \Gamma(\lambda^A)$, where $\Gamma(\lambda^A)$ is a decreasing concave function, and set $\alpha = \lambda^A / (\lambda^A + \Gamma(\lambda^A))$. The interpretation is that news outlets that are more balanced have access to more factual information, for example because they employ journalists focused more on hard evidence.³⁷ In Appendix E in the Supplemental Material, we extend our characterization to this more general framework and show that it leads to choices of moderately biased media. Here, we illustrate the optimal choice in Panel (b) of Figure 5 with Γ given by: $(\lambda^A + \Gamma(\lambda^A)) + 3\sqrt{(\lambda^A)^2 + (\Gamma(\lambda^A))^2} = 4$.

The overall structure of the optimal policy is similar, with extreme voters choosing media corresponding to their own bias and moderates doing the opposite. Regarding the trade-off between informativeness and bias, the model predicts that more moderate voters have a stronger preference for informativeness, whereas voters with more extreme beliefs, prefer more biased outlets. This reinforces the echo-chamber effect as extreme voters turn to more and more biased media as they become more polarized over time. Moderate voters who becomes less and less certain over time shift toward more and more balanced

³⁷One could consider the opposite case where more biased outlets have access to more factual information which would reinforce the demand for skewed media that is already present in the baseline model. Formally, a convex function Γ will have no effect on the optimal choices, as long as the DM can divide attention across different news media in linear fashion.

outlets. Overall, absent breakthrough news, demand for outlets with moderate biases falls over time, as extreme voters become more extreme and switch to more extreme media; and moderate voters become more moderate and eventually subscribe to balanced media.

5 Generalized Model with Non-Conclusive Evidence

In the baseline model, we have assumed that the DM can access fully revealing signals. We now generalize the model to allow for the signals to be noisy. A -signals now arrive at rate $\bar{\lambda}$ in state A , and at rate $\underline{\lambda}$ in state B , where we assume $\bar{\lambda} > \underline{\lambda} \geq 0$. Importantly, the DM does not observe the state generating these signals. Similarly, B -signals arrive at rate $\bar{\lambda}$ in state B , and at rate $\underline{\lambda}$ in state A .³⁸ If $\underline{\lambda} = 0$, the signals are fully revealing, and we obtain the baseline model. We impose the following assumption

Assumption 1. *Either $r = 0$, or $u_a^A \bar{\lambda} + u_b^A \underline{\lambda} > 0$ and $u_b^B \bar{\lambda} + u_a^B \underline{\lambda} > 0$.*

This assumption means that losses associated with the wrong action are not too large. Without this assumption, delayed action may be beneficial because this leads to discounting of losses. The DM would prefer to learn as slowly as possible. Assumption 1 guarantees that for all $c \geq 0$, the DM has a preference to speed up learning.³⁹

With noisy signals, the DM may find it optimal to wait for more than one signal before taking an action. Our characterization from the baseline model generalizes naturally, if the DM finds it optimal *not* to wait for multiple signals. In this case, we say the model satisfies the *single experimentation property* (SEP). We provide a necessary and sufficient condition for SEP in terms of a critical cost level and characterize the optimal strategy under SEP. In Section 6, we illustrate what the optimal strategy looks like when SEP does not hold.

5.1 Optimal Strategy under SEP

As in the baseline model, the DM's attention strategy determines which signals are actually observed. If she allocates a fraction $\alpha \in [0, 1]$ of her attention to seeking A -signals and $\beta = 1 - \alpha$ to seeking B -signals, she receives an A -signal with arrival rate $\alpha \lambda^A(p)$, where $\lambda^A(p) := p \bar{\lambda} + (1 - p) \underline{\lambda}$. An A -signal leads to a posterior $q^A(p) := \bar{\lambda} p / (\bar{\lambda} p + \underline{\lambda} (1 - p)) > p$. Similarly B -signals arrive at rate $(1 - \alpha) \lambda^B(p)$ where $\lambda^B(p) := p \underline{\lambda} + (1 - p) \bar{\lambda}$, and lead to posterior $q^B(p) := \underline{\lambda} p / (\bar{\lambda} p + \underline{\lambda} (1 - p)) < p$. If the DM does not receive any signal, her belief evolves according to $\dot{p}_t = -(2\alpha_t - 1) \delta p_t (1 - p_t)$, where $\delta := \bar{\lambda} - \underline{\lambda}$ denotes the difference in arrival rates between signal and noise. We now state the optimal strategy, with

³⁸Allowing the DM to choose any signal with state contingent arrival rates $(\lambda_A, \lambda_B) \in [\underline{\lambda}, \bar{\lambda}]^2$ does not change the analysis; the DM will optimally choose only between $(\lambda_A, \lambda_B) = (\underline{\lambda}, \bar{\lambda})$ and $(\lambda_A, \lambda_B) = (\bar{\lambda}, \underline{\lambda})$.

³⁹We could instead impose weaker conditions that depend on c . In the proofs we use $\hat{U}(p) > \bar{U}(p)$ for all p , $c + rU_b(\underline{p}^*) \geq 0$ and $c + rU_a(\underline{p}^*) \geq 0$. These are implied by Assumption 1.

the detailed analysis and explicit definitions relegated to Appendix A. As before, the optimal policy is a combination of immediate action, contradictory learning and confirmatory learning.

Theorem 2. *For any $(r, u_x^\omega, \bar{\lambda}, \underline{\lambda})$, there exists a critical cost-level $\underline{c}^{SEP} = \underline{c}^{SEP}(r, u_x^\omega, \bar{\lambda}, \underline{\lambda}) \geq 0$ such that SEP holds if and only if $c \geq \underline{c}^{SEP}$.⁴⁰ Given $c \geq \underline{c}^{SEP}$, there exists $\bar{c} = \bar{c}(r, u_x^\omega, \bar{\lambda}, \underline{\lambda}) > 0$ such that the optimal strategy is characterized as follows:*

- (a) *Suppose $c \geq \bar{c}$. Then the DM takes an optimal action $x^*(p)$ without information acquisition.*
- (b) *Suppose $c \in [\underline{c}^{SEP}, \bar{c})$. For $p \in (\underline{p}^*, \bar{p}^*)$, the optimal strategy takes either the contradictory form (3.1) or the mixed contradictory and confirmatory form (3.2). For $p \notin (\underline{p}^*, \bar{p}^*)$, it is optimal to take an immediate action. In particular, (3.1) is optimal for c in a neighborhood of \bar{c} , and (3.2) is optimal for c in a neighborhood of \underline{c}^{SEP} if r and $\underline{\lambda}$ are both sufficiently low.*

5.2 Implications for Stochastic Choice and Response Time

The general model exhibits rich implications for the choice process. We explore them and compare them with some well-known benchmarks. In classic models such as Luce’s logit model, stochastic choice can be interpreted as arising from unobserved random utility. In contrast, in models of information acquisition such as Wald, RI, and DDM, as well as our model, choice is stochastic because the DM learns the payoff-relevant state with noise, so that the choice probabilities depend on the DM’s prior. This is not the case in random-utility models. There is a particular structure to this dependence and we first explore the predictions for the accuracy and delay of decisions across DMs with different prior beliefs. Second, we explore the predicted time pattern of choice in a single DM with a fixed prior.

Accuracy of decisions. A measure of the accuracy of the decision is the average posterior belief that the DM holds when taking a particular action. This measure captures the (subjective) probability that an action is optimal conditional on being chosen. In the RI model, the accuracy does not vary with the DM’s initial belief for the range of priors where the DM acquires information (Matejka and McKay, 2015). This property is also shared by the DDM in which the DM makes a decision if and only if her belief drifts to one of the stopping boundaries that are constant over time.⁴¹ The accuracy, or the belief, at the boundary therefore determines the accuracy of the decision.

⁴⁰We have $\underline{c}^{SEP} = 0$ if $r > 0$ and $\underline{\lambda}$ is sufficiently close to zero.

⁴¹This holds in our framework with binary states and actions (Peskir and Shiryaev, 2006). Below we will discuss the DDM with a continuum of states considered by Fudenberg, Strack, and Strzalecki (2017) which generates prior-dependence.

In our model, by contrast, the decision accuracy varies with the initial belief. The reason is that decisions are not only taken when the belief has reached the bound of the experimentation region. For both the contradictory and the confirmatory region, decisions are also taken after signals are observed, in which case the belief will land in the interior of the stopping region. Moreover, since our optimal strategy is a combination of the contradictory strategy, which has a drift to bound structure, and the confirmatory strategy, where actions are only taken if signals are observed, the accuracy depends on the initial belief in a complex way, as depicted in Panel (a) of Figure 7, which also reports the accuracy in the RI-model (the constant dashed lines).

Most striking is the difference between the contradictory region and the confirmatory region. The higher initial uncertainty in the latter is associated with a significantly higher accuracy of the final action. To make this precise, we say that a DM with belief p makes a **uniformly more accurate decision** than a DM with belief p' if the following holds: *For any two histories h and h' , resulting from p and p' , respectively, where the DM takes the same action $x \in \{a, b\}$, the posterior belief at h must be more extreme in favor of action x than the posterior belief at h' .*

Proposition 3 (“Skepticism fosters accuracy”). *Suppose $c \in [\underline{c}^{SEP}, \bar{c})$ and $(\underline{p}, \bar{p}) \neq \emptyset$. Then, a DM with any prior belief $p \in (\underline{p}, \bar{p})$ makes a uniformly more accurate decision than a DM with any prior belief $p' \in (\underline{p}^*, \underline{p}) \cup (\bar{p}, \bar{p}^*)$.*

Proposition 3 makes precise the intuition why the DM employs the confirmatory strategy when she is uncertain and the contradictory strategy when she is more certain. As the proposition shows, the former strategy involves more information acquisition than the latter. Since the value of information is higher when the DM is initially more uncertain, she employs a strategy that is more effective in fuller learning (i.e., the confirmatory strategy). As shown next, this has a cost:

Delay. Let $\tau(p)$ denote the *expected delay* of taking an action, where p is the prior belief of the DM. We show that $\tau(p)$ is single-peaked as illustrated in Panel (b) of Figure 7.

Proposition 4 (“Skepticism entails delay”). *Suppose $c \in [\underline{c}^{SEP}, \bar{c})$. Then $\tau(p)$ is quasi-concave with maximum at $p' \in (\underline{p}^*, \bar{p}^*)$ if either of the following holds:*

- (a) *The contradictory strategy is optimal. In this case $p' = \check{p}$.*
- (b) *The optimal strategy combines contradictory and confirmatory evidence, r is sufficiently close to zero, and $p^* \in (\underline{\lambda}/(\underline{\lambda} + \bar{\lambda}), \bar{\lambda}/(\underline{\lambda} + \bar{\lambda}))$. In this case $p' \in [\underline{p}, \bar{p}]$.*

Clearly, in the contradictory region, the delay increases in the distance between the prior belief and the boundary of the experimentation region. For the confirmatory region, we show that the delay is concave in p . If payoffs are symmetric, the longest expected delay arises at $p^* = 1/2$. Intuitively, at $p \neq 1/2$ a decision is made faster since the DM

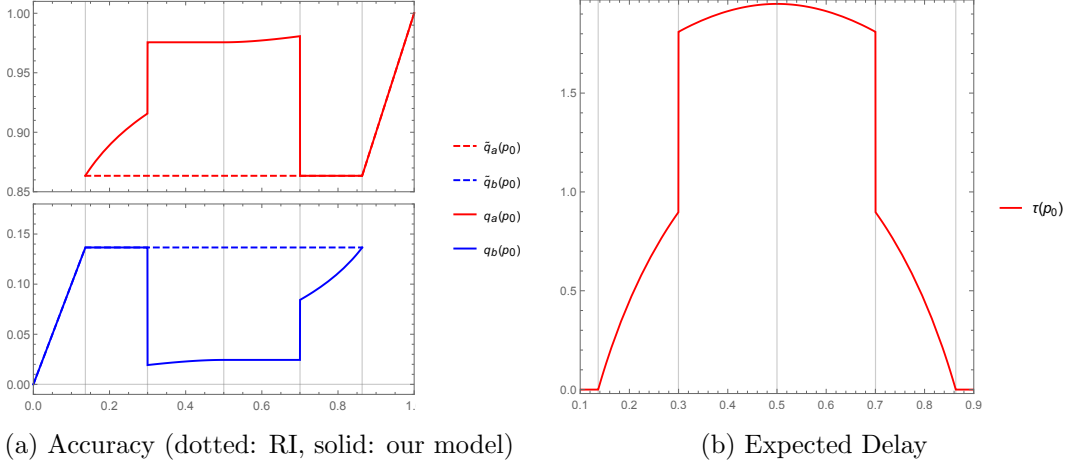


Figure 7: Parameters: $\bar{\lambda} = 1$, $\underline{\lambda} = .025$, $r = 0$, $c = 0.115$, $u_a^A = u_b^B = 1$, $u_b^A = u_a^B = 0$. The weight of the cost function in the RI-model is $\mu = 1.025$, which leads to identical experimentation regions in both models. Grid lines at: \underline{p}^* , \underline{p} , p^* , \bar{p} , \bar{p}^* (left to right).

focuses her attention on the signal that is more likely to arrive.⁴² Finally, we show using a revealed preference argument that at \underline{p} and \bar{p} , the delay in the contradictory strategy must be shorter to compensate for the less accurate decision.

Similar results hold in the RI model and DDM. The difference is that in our model, the two modes of learning lead to discontinuities in the average delay when the posterior moves from the confirmatory to the contradictory region.⁴³

Speed-Accuracy Complementarity.

So far, we have analyzed properties of the stochastic choice as a function of the prior belief. Next, we obtain a prediction on how, for a given prior, the accuracy varies as a function of decision time.

Proposition 5. *Suppose $c \in [\underline{c}^{SEP}, \bar{c}]$. For a DM with belief in $p_0 \in [\underline{p}^*, \underline{p}] \cup [p^*, \bar{p}]$, conditional on taking a , a later decision results in lower accuracy. For a DM with $p_0 \in [\underline{p}, p^*] \cup [\bar{p}, \bar{p}^*]$, conditional on choosing b , a later decision results in lower accuracy.*

We obtain *speed-accuracy complementarity*, meaning later decisions tend to be less accurate. The reason is that (regardless of the prior), the DM becomes less and less convinced of the state she is seeking evidence for. Hence, if she discovers, say A -evidence, after a long period of waiting, she updates to a lower posterior than if she discovers A -evidence quickly (formally, $q^A(p_t) < q^A(p_0)$).

Whether a long delay produces a more or less accurate decision has been of considerable interest in cognitive psychology. In a typical experiment in this area, subjects make

⁴²In the asymmetric case, p^* differs from $1/2$ so that the DM sometimes focuses attention on the less likely signals. τ is still concave in this case but the peak may differ from both p^* and $1/2$.

⁴³In the static RI model, we can interpret the (average) cost incurred as delay and the result follows from concavity of mutual information. In the DDM, the result follows straightforwardly from the convexity of the value function (see Peskir and Shiryaev, 2006, Theorem 21.1).

choices over time—often of perceptual judgment type (see [Ratcliff and McKoon, 2008](#), for a survey). What emerges as a pattern is that decisions with a longer response time tend to be less accurate. This phenomenon, known as *speed-accuracy complementarity*, is difficult to explain via the DDM often used to fit the data, unless the stopping boundaries collapse over time.⁴⁴ [Fudenberg, Strack, and Strzalecki \(2017\)](#) justify collapsing boundaries by introducing rich uncertainty in terms of payoff differences between two actions, and produce a less accurate decision over time.⁴⁵ We obtain a similar prediction in a Poisson model but the mechanism behind our result is different.⁴⁶

6 Repeated Experimentation

So far we have focused on the situation in which the SEP holds, which requires the cost of information acquisition to be sufficiently high (Theorem 2). We now discuss cases where the SEP fails. In the examples we solve, the optimal solution retains the same structure as in Theorems 1 and 2.⁴⁷ In particular, the optimal policy combines immediate action, confirmatory and contradictory evidence. At the boundaries of the experimentation region, it is still optimal to seek contradictory signals, and the confirmatory region appears if the cost of information acquisition is sufficiently low. Unlike when SEP holds, however, the DM may optimally switch between different modes of learning, as discovery of evidence while following one strategy may lead to a posterior for which a different learning strategy is prescribed.

Figure 8 depicts an optimal policy that exhibits one such pattern.⁴⁸ The colored bars indicate originating beliefs on top, and the corresponding regions where the posterior lands after receiving a signal on the bottom. For instance, discovery made in top I will

⁴⁴The DDM assumes that a subject’s perception follows a diffusion process with drift reflecting the underlying stimulus. In DDMs with constant boundaries, the threshold beliefs for decisions do not change over time, so a decision is made with the same accuracy over time, and *speed-accuracy independence* follows. If the boundaries collapse over time, then a later decision is made with less extreme or more uncertain beliefs, so speed-accuracy complementarity can be explained. See [Drugowitsch, Moreno-Bote, Churchland, Shadlen, and Pouget \(2012\)](#) and [Hawkins, Forstmann, Wagenmakers, Ratcliff, and Brown \(2015\)](#) for discussions of the DDM with collapsing boundaries and the related evidence.

⁴⁵Given the uncertainty about payoff differences, a long passage of time without observing strong evidence in favor of either action indicates that payoff differences are likely to be small. This reduces the incentive to acquire further information, causing the DM to stop with lower accuracy.

⁴⁶The extent to which Poisson signals accurately describe decision making in these experiments, or the underlying neurophysiology, is not clear. Whether brain processes stimuli through a gradual process described by DDM or a discrete process described by a Poisson model is subject of active research in neuroscience (see [Latimer, Yates, Meister, Huk, and Pillow \(2015\)](#) for instance). Within a Poisson model, our reliance on SEP further limits the applicability of our model in explaining these experiments. Nevertheless, our finding points to the potential usefulness of a Poisson-based approach; we leave for future research extension of the current model toward this direction.

⁴⁷Providing a general characterization is difficult since the number of possible strategies increases. Our analysis of examples studied in Appendix F in the Supplemental Material reveals both the difficulties and the potential ways to overcome them.

⁴⁸A plot of the value function and posterior as well as other cases can be found in Appendix F of Supplemental Material.

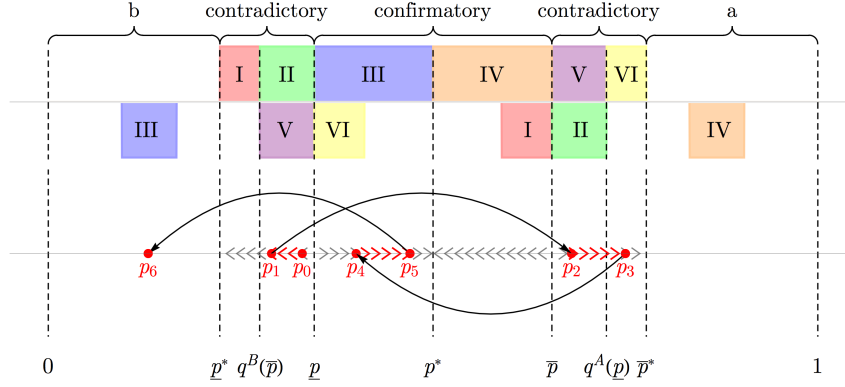


Figure 8: Repeated Experimentation. $\bar{\lambda} = 1$, $\underline{\lambda} = .2$, $r = 0$, $c = .1$, $u_a^A = u_b^B = 1$, $u_b^A = u_a^B = 0$.

lead to a posterior in bottom I. In this example, discovery made in a contradictory region (I-II and V-VI on the top) leads to a posterior in another part of the contradictory region or in the confirmatory region, that is, it will never lead to an immediate action.⁴⁹

The bottom graph of Figure 8 illustrates a possible evolution of the belief under the optimal strategy. Suppose the DM has a prior belief p_0 in top II region. She then looks for A -evidence, as prescribed by contradictory learning. Her belief drifts down until she discovers A -evidence. Suppose she does so at p_1 , still within region II. Her belief then jumps up to p_2 (in bottom II, or top V, region). Now she looks for B -evidence, as prescribed by contradictory learning for that belief. Her belief drifts up until she finds B -evidence. Suppose she does so at p_3 in top VI region. Her belief then jumps down to p_4 in bottom VI, or top III, region, a confirmatory region where she again looks for B -evidence. Suppose she discovers B -evidence at p_5 . This time her belief jumps down to p_6 , in bottom region III, which is outside the experimentation region. She now picks action b , completing the journey.

7 Conclusion

We have studied a model in which a decision maker may allocate her limited attention to collecting different types of evidence that support alternative actions. Assuming Poisson arrival of evidence, we have shown that the optimal policy combines *immediate action*, *contradictory learning*, and *confirmatory learning* for different prior beliefs. We have then used this characterization to yield rich predictions about information acquisition and choices several in real-world settings.

We envision several avenues of extending the current work. First, our model is relatively tractable (e.g., the value function and optimal policy are available in closed form).

⁴⁹As in the case where SEP holds, a discovery made in the confirmatory region (top regions III-IV) leads to a posterior in the stopping region and an immediate action.

Therefore, we hope that this framework will be useful to integrate dynamic choice of information in applied theory models. This includes our applications of jury and media choice, which could be extended beyond our analysis in the present paper. Dynamic information choice might be also embedded into a principal-agent setup in which a principal tries to induce an agent to acquire information or a strategic setup such as R&D competition where different firms choose from alternative innovation approaches over time.

Second, one may relax the “lumpiness” of information arrival to allow for more gradual information acquisition. The previous section about repeated experimentation already points to an extension in this direction, and suggests that our characterization is robust to such a generalization. A complete analysis of this case will be useful for applications in which decision makers learn gradually, such as a researcher who makes day-to-day decisions about the next steps in a project, as opposed to a manager who decides only based on reports that are made once a month. We leave this for future research.

A Proofs of Theorems 1 and 2

We present all Lemmas used in the proofs of Theorems 1 and 2 for the general model with non-conclusive evidence. The results specialize to the baseline model by setting $\bar{\lambda} = \lambda$ and $\underline{\lambda} = 0$. We start by formally stating the DM’s problem for the general model in Section 5. We analyze the problem assuming that the DM is forced to take an action after receiving the first signal. This is without loss in the baseline model. In the general model we will show it is without loss if and only if $c \geq \underline{c}^{SEP}$ for some cutoff $\underline{c}^{SEP} \geq 0$.

A.1 The DM’s problem

As in the baseline model, the DM chooses an attention strategy (α_t) and a time $T \in [0, \infty]$ at which she will stop acquiring information if she does not observe any signal up to time T . Her problem is thus given by

$$V^*(p_0) = \max_{(\alpha_\tau), T} \int_0^T e^{-rt} P_t(p_0, (\alpha_\tau)) [\alpha_t \lambda^A(p_t) U(q^A(p_t)) + \beta_t \lambda^B(p_t) U(q^B(p_t)) - c] dt + e^{-rT} P_T(p_0, (\alpha_\tau)) U(p_T), \quad (\text{A.1})$$

$$\text{s.t. } \dot{p}_t = -(\alpha_t - \beta_t) \delta p_t (1 - p_t), \quad (\text{A.2})$$

where $P_t(p_0, (\alpha_\tau)) := p_0 e^{-(\bar{\lambda} \int_0^t \alpha_s ds + \underline{\lambda} \int_0^t \beta_s ds)} + (1 - p_0) e^{-(\underline{\lambda} \int_0^t \alpha_s ds + \bar{\lambda} \int_0^t \beta_s ds)}$, and $\beta_t = 1 - \alpha_t$. Given that the problem is autonomous, the optimal α_t only depends on the belief at time t . Similarly, the decision to stop and take an action only depends on p_t .

The HJB equation for this problem is a variational inequality

$$\max \left\{ -c - rV(p) + \max_{\alpha \in [0,1]} F_\alpha(p, V(p), V'(p)), U(p) - V(p) \right\} = 0, \quad (\text{A.3})$$

where

$$F_\alpha(p, V(p), V'(p)) := \left\{ \begin{array}{l} \alpha \lambda^A(p) (U(q^A(p)) - V(p)) + (1 - \alpha) \lambda^B(p) (U(q^B(p)) - V(p)) \\ - (2\alpha - 1) \delta p (1 - p) V'(p) \end{array} \right\}. \quad (\text{A.4})$$

In the “experimentation region” where $V(p) > U(p)$ the HJB equation reduces to

$$c + rV(p) = F(p, V(p), V'(p)) \left(:= \max_{\alpha \in [0,1]} F_\alpha(p, V(p), V'(p)) \right). \quad (\text{A.5})$$

If $V(p) = U(p)$, then $T(p) = 0$ is optimal and we must have $c + rV(p) \geq F(p, V(p), V'(p))$.

In the following, we will construct a candidate value function that satisfies (A.3) for all points of differentiability. This would be sufficient if the candidate were differentiable everywhere. Since our candidate function has kinks, we show instead that it is a *viscosity solution* of (A.3), a necessary and sufficient for condition the value function according to the verification theorem we invoke (see Proposition 7).

Note that $F_\alpha(\cdot)$ is linear in α . Therefore, the optimal policy is a bang-bang solution and we have $\alpha^*(p) \in \{0, 1\}$ except for posteriors where the derivative of the objective vanishes. Whenever it is differentiable, a candidate value function must then solve one of the two ODEs: $c + rV_0(p) = F_0(p, V_0(p), V_0'(p))$ and $c + rV_1(p) = F_1(p, V_1(p), V_1'(p))$.

For the construction of the candidate value function, we shall assume

$$U(q^A(p)) = U_a(q^A(p)) \quad \text{and} \quad U(q^B(p)) = U_b(q^B(p)), \quad (\text{C})$$

meaning the DM finds it optimal to choose actions a and b when receiving an A -signal and a B -signal, respectively. This assumption will be justified in Proposition 7 and in Section A.3, where we present a sufficient condition for (C) to hold at the optimal policy.

Given Condition (C), the ODEs above can be rewritten as:

$$c + rV_0(p) = \lambda^B(p) (U_b(q^B(p)) - V_0(p)) + \delta p (1 - p) V_0'(p), \quad (\text{A.6})$$

$$c + rV_1(p) = \lambda^A(p) (U_a(q^A(p)) - V_1(p)) - \delta p (1 - p) V_1'(p). \quad (\text{A.7})$$

Solutions to these ODEs with boundary condition $V(x) = W$ are well-defined if $x \in (0, 1)$, and denoted by $V_0(p; x, W)$ and $V_1(p; x, W)$, respectively.⁵⁰

⁵⁰ $V_0(p; x, W)$ and $V_1(p; x, W)$ are uniquely defined if $x \in (0, 1)$ because (A.6) and (A.7) satisfy local Lipschitz continuity for all $p \in (0, 1)$.

A.2 Two Benchmarks and a Condition for Experimentation

We define two benchmark functions:

$$\bar{U}(p) := \frac{\frac{1}{2}\lambda^A(p)U_a(q^A(p)) + \frac{1}{2}\lambda^B(p)U_b(q^B(p)) - c}{r + \frac{1}{2}\bar{\lambda} + \frac{1}{2}\underline{\lambda}}, \quad (\text{A.8})$$

$$\hat{U}(p) := \frac{\lambda^A(p)U_a(q^A(p)) + \lambda^B(p)U_b(q^B(p)) - c}{r + \bar{\lambda} + \underline{\lambda}}. \quad (\text{A.9})$$

While these functions are defined for all $p \in [0, 1]$, more intuitive explanations are given in case (C) holds for p . In that case, $\bar{U}(p)$ corresponds to the value of the stationary strategy described in the main text, where $\alpha_t = \beta_t = 1/2$ for all t and the DM takes an optimal action after receiving a signal. Meanwhile, $\hat{U}(p)$ corresponds to the value of a hypothetical “full attention strategy,” where $\alpha_t = \beta_t = 1$ for all t and the DM takes an optimal action after receiving a signal. Since limited attention prevents the DM from achieving this value in our model, $\hat{U}(p)$ serves only as an analytical device for proofs.

We make several observations that will be used later.

Lemma 1. (a) If Assumption 1 is satisfied, $\bar{U}(p) < \hat{U}(p)$ for all $p \in [0, 1]$.

(b) $\bar{U}(p)$ and $\hat{U}(p)$ are linear in p , $U'_b(p) < \bar{U}'(p) < U'_a(p)$, and $U'_b(p) < \hat{U}'(p) < U'_a(p)$.

(c) $\bar{U}(p), \hat{U}(p) < U(p)$ at $p \in \{0, 1\}$; and for all $p \in [0, 1]$, $\bar{U}(p)$ and $\hat{U}(p)$ are strictly decreasing without bound in c .

Proof. Parts (b) and (c) are obtained from straightforward derivations. For (a) suppose first that $r = 0$. In this case $\bar{U}(p) = \hat{U}(p) - c/(\bar{\lambda} + \underline{\lambda}) < \hat{U}(p)$. If $r > 0$, simple algebra shows that $\bar{U}(0) < \hat{U}(0)$ and $\bar{U}(1) < \hat{U}(1)$ are equivalent to the inequalities in Assumption 1. The linearity of \bar{U} and \hat{U} implies $\bar{U}(p) < \hat{U}(p)$ for all $p \in [0, 1]$. \square

Intuitively, $\hat{U}(p)$ is the most the DM can get from experimentation.⁵¹ By Lemma 1, therefore, experimentation can be optimal for some p only if

$$\hat{U}(\hat{p}) > U(\hat{p}). \quad (\text{EXP})$$

We will show in Proposition 7 that experimentation is optimal for some p if and only if (EXP) is satisfied.

A.3 Contradictory Strategy

Recall the structure of the contradictory strategy given by (3.1). For our purpose, it is convenient to define the contradictory strategy to consist of a “left policy” which prescribes an immediate action b for $p \leq \underline{p}^*$ and an attention strategy $\alpha = 1$ for $p > \underline{p}^*$, and a “right

⁵¹Indeed, Lemma 8 in Appendix B.4 in the Supplemental Material shows that $\max\{U(p), \hat{U}(p)\}$ is an upper bound for the value function if SEP holds.

policy” that prescribes an immediate action a for $p \geq \bar{p}^*$ and an attention strategy $\alpha = 0$ for $p < \bar{p}^*$, where \underline{p}^* and \bar{p}^* are stopping boundaries. We postulate, and later verify, that the stopping boundaries satisfy value matching: $V(\underline{p}^*) = U_b(\underline{p}^*)$ and $V(\bar{p}^*) = U_a(\bar{p}^*)$; and smooth pasting: $V'(\underline{p}^*) = U'_b(\underline{p}^*)$ and $V'(\bar{p}^*) = U'_a(\bar{p}^*)$. These two conditions, together with (A.6) and (A.7), pin down the boundaries (we omit the algebra):

$$\underline{p}^* = \frac{(u_b^B - u_a^B) \underline{\lambda} + u_b^B r + c}{r(u_b^B - u_b^A) + (u_a^A - u_b^A) \bar{\lambda} + (u_b^B - u_a^B) \underline{\lambda}}, \quad (\text{A.10})$$

$$\bar{p}^* = \frac{(u_b^B - u_a^B) \bar{\lambda} - u_a^B r - c}{r(u_a^A - u_b^A) + (u_b^B - u_a^B) \bar{\lambda} + (u_a^A - u_b^A) \underline{\lambda}}, \quad (\text{A.11})$$

whenever these values are within $[0, 1]$. Otherwise, we set $\underline{p}^* := 1$ and $\bar{p}^* := 0$.

With aid of the solutions to the ODEs (A.6) and (A.7), we can then define the value of the left contradictory policy:

$$\underline{V}_{ct}(p) := \begin{cases} U_b(p), & \text{if } p \leq \underline{p}^*, \\ V_1(p; \underline{p}^*, U_b(\underline{p}^*)), & \text{if } p > \underline{p}^*, \end{cases}$$

and the value of the right contradictory policy:

$$\bar{V}_{ct}(p) := \begin{cases} U_a(p), & \text{if } p \geq \bar{p}^*, \\ V_0(p; \bar{p}^*, U_a(\bar{p}^*)), & \text{if } p < \bar{p}^*. \end{cases}$$

Expressions for $\underline{V}_{ct}(p)$ and $\bar{V}_{ct}(p)$ can be found in Appendix B.1 in the Supplemental Material. Combining these functions, we define the value of the contradictory strategy as $V_{ct}(p) := \max \{ \underline{V}_{ct}(p), \bar{V}_{ct}(p) \}$. Without further analysis, it is not clear when $V_{ct}(p)$ is the value of a strategy of the form (3.1). This will be clarified in Proposition 6.

To derive some properties of \underline{V}_{ct} and \bar{V}_{ct} , we first state a crossing condition with \hat{U} .

Lemma 2. *If $V_1(p)$ satisfies (A.7) and $V_1(p) = \hat{U}(p)$ at some $p \in (0, 1)$ then*

$$q^B(p) < (=) \underline{p}^* \implies V'_1(p) < (=) \hat{U}'(p). \quad (\text{A.12})$$

Likewise, if $V_0(p)$ satisfies (A.6) and $V_0(p) = \hat{U}(p)$ at some $p \in (0, 1)$ then

$$q^A(p) > (=) \bar{p}^* \implies V'_0(p) > (=) \hat{U}'(p). \quad (\text{A.13})$$

The proof follows from straightforward algebra. Note that $q^B(p) \leq \underline{p}^*$ and $q^A(p) \geq \bar{p}^*$ are necessary to ensure that observing a signal causes the DM’s posterior belief to “land” in the stopping region. Therefore, it is necessary for SEP that these conditions hold for

$p \in [\underline{p}^*, \bar{p}^*]$. This is guaranteed by the following condition:⁵²

$$q^A(\underline{p}^*) \geq \bar{p}^* \quad \text{and} \quad q^B(\bar{p}^*) \leq \underline{p}^*. \quad (\text{N-SEP})$$

Note also that (N-SEP) implies that condition (C) holds for all $p \in [\underline{p}^*, \bar{p}^*]$.

The following lemma summarizes the properties of the contradictory strategy:

- Lemma 3.** (a) $\underline{V}_{ct}(p)$ and $\bar{V}_{ct}(p)$ are continuously differentiable and convex on $(0, 1)$;
(b) $\underline{V}_{ct}(p)$ is strictly convex and $\underline{V}_{ct}(p) > U_b(p)$ on $(\underline{p}^*, 1]$, and $\bar{V}_{ct}(p)$ is strictly convex and $\bar{V}_{ct}(p) > U_a(p)$ on $[0, \bar{p}^*)$. $V_{ct}(p) > U(p)$ for $p \in (\underline{p}^*, \bar{p}^*)$.
(c) If $\underline{p}^*, \bar{p}^* \in (0, 1)$, they satisfy

$$U_b(\underline{p}^*) = \widehat{U}(\underline{p}^*), \quad \text{and} \quad U_a(\bar{p}^*) = \widehat{U}(\bar{p}^*). \quad (\text{A.14})$$

- (d) Suppose (EXP) and (N-SEP) holds. Then, $\underline{V}_{ct}(p) < \widehat{U}(p)$ for $p \in (\underline{p}^*, \bar{p}^*]$, $\bar{V}_{ct}(p) < \widehat{U}(p)$ for $p \in [\underline{p}^*, \bar{p}^*)$, and $V_{ct}(p) = U(p) > \widehat{U}(p)$ for $p \notin [\underline{p}^*, \bar{p}^*]$.

Proof. Parts (a)-(c) follow from straightforward algebra. Next we prove the first statement in Part (d). Suppose by contradiction that $\underline{V}_{ct}(p) \geq \widehat{U}(p)$ for some $p \in (\underline{p}^*, \bar{p}^*]$. By smooth pasting and Lemma 1.(b), $\underline{V}'_{ct}(\underline{p}^*) = U'_b(\underline{p}^*) < \widehat{U}'(\underline{p}^*)$; and by value matching and Lemma 3.(c), $\underline{V}_{ct}(\underline{p}^*) = U_b(\underline{p}^*) = \widehat{U}(\underline{p}^*)$. This implies that \bar{V}_{ct} crosses \widehat{U} from above at \underline{p}^* . Therefore, $\underline{V}_{ct}(p) \geq \widehat{U}(p)$ for some $p \in (\underline{p}^*, \bar{p}^*]$ implies that there exists $p' \in (\underline{p}^*, \bar{p}^*]$ such that $\underline{V}_{ct}(p') = \widehat{U}(p')$ and $\underline{V}'_{ct}(p') > \widehat{U}'(p')$, where the strict inequality follows from strict convexity of \underline{V}_{ct} . Note however, that $\underline{V}_{ct}(p') = V_1(p')$ satisfies (A.7) and since (N-SEP) implies $q^B(p') \leq \underline{p}^*$, (A.12) in Lemma 2 implies that $\underline{V}'_{ct}(p') \leq \widehat{U}'(p')$ which is a contradiction. We thus conclude that $\underline{V}_{ct}(p) < \widehat{U}(p)$ for $p \in (\underline{p}^*, \bar{p}^*]$. The second statement follows by a symmetric argument. For the last statement, consider $p \notin [\underline{p}^*, \bar{p}^*]$. By construction, $V_{ct}(p) = U(p)$ and by Lemma 1, $\widehat{U}(p) < U(p)$. \square

Lemma 1 implies that there are two points $p \in (0, 1)$ where $\widehat{U}(p)$ and $U(p)$ intersect if (EXP) is satisfied. By (A.14), the intersection points are given by \underline{p}^* and \bar{p}^* . This means that, given (EXP), we must have $0 < \underline{p}^* < \hat{p} < \bar{p}^* < 1$. If (EXP) does not hold, $\widehat{U}(p)$ and $U(p)$ do not intersect, $\underline{p}^* > \bar{p}^*$, and we have $V_{ct}(p) = U(p)$ for all $p \in [0, 1]$.

A.4 Confirmatory Strategy

The confirmatory strategy prescribes $\alpha = 1$ for $p > p^*$, $\alpha = 0$ for $p < p^*$ and $\alpha = 1/2$ at p^* , for some belief $p^* \in (0, 1)$. Therefore, the value of this strategy, denoted by V_{cf} , must satisfy (A.7) for all $p > p^*$ and (A.6) for all $p < p^*$. Moreover, we must have

⁵²It is straightforward to verify that $q^A(\underline{p}^*) \geq \bar{p}^*$ if and only if $q^B(\bar{p}^*) \leq \underline{p}^*$. We state (N-SEP) in a redundant way since it is more intuitive.

$V_{cf}(p^*) = \bar{U}(p^*)$. Therefore we set

$$V_{cf}(p) := \begin{cases} \underline{V}_{cf}(p) := V_0(p; p^*, \bar{U}(p^*)), & \text{if } p \leq p^*, \\ \bar{V}_{cf}(p) := V_1(p; p^*, \bar{U}(p^*)), & \text{if } p > p^*. \end{cases} \quad (\text{A.15})$$

To pin down p^* , we assume smooth pasting at p^* and $V'_{cf}(p^*) = \bar{U}'(p^*)$. Inserting $V_0(p^*) = \bar{U}(p^*)$ and $V'_0(p^*) = \bar{U}'(p^*)$ in (A.6) yields⁵³

$$p^* = \frac{(u_b^B - u_a^B) \bar{\lambda} \underline{\lambda} + r (u_b^B \bar{\lambda} - u_a^B \underline{\lambda}) + c (\bar{\lambda} - \underline{\lambda})}{(u_a^A - u_b^A) \bar{\lambda} \underline{\lambda} + r (u_a^A \bar{\lambda} - u_b^A \underline{\lambda}) + (u_b^B - u_a^B) \underline{\lambda} \bar{\lambda} + r (u_b^B \bar{\lambda} - u_a^B \underline{\lambda}) + 2c (\bar{\lambda} - \underline{\lambda})}. \quad (\text{A.16})$$

Explicit expressions for $\underline{V}_{cf}(p)$ and $\bar{V}_{cf}(p)$ can be found in Appendix B.1 in the Supplemental Material. We establish several properties of $V_{cf}(p)$ that will be used later.

Lemma 4. (a) $V_{cf}(p)$ is continuously differentiable and strictly convex on $(0, 1)$, and

$$V_{cf}(p) \geq \bar{U}(p) \text{ for all } p \in [0, 1] \text{ with strict inequality for } p \neq p^*.$$

(b) Suppose (N-SEP) holds. Then, $V_{cf}(p) < \max\{\widehat{U}(p), U(p)\}$ for all $p \in [0, 1]$.

Proof. Part (a) follows from straightforward algebra. For part (b), we first prove that $V_{cf}(p) < \widehat{U}(p)$ for $p \in [p^*, \bar{p}^*]$ assuming (EXP). Since $V_{cf}(p) = V_1(p)$, and $V_{cf}(p^*) = \bar{U}(p^*) < \widehat{U}(p^*)$, we can use Lemma 2 in a similar way as in Lemma 3.(d) to show that $V_{cf}(p) < \widehat{U}(p)$ for $p \in [p^*, \bar{p}^*]$. A symmetric argument establishes $V_{cf}(p) < \widehat{U}(p)$ for $p \in [p^*, p^*]$ assuming (EXP). Consider next $V_{cf}(p)$ for $p \in [0, p']$, where $p' := \min\{p^*, \underline{p}^*\}$. Observe $V_{cf}(0) < U_b(0) = U(0)$ and $V_{cf}(p') < \widehat{U}(p') \leq U(p')$. The latter was shown above if $p' = \underline{p}^*$ and (EXP) holds. If $p' = p^* < \underline{p}^*$ (which is the case if (EXP) fails), we have $V_{cf}(p') = \bar{U}(p') < \widehat{U}(p') < U(p')$. Since $V_{cf}(\cdot)$ is strictly convex and since $U(\cdot) = U_b(\cdot)$ is linear in that range, we have $V_{cf}(p) < U(p)$ for $p \in [0, p']$. A symmetric argument shows $V_{cf}(p) < U(p)$ for $p \in [\max\{p^*, \bar{p}^*\}, 1]$. Combining the results yields (b). \square

A.5 Solution Candidate

Our solution candidate is the upper envelope of $V_{ct}(p)$ and $V_{cf}(p)$, denoted by $V_{Env}(p) := \max\{V_{ct}(p), V_{cf}(p)\}$. Towards showing that this candidate is the optimal value function of the problem (A.1), we characterize its structure in Proposition 6. To this end, we first establish the following lemma.

Lemma 5 (Crossing Lemma). Let V_0 satisfy (A.6) and V_1 satisfy (A.7). If $V_0(p) = V_1(p) > (=)\bar{U}(p)$ at some $p \in (0, 1)$, then $V'_0(p) > (=)V'_1(p)$.

Proof. Suppose $V_0(p) = V_1(p) = V(p)$ for some $p \in (0, 1)$. Solving (A.6) and (A.7) for $V'_0(p)$ and $V'_1(p)$ and some algebra yields $p(1-p)\delta(V'_0(p) - V'_1(p)) / (2r + \bar{\lambda} + \underline{\lambda}) = V(p) - \bar{U}(p)$. Therefore $\text{sgn}(V'_0(p) - V'_1(p)) = \text{sgn}(V(p) - \bar{U}(p))$. \square

⁵³The same expression is obtained from inserting $V_1(p^*) = \bar{U}(p^*)$ and $V'_1(p^*) = \bar{U}'(p^*)$ in (A.7).

Proposition 6 (Structure of V_{Env}). Suppose (N-SEP) holds.

(a) If (EXP) holds and $V_{ct}(p^*) \geq V_{cf}(p^*)$, then there exists a unique $\check{p} \in (\underline{p}^*, \bar{p}^*)$ such that $\underline{V}_{ct}(\check{p}) = \bar{V}_{ct}(\check{p})$ and

$$V_{Env}(p) = V_{ct}(p) = \begin{cases} \underline{V}_{ct}(p), & \text{if } p < \check{p}, \\ \bar{V}_{ct}(p), & \text{if } p \geq \check{p}. \end{cases}$$

(b) If (EXP) holds and $V_{ct}(p^*) < V_{cf}(p^*)$, then $p^* \in (\underline{p}^*, \bar{p}^*)$, and there exist a unique $\underline{p} \in (\underline{p}^*, p^*)$ such that $V_{ct}(\underline{p}) = V_{cf}(\underline{p})$, and a unique $\bar{p} \in (p^*, \bar{p}^*)$ such that $V_{ct}(\bar{p}) = V_{cf}(\bar{p})$ and

$$V_{Env}(p) = \begin{cases} \underline{V}_{ct}(p), & \text{if } p < \bar{p}, \\ V_{cf}(p), & \text{if } p \in [\underline{p}, \bar{p}], \\ \bar{V}_{ct}(p), & \text{if } p > \bar{p}. \end{cases}$$

(c) If (EXP) is violated, $V_{Env}(p) = U(p)$ for all $p \in [0, 1]$.

Proof. Part (a): We first prove that $V_{ct}(p) \geq V_{cf}(p)$ for all $p \in [0, 1]$. Since $V_{ct}(p) \geq U(p) > V_{cf}(p)$ for $p \notin [\underline{p}^*, \bar{p}^*]$, it suffices to show $V_{ct}(p) \geq V_{cf}(p)$ for $p \in [\underline{p}^*, \bar{p}^*]$. To this end, suppose first $p^* > \underline{p}^*$. Recall from Lemmas 3 and 4 that $\underline{V}_{ct}(\underline{p}^*) = \widehat{U}(\underline{p}^*) > V_{cf}(\underline{p}^*)$. Since $V_{cf}(\cdot) \geq \bar{U}(\cdot)$, by the Crossing Lemma 5, \underline{V}_{ct} can cross V_{cf} only from above on $[\underline{p}^*, p^*)$. If $\underline{V}_{ct}(p^*) > V_{cf}(p^*)$ this implies that $V_{cf}(p) < \underline{V}_{ct}(p) \leq V_{ct}(p)$ for all $p \in [\underline{p}^*, p^*)$. If $\underline{V}_{ct}(p^*) = V_{cf}(p^*)$ then we obtain $\underline{V}_{ct}''(p^*) > V_{cf}''(p^*)$, which implies $\underline{V}_{ct}(p) > V_{cf}(p)$ for $p < p^*$ in a neighborhood of p^* ; again the fact that $\underline{V}_{ct}(p)$ can cross $V_{cf}(p)$ only from above implies that $V_{cf}(p) < \underline{V}_{ct}(p) \leq V_{ct}(p)$ for all $p \in [\underline{p}^*, p^*)$. If $\underline{V}_{ct}(p^*) < V_{cf}(p^*)$, then $\bar{V}_{ct}(p^*) = V_{ct}(p^*) \geq V_{cf}(p^*)$. Since both $\bar{V}_{ct}(p)$ and $V_{cf}(p^*)$ satisfy (A.6), we must have $V_{cf}(p) \leq \bar{V}_{ct}(p) \leq V_{ct}(p)$ for all $p \in [\underline{p}^*, p^*]$. Either way, we have proven that $V_{cf}(p) \leq V_{ct}(p)$ for all $p \in [\underline{p}^*, p^*]$. A symmetric argument proves that $V_{cf}(p) \leq V_{ct}(p)$ for all $p \in [p^*, \bar{p}^*]$ in case $p^* < \bar{p}^*$.

We have now proven that $V_{ct}(p) \geq V_{cf}(p)$ for all $p \in [0, 1]$. Recall from Lemma 3 that $\underline{V}_{ct}(\underline{p}^*) = \widehat{U}(\underline{p}^*) > \bar{V}_{ct}(\underline{p}^*)$ and $\bar{V}_{ct}(\bar{p}^*) = \widehat{U}(\bar{p}^*) > \underline{V}_{ct}(\bar{p}^*)$. By the intermediate value theorem, there exists $\check{p} \in (\underline{p}^*, \bar{p}^*)$ where $\underline{V}_{ct}(\check{p}) = \bar{V}_{ct}(\check{p})$. Since $V_{cf}(\cdot) \geq \bar{U}(\cdot)$, and $V_{ct}(p) \geq V_{cf}(p)$ we have $\underline{V}_{ct}(\check{p}) = \bar{V}_{ct}(\check{p}) \geq \bar{U}(\check{p})$, with strict inequality for any $p \neq p^*$. Hence, by the Crossing Lemma 5, \underline{V}_{ct} must cross \bar{V}_{ct} from above at \check{p} . This means that the intersection point \check{p} is unique and the structure stated in part (a) obtains.

Part (b): We first prove that $p^* \in (\underline{p}^*, \bar{p}^*)$. By Lemma 3, $V_{ct}(p) \geq U(p)$ for all $p \in [0, 1]$. This implies $V_{cf}(p^*) > U(p^*)$, and since $V_{cf}(p^*) = \bar{U}(p^*) < \widehat{U}(p^*)$, and since by Lemma 3.(d) $\widehat{U}(p) \leq U(p)$ for $p \notin (\underline{p}^*, \bar{p}^*)$, we must have $p^* \in (\underline{p}^*, \bar{p}^*)$. Next, by Lemma 4.(b), $V_{cf}(\underline{p}^*) < \widehat{U}(\underline{p}^*) = \underline{V}_{ct}(\underline{p}^*)$. Therefore, $V_{cf}(p)$ and $\underline{V}_{ct}(p)$ intersect at some $\underline{p} \in (\underline{p}^*, p^*)$ and by the Crossing Lemma 5, the intersection is unique. Moreover, for $p < p^*$, we have $V_{cf}(p) > \bar{V}_{ct}(p)$ since both satisfy (A.6), and hence $\bar{V}_{ct}(p) < \underline{V}_{ct}(p)$ for

all $p \in (\underline{p}^*, \underline{p})$. This proves the result for $p \leq \underline{p}^*$. For $p > \underline{p}^*$ the arguments are symmetric.

Part (c): If **(EXP)** fails, then $V_{ct}(p) = U(p) \geq \widehat{U}(p)$ for all $p \in [0, 1]$. By Lemma 4.(b), $V_{cf}(p) < \max\{U(p), \widehat{U}(p)\} = U(p)$. Hence, $V_{Env}(p) = U(p)$ for all $p \in [0, 1]$. \square

A.6 Verification of the Candidate

We now show that V_{Env} is the value function of the DM's problem in (A.1). We need two additional lemmas for this result.

Lemma 6 (Unimprovability). *Suppose **(N-SEP)** and **(EXP)** are satisfied.*

- (a) *If V_0 satisfies (A.6) and $V_0(p) \geq \bar{U}(p)$ for some $p \in [\underline{p}^*, \bar{p}^*]$, then V_0 satisfies (A.5) at p , and $\alpha = 0$ is a maximizer.*
- (b) *If V_1 satisfies (A.7) and $V_1(p) \geq \bar{U}(p)$ for some $p \in [\underline{p}^*, \bar{p}^*]$, then V_1 satisfies (A.5) at p , and $\alpha = 1$ is a maximizer.*

The maximizers are unique if $V.(p) > \bar{U}(p)$.

Proof. Consider first the case that $V_0(p)$ satisfies (A.6). With $V = V_0(p)$, and substituting $V' = V'_0(p)$ from (A.6), we have

$$\frac{\partial F_\alpha(p, V_0(p), V'_0(p))}{\partial \alpha} = (2r + \bar{\lambda} + \underline{\lambda}) (\bar{U}(p) - V_0(p)).$$

where we have used **(C)** which holds for $p \in [\underline{p}^*, \bar{p}^*]$ if **(N-SEP)** is satisfied. This implies that $\alpha = 0$ is a maximizer if $V_0(p) \geq \bar{U}(p)$, and the unique maximizer if the inequality is strict. This proves Part (a). The proof of Part (b) follows from a similar argument. \square

The next lemma shows that under Assumption 1, the LHS of the HJB equation is positive whenever $V_{Env}(p) = U(p)$.

Lemma 7. *If Assumption 1 is satisfied, then $c + rU_b(p) \geq 0$ for $p \leq \underline{p}^*$, and $c + rU_a(p) \geq 0$ for $p \geq \bar{p}^*$.*

Proof. We show $c + rU_b(p) \geq 0$. If $u_b^A \geq 0$, we have $U_b(p) \geq 0$ for all p so the result follows. Next consider the case that $u_b^A < 0$. In this case $U'_b(p) < 0$ and hence it is sufficient to show $rU_b(\underline{p}^*) \geq 0$. For $r > 0$ this can be rearranged to

$$u_b^B (u_a^A \bar{\lambda} + u_b^A \underline{\lambda}) - u_b^A (u_b^B \bar{\lambda} + u_a^B \underline{\lambda}) \geq 0,$$

which holds under Assumption 1 if $u_b^A < 0$. $c + rU_a(p) \geq 0$ is obtained similarly. \square

We are now ready to verify the optimality of our candidate value function.

Proposition 7. *Given **(N-SEP)** and Assumption 1, $V^*(p) = V_{Env}(p)$ for all $p \in [0, 1]$.*

Proof. In this proof we write $V(p) = V_{Env}(p)$. Theorem III.4.11 in [Bardi and Capuzzo-Dolcetta \(1997\)](#) characterizes the value function of a dynamic programming problem with an optimal stopping decision as in (A.1) as the (unique) viscosity solution of the HJB equation.⁵⁴ For all $p \in (0, 1)$ where $V(p)$ is differentiable, this requires that $V(p)$ satisfy (A.3). We start by considering points of differentiability where $V(p) = U(p)$, i.e., $p \notin (\underline{p}^*, \bar{p}^*)$.⁵⁵ First consider $p < \hat{p}$, which implies $p \leq \underline{p}^*$. If $q^A(p) \geq \hat{p}$, some algebra shows that $F_1(p, U_b(p), U'_b(p)) \leq c + rU_b(p)$ if and only if $p \leq \underline{p}^*$. If instead $q^A(p) < \hat{p}$, we have $F_1(p, U_b(p), U'_b(p)) = 0 \leq c + rU_b(p)$ (where the inequality follows from Lemma 7). Also, if $p < \hat{p}$, $F_0(p, U_b(p), U'_b(p)) = 0 \leq c + rU_b(p)$. Hence (A.3) holds for all $p < \hat{p}$, $p \notin (\underline{p}^*, \bar{p}^*)$. A similar argument shows that (A.3) holds for all $p > \hat{p}$, $p \notin (\underline{p}^*, \bar{p}^*)$. We have thus shown that $V(p)(= U(p))$ satisfies (A.3) for all $p \notin (\underline{p}^*, \bar{p}^*)$ where it is differentiable. For the case that (EXP) is violated, this implies that $V(p)$ satisfies (A.3) for all $p \neq \hat{p}$.

Next, we consider points of differentiability $p \in (\underline{p}^*, \bar{p}^*)$. By (A.14), $\underline{p}^* < \bar{p}^*$ if and only if (EXP) is satisfied. Lemma 3 implies that $V_{ct}(p) > U(p)$ for all $p \in (\underline{p}^*, \bar{p}^*)$ and hence $V(p) > U(p)$. This implies that (A.3) is equivalent to (A.5) for all $p \in (\underline{p}^*, \bar{p}^*)$. Since $V(p)$ satisfies (A.6) or (A.7) at points of differentiability, and $V(p) \geq V_{cf}(p) \geq \bar{U}(p)$, the Unimprovability Lemma 6 implies that $V(p)$ satisfies (A.5).

We have shown that $V(p)$ satisfies (A.3) for all points of differentiability. For $V(p)$ to be a viscosity solution it remains to show that for all points of non-differentiability,

$$\max \{-c - rV(p) + F(p, V(p), \rho), U(p) - V(p)\} \leq 0, \quad (\text{A.17})$$

for all $\rho \in [V'_-(p), V'_+(p)]$; and the opposite inequality holds for all $\rho \in [V'_+(p), V'_-(p)]$. By Proposition 6, non-differentiability arises at \hat{p} if (EXP) fails; at \check{p} if (EXP) holds and $V_{ct}(p^*) \geq V_{cf}(p^*)$; and at \underline{p} and \bar{p} if (EXP) holds and $V_{ct}(p^*) < V_{cf}(p^*)$. Since $V(p) \geq \bar{U}(p)$, the Crossing Lemma 5 implies that $V(p)$ has convex kinks at all these points so that $V'_-(p) \leq V'_+(p)$. Therefore it suffices to check (A.17) for all $\rho \in [V'_-(p), V'_+(p)]$. F_α is linear in α (see (A.4)), so it suffices to consider $\alpha \in \{0, 1\}$. For $\alpha = 1$ we have $F_1(p, V(p), \rho) \leq F_1(p, V(p), V'_-(p))$ and for $\alpha = 0$ we have $F_0(p, V(p), \rho) \leq F_0(p, V(p), V'_+(p))$. Therefore if $U(p) \leq V(p)$, which holds for our candidate solution by construction,

$$c + rV(p) \geq \max \{F_1(p, V(p), V'_-(p)), F_0(p, V(p), V'_+(p))\} \quad (\text{A.18})$$

implies that (A.17) holds for all for $\rho \in [V'_-(p), V'_+(p)]$. We distinguish three cases.

Case A: (EXP) is violated. In this case we only have to consider $p = \hat{p}$. We have $V(\hat{p}) = U_a(\hat{p}) = U_b(\hat{p})$, $V'_-(\hat{p}) = U'_b(\hat{p})$ and $V'_+(\hat{p}) = U'_a(\hat{p})$. Now suppose by contradiction

⁵⁴To formally apply their theorem, we have to use P_t as a second state-variable and define a value function $v(p, P) = PV(p)$. Since v is continuously differentiable in P , it is straightforward to apply the result directly to $V(p)$.

⁵⁵If (EXP) is violated so that $\underline{p}^* > \bar{p}^*$, we have $(\underline{p}^*, \bar{p}^*) = \emptyset$ so that $p \notin (\underline{p}^*, \bar{p}^*)$ if and only if $p \in [0, 1]$.

that (A.18) is violated at \hat{p} . Suppose that $c + rU_b(\hat{p}) < F_1(\hat{p}, U_b(\hat{p}), U'_b(\hat{p}))$ (the other case is similar). Some algebra shows that this implies $\underline{p}^* < \hat{p}$ which by (A.14) contradicts that (EXP) is violated. Therefore $V(p)$ is a viscosity solution of (A.3) if (EXP) is violated.

Case B: (EXP) is satisfied and $V_{ct}(p^*) \geq V_{cf}(p^*)$. Consider $p = \check{p}$. (A.18) becomes

$$c + r\underline{V}_{ct}(\check{p}) = c + r\overline{V}_{ct}(\check{p}) \geq \max \left\{ F_1(\check{p}, \underline{V}_{ct}(\check{p}), \underline{V}'_{ct}(\check{p})), F_0(\check{p}, \overline{V}_{ct}(\check{p}), \overline{V}'_{ct}(\check{p})) \right\}.$$

By the Unimprovability Lemma 6, this holds with equality since $\underline{V}_{ct}(p)$ satisfies (A.7), $\overline{V}_{ct}(p)$ satisfies (A.6) at \check{p} , and Condition (C) holds at \check{p} . Lemma 3 implies that $V_{ct}(p) > U(p)$ for all $p \in (\underline{p}^*, \overline{p}^*)$ and hence $V(\check{p}) > U(\check{p})$. (A.17) is thus satisfied at \check{p} .

Case C: (EXP) is satisfied and $V_{ct}(p^*) < V_{cf}(p^*)$. The proof is similar to Case B.

We have thus shown that $V(p)$ is a viscosity solution of (A.3) which is sufficient for $V(p)$ to be the value function of problem (A.1). \square

A.7 Proof of Theorem 1

We show that Theorem 1 holds with the cutoffs (where $a \vee b = \max\{a, b\}$ and $a \wedge b = \min\{a, b\}$):

$$\begin{aligned} \bar{c} &:= 0 \vee \frac{\lambda (u_a^A - u_b^A) (u_b^B - u_a^B) - r (u_a^A u_b^B - u_b^A u_a^B)}{(u_a^A + u_b^B) - (u_b^A + u_a^B)}, & (\text{A.19}) \\ \underline{c} &:= 0 \vee \begin{cases} \bar{c} \wedge \min \left\{ \frac{r+\lambda}{1+(\frac{2r+\lambda}{\lambda})^{\lambda/r}} (u_a^A - u_b^A) - ru_a^A, \frac{r+\lambda}{1+(\frac{2r+\lambda}{\lambda})^{\lambda/r}} (u_b^B - u_a^B) - ru_b^B \right\} & \text{if } r > 0, \\ \bar{c} \wedge \lambda \min \left\{ (u_a^A - u_b^A), (u_b^B - u_a^B) \right\} / (1 + e^2) & \text{if } r = 0. \end{cases} & (\text{A.20}) \end{aligned}$$

Proof of Theorem 1. In the baseline model $\underline{\lambda} = 0$. Hence (N-SEP) holds and (EXP) simplifies to

$$c ((u_a^A + u_b^B) - (u_b^A + u_a^B)) + r (u_a^A u_b^B - u_b^A u_a^B) < \lambda (u_a^A - u_b^A) (u_b^B - u_a^B).$$

By Propositions 6 and 7, no experimentation is optimal if (EXP) is violated, which holds if and only if $c \geq \bar{c}$. This proves part (a).

Conversely if $c \leq \bar{c}$, (EXP) is satisfied. Next we show that

$$c \geq \underline{c} \iff \max \{ \underline{V}_{ct}(p^*), \overline{V}_{ct}(p^*) \} \geq \overline{U}(p^*). \quad (\text{A.21})$$

By Proposition 6 and 7, and the Crossing Lemma 6, this implies that the policies stated in Parts (b) and (c) of Theorem 1 are optimal.

We first assume $r > 0$. The closed-form solutions for $\underline{V}_{ct}(p^*)$ and $\overline{V}_{ct}(p^*)$ (see Appendix

B.1 in the Supplemental Material), can be used to show that

$$\begin{aligned} & \max \{ \underline{V}_{ct}(p^*), \overline{V}_{ct}(p^*) \} \geq \overline{U}(p^*) \\ \iff & \left(\max \left\{ \frac{c + ru_a^A}{\lambda u_a^A - (r + \lambda)u_b^A - c}, \frac{c + ru_b^B}{\lambda u_b^B - (r + \lambda)u_a^B - c} \right\} \right)^{\frac{r}{\lambda}} \geq \frac{\lambda}{2r + \lambda} \iff c \geq \underline{c}. \end{aligned}$$

This proves (A.21) for $r > 0$. Taking the limit $r \rightarrow 0$ yields the result for $r = 0$.⁵⁶

By definition we have $\bar{c} \geq \underline{c} \geq 0$. It remains to show that the inequalities are strict for r sufficiently small. Fix any (λ, u_x^ω) . We first show that $\bar{c} > 0$ implies $\bar{c} > \underline{c}$. If $c = \bar{c} > 0$, then (EXP) is violated and by Lemma 4.(b), $V_{cf}(p) < U(p)$ for all $p \in [0, 1]$. Lowering c slightly to $c' = \bar{c} - \varepsilon$, $\varepsilon > 0$, by continuity of V_{cf} in c , we still have $V_{cf}(p) < U(p)$ for all $p \in [0, 1]$. Therefore we must have $c' > \underline{c}$. Since $c' < \bar{c}$ by definition, we have shown that $\underline{c} < \bar{c}$. Clearly, $\bar{c} > 0$ if r is sufficiently small, since (EXP) is satisfied if $c = r = 0$. Finally, it follows from (A.20) that $\underline{c} > 0$ for r close to zero, which completes the proof. \square

A.8 Proof of Theorem 2

Proof of Theorem 2. Proposition 8 in Appendix B.4 in the Supplemental Material shows that SEP holds if and only if (N-SEP) is satisfied. (N-SEP) can be rearranged to

$$C_0 + C_1 c - \delta c^2 \leq 0, \quad (\text{A.22})$$

where

$$\begin{aligned} C_0 &= (-\bar{\lambda}u_a^A u_b^B + \underline{\lambda}u_b^A u_a^B) r^2 - 2(u_a^A u_b^B - u_b^A u_a^B) \bar{\lambda} \underline{\lambda} r + \delta \bar{\lambda} \underline{\lambda} (u_a^A - u_b^A) (u_b^B - u_a^B), \\ C_1 &= -r [(u_a^A + u_b^B) \bar{\lambda} - (u_a^B + u_b^A) \underline{\lambda}] - 2((u_a^A + u_b^B) - (u_b^A + u_a^B)) \bar{\lambda} \underline{\lambda} < 0. \end{aligned}$$

Since the LHS of (A.22) is decreasing in c (for $c \geq 0$), there exists a cutoff \underline{c}^{SEP} such that (N-SEP) holds if and only if $c \geq \underline{c}^{SEP}$. If $r > 0$ and $\underline{\lambda} = 0$, C_0 simplifies to $-r^2 \bar{\lambda} u_a^A u_b^B < 0$ so $\underline{c}^{SEP} = 0$ for $\underline{\lambda}$ in a neighborhood of zero if $r > 0$.

Now suppose $c \geq \underline{c}^{SEP}$ so that SEP holds. (EXP) defines a cutoff \bar{c} such that $V(p) = U(p)$ for all p if and only if $c \geq \bar{c}$ by the same argument as in the proof of Theorem 1. Proposition 7 shows that if $c < \bar{c}$, the optimal strategy must satisfy either (3.1) or (3.2). By a similar argument as in the proof of Theorem 1, for c sufficiently close to \bar{c} , we have $V_{cf}(p) < U(p)$ for all $p \in [0, 1]$. Therefore $V_{ct}(p^*) > V_{cf}(p^*)$ and by Proposition 6, $V(p) = V_{ct}(p)$ and the optimal policy is of the form (3.1). Finally, as $\underline{\lambda} \rightarrow 0$, V_{cf} and V_{ct} converge to the values in the baseline model and if r is sufficiently small, $\underline{c}^{SEP} < \underline{c}$ so that the optimal policy is of the form (3.2) for $c \in (\underline{c}^{SEP}, \underline{c})$. \square

⁵⁶This can also be obtained directly by solving the ODEs for $r = 0$ to obtain $\underline{V}_{ct}(p^*)$ and $\overline{V}_{ct}(p^*)$.

References

- ARROW, K. J., D. BLACKWELL, AND M. A. GIRSHICK (1949): “Bayes and Minimax Solutions of Sequential Decision Problems,” *Econometrica*, pp. 213–244. [2](#)
- BARDI, M., AND I. CAPUZZO-DOLCETTA (1997): *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhäuser. [15](#), [36](#)
- BURKE, J. (2008): “Primetime Spin: Media Bias and Belief Confirming Information,” *Journal of Economics and Management Strategy*, 17(3), 633–665. [18](#)
- CHAN, J., A. LIZZERI, W. SUEN, AND L. YARIV (2016): “Deliberating Collective Decisions,” *Review of Economic Studies*, forthcoming. [5](#), [15](#)
- CHE, Y.-K., AND K. MIERENDORFF (2017): “Supplemental Material for ‘Optimal Sequential Decision with Limited Attention’,” <http://www.homepages.ucl.ac.uk/~uctpkmi/papers/CM-Supp-2017-07-02.pdf>. [6](#)
- DAMIANO, E., H. LI, AND W. SUEN (2017): “Learning while Experimenting,” unpublished. [5](#)
- DRUGOWITSCH, J., R. MORENO-BOTE, A. CHURCHLAND, M. SHADLEN, AND A. POUGET (2012): “The Cost of Accumulating Evidence in Perceptual Decision Making,” *Journal of Neuroscience*, 32, 3612?–3628. [26](#)
- FRANCETICH, A. (2014): “Managing Multiple Research Projects,” unpublished manuscript, Department of Decision Sciences and IGIER, Bocconi University. [5](#)
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2017): “Speed, Accuracy, and the Optimal Timing of Choices,” Available at SSRN 2602927. [4](#), [5](#), [6](#), [23](#), [26](#)
- GENTZKOW, M., AND J. SHAPIRO (2008): “Competition and Truth in the Market for News,” *Journal of Economic Perspective*, 22, 133–154. [18](#)
- HAWKINS, G., B. FORSTMANN, E.-J. WAGENMAKERS, R. RATCLIFF, AND S. BROWN (2015): “Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making,” *Journal of Neuroscience*, 35, 2476–2484. [26](#)
- HENRY, E., AND M. OTTAVIANI (2017): “Research and Approval Process: The Organization of Persuasion,” *Economic Journal*, forthcoming. [5](#)
- KE, T. T., AND J. M. VILLAS-BOAS (2016): “Optimal Learning before Choice,” unpublished manuscript, MIT. [6](#)
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 73(1), 39–68. [5](#)
- KLEIN, N., AND S. RADY (2011): “Negatively Correlated Bandits,” *The Review of Economic Studies*, 78(2), 693–732. [5](#)
- LATIMER, K., J. YATES, M. MEISTER, A. HUK, AND J. PILLOW (2015): “Single-trial spike trains in parietal cortex reveal discrete steps during decision-making,” *Science*, 349, 184–187. [26](#)
- MATEJKA, F., AND A. MCKAY (2015): “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” *The American Economic Review*, 105(1), 272–298. [6](#), [23](#)
- MAYSKAYA, T. (2016): “Dynamic Choice of Information Sources,” unpublished

- manuscript, California Institute of Technology. 6
- MEYER, M. A. (1991): “Learning from Coarse Information: Biased Contests and Career Profiles,” *Review of Economic Studies*, 58(1), 15–41. 18
- MOSCARINI, G., AND L. SMITH (2001): “The Optimal Level of Experimentation,” *Econometrica*, 69(6), 1629–1644. 5
- MULLAINATHAN, S., AND A. SHLEIFER (2005): “The Market for News,” *American Economic Review*, 95(4), 1031–1053. 18
- NIKANDROVA, A., AND R. PANCS (2017): “Dynamic Project Selection,” *Theoretical Economics*, forthcoming. 5
- OLIVEROS, S. (2015): “Demand for Slant: How Abstention Shapes Voters’ Choice of News Media,” *Economic Journal*, 125(587), 1327–1368. 18
- PEREGO, J., AND S. YUKSEL (2016): “Media Competition and the Source of Disagreement,” unpublished manuscript, NYU and UCSB. 18
- PERSICO, N. (2004): “Committee Design with Endogenous Information,” *Review of Economic Studies*, 71, 165–191. 15
- PESKIR, G., AND A. SHIRYAEV (2006): *Optimal Stopping and Free-Boundary Problems*. Birkhäuser Verlag. 5, 23, 25
- RATCLIFF, R., AND G. MCKOON (2008): “The Diffusion Decision Model: Theory and Data for Two-choice Decision Tasks,” *Neural computation*, 20(4), 873–922. 5, 26
- SIMS, C. A. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3), 665–690. 6
- STEINER, J., C. STEWART, AND F. MATEJKA (2015): “Rational Inattention Dynamics: Inertia and Delay in Decision-Making,” Discussion paper, CEPR Discussion Papers. 6
- STEPHENSON, M. C. (2010): “Information Acquisition and Institutional Design,” *Harvard Law Review*, 124, 1422–1483. 15
- SUEN, W. (2004): “The Self-Perpetuation of Biased Beliefs,” *The Economic Journal*, 114(495), 337–396. 18
- WALD, A. (1947): “Foundations of a General Theory of Sequential Decision Functions,” *Econometrica*, pp. 279–313. 2
- WOODFORD, M. (2016): “Optimal Evidence Accumulation and Stochastic Choice,” unpublished manuscript, Columbia University. 6
- YOON, Y. (2016): “Why Are Information Senders Biased? Indirect Demand for Information Bias,” unpublished manuscript, Nanjing Audit University. 18
- ZHONG, W. (2017): “Optimal Dynamic Information Acquisition,” unpublished manuscript, Columbia University. 5