Disinformation in the Wald Model†

Florian Brandl
University of Bonn

Xianwen Shi
University of Toronto

In the classical sequential sampling model of Wald (1945), a decision maker (Alice) learns a binary state from a noisy signal. We study the effects of disinformation by introducing an adversary (Bob) who can pay a cost to distort the signal. Both players are Bayesian, ex-ante symmetrically informed, and share a common prior about the state. Alice wants to choose an action that matches the state, while Bob prefers her to choose a high action regardless of the state. We show that disinformation invariably reduces Alice's welfare and decision accuracy, but its effect on her decision time is ambiguous. Although Bob has an incentive to engage in distortion, it may backfire on him in equilibrium. We also analyze how the distribution of Bob's distortion cost affects the equilibrium strategies and outcomes of both players. The basis for our results are novel insights into the classic sequential sampling problem with more than two states.

1. Introduction

A major problem we face is the spread of disinformation. Disinformation is false or misleading information deliberately created to influence public perception or individual behavior.[1] Disinformation affects various domains, such as politics, public health, and social issues. For instance, foreign actors use disinformation tactics to meddle with other countries' democratic processes and advance their own agendas; anti-vaccine groups disseminate false information about the safety and effectiveness of vaccines to deter people from immunizing themselves and their families; and fossil fuel lobbyists fund disinformation campaigns to create uncertainty about the effect of greenhouse gas emissions on global warming with the intent of prolonging the use of fossil fuels. Disinformation campaigns do not always succeed in deceiving their intended targets. Some targets are aware of the risk of being misled and take steps to protect themselves from false or biased information. For instance, they critically evaluate information and discount information that favors the agenda of the disinformation source. If the evidence is inconclusive, they may seek more information until a desired level of confidence is reached.

†Preliminary draft. Comments welcome.

[1] Misinformation, on the other hand, refers to false or misleading information without the intention to mislead.

Common problems of disinformation share several distinctive features. First, prior to the disinformation campaign, the manipulator and the receiver of disinformation are similarly (un)informed about the ground truth. For example, the interfering foreign government may not have better insights than domestic voters about which party or candidate is best for the country; and anti-vaccine activists and individuals considering vaccination have similar access to information about the safety and effectiveness of vaccines. Second, the manipulator has a clear motive, and their interest is not aligned with the receiver. Third, although the receiver has no control over the data quality, they can choose how much information to collect, for example, by delaying their decision and waiting for more information to arrive. Finally, it is costly for the receiver to collect information and for the manipulator to launch a disinformation campaign.

One way to approach the problem of disinformation is to focus on the psychological factors that influence how people process and evaluate information, such as confirmation bias or motivated reasoning. However, this perspective does not account for the strategic interactions between the players who produce and consume disinformation. We assume that both the manipulator and the receiver are rational and Bayesian. We then ask: How does the decision maker optimally balance the cost of information gathering and the benefit of better decision-making under the influence of disinformation? To what extent can the manipulator achieve their desired outcome by distorting information? How does disinformation affect the quality and timing of the decision maker's choices? Can the (negative) effects of disinformation be mitigated by raising the cost of manipulation?

To answer these questions, we develop a model of disinformation where a decision maker (Alice) solves Wald's sequential sampling problem (Wald, 1945) when the data-generating process may be influenced by a manipulator (Bob). We consider a variant of the drift-diffusion model in which the data is generated from a Brownian motion with a drift that depends on an underlying state and Bob's manipulation action. In that model, both players share a common prior belief over a binary state. In the beginning, Bob's manipulation cost is drawn from a known distribution and observed by Bob. Bob once and for all decides whether to manipulate, and his action is hidden from Alice. Manipulation increases the drift of the Brownian motion by a fixed amount so that the drift takes one of four values depending on the state and Bob's action. Alice can learn about the state and Bob's action by observing the Brownian motion at a constant flow cost, and she has to decide when to stop sampling and choose an action, either high or low, based on her observations. Alice prefers her action to be aligned with the state, while Bob wants her to choose the high action in either state. The game ends, and payoffs are realized as soon as Alice acts.

In our model of disinformation, an equilibrium always exists. In any equilibrium, Bob manipulates if his cost is below a cutoff. Alice collects data as long as her belief about the state lies between two cutoff paths, and she takes the high (low, respectively) action when the upper (lower, respectively) cutoff is first reached. In contrast to the standard Wald problem, in which the goal is to discern two possible drifts and the optimal cutoffs are constant over time, here, the upper cutoff path decreases over time while the lower cutoff path increases over time. Alice learns not only about the state but also about Bob's action. Until one of the cutoffs has been reached, Alice's updates towards distributions of the drift for which the process is less informative, which decreases her continuation value. Hence, the continuation region becomes narrower as time proceeds. Disinformation always hurts Alice. Since both parties are symmetrically informed about the state, Bob's manipulation does not convey any information about the state and serves as a pure obfuscation for Alice's learning. This is in contrast to the signaling model, where the strategic sender is informed, and hence, the receiver may benefit from signaling in equilibrium. Disinformation hurts Alice because it always lowers her decision accuracy, although the expected decision time may be higher or lower compared to the case of no disinformation. Bob may or may not gain from engaging in disinformation. Increasing the manipulation cost leads to equilibria in which Bob manipulates less frequently. This suggests that higher barriers to manipulation effectively deter manipulation and thereby prevent harm to the receiver.

2. Related literature

This paper investigates the effects of disinformation on rational and fully Bayesian individuals. Disinformation is a pervasive phenomenon that can have harmful consequences for various domains, such as politics, health, and climate change (see, e.g., Benkler et al., 2018; Melchior and Oliveira, 2022; Gwiazdon and Brown, 2023). Previous research has mainly explored the psychological biases that make people susceptible to disinformation, and proposed interventions such as Internet literacy education, critical thinking, and digital citizenship. In contrast, we take a different approach: we model all players as rational and Bayesian, using a sequential sampling model based on Wald (1945) to analyze the equilibrium and welfare implications of disinformation.

Our paper is related to the theoretical literature on fake news and online manipulation (see, e.g., Dellarocas, 2006; Glazer et al., 2021, and reference therein). Dellarocas (2006) studies a static signaling model of online forum manipulation and shows that a separating equilibrium exists in which only high-quality sellers buy fake reviews, and consumers benefit from these sellers' strategic manipulation. Glazer et al. (2021) examine whether a strategic rating platform can design a dynamic reporting policy about (potentially fake) messages to inform the receiver better and show that any manipulation by the platform always hurts the receiver. The result of the latter is similar to the one obtained in this paper, but the settings of the two models are very different. There is no third party like platform in our model which is based on the classical sequential sampling model by Wald (1945).

Since the pioneering work of Wald (1945), Wald and Wolfowitz (1948), and Arrow et al. (1949), the optimal stopping problems of sequential hypothesis testing have been extensively studied in statistics . For a survey, see Lai (1997), and for a textbook treatment, see Chernoff (1972) or Shiryaev (2008). However, most of the literature relies on a Gaussian or two-point prior distribution. Ekström and Vaicenavicius (2015) generalize the sequential hypothesis testing problem to a general prior distribution and establish properties of the optimal solution. Our model also adopts a two-point prior, but the presence of manipulation makes it equivalent to a problem with a four-point prior. Drawing on insights from Ekström and Vaicenavicius (2015), our analysis of the single-agent benchmark derives new comparative statics with respect to the prior distribution.

There is a growing literature in economic theory that applies the Wald model to study optimal stopping problems in both a single-agent and a strategic multi-agent setting. For example, Fudenberg et al. (2018) use the Wald model with a Brownian prior distribution to study the joint distribution of the accuracy and timing of decisions and show that earlier decisions are more accurate. Their insight about speed and accuracy is generalized by Liang et al. (2022), where an agent optimally allocates attention among several different information sources to learn about a Gaussian multidimensional state.[2] By contrast, we compare the accuracy under optimal stopping across different prior distributions and show that if a prior distribution is obtained from another by adding independent noise, the expected accuracy is lower under the former prior than under the latter. This result, with independent interest, is essential for us to show that Alice is always worse off from disinformation. The Wald model has also been used to study strategic interaction among different agents, such as persuasion (Henry and Ottaviani, 2019),[3] committee deliberation (Chan et al., 2018), and delegation (McClellan, 2022). In these models, the upper and lower cutoff in the stopping problem are controlled by different agents, while in our model, the same agent (Alice) controls both cutoffs.

Our model assumes that both Alice and Bob are uninformed about the state, differentiating it from recent papers on stopping problems with manipulated signals. These papers take either the perspective of dynamic signaling, where the informed sender's signals may reveal their private information (see, e.g., Daley and Green, 2012; Dilmé, 2019; Gryglewicz and Kolb, 2022; Cetemen and Margaria, 2023), or the perspective of reputation formation where the signals generated by the informed agent's action may reveal the agent's private type (see, e.g., Ekmekci and Maestri, 2022; Ekmekci et al., 2022). In all these papers, agents are ex-ante asymmetrically informed, and the informed agent's manipulated signals may reveal what they know. By contrast, in our model, both agents are symmetrically uninformed, so Bob's manipulation can only add noise to the signals that Alice observes, making her worse off.

Finally, our paper connects to the computer science literature on adversarial learning and robust estimation (see, e.g., Diakonikolas et al., 2019; Lai et al., 2016; Charikar et al., 2017). These papers adopt a worst-case analysis—the adversary has the exact opposite preferences as the decision maker, the decision maker solves maxmin problem, and there is no cost for

[2] Che and Mierendorff (2019) apply a Wald model with Poisson signals to explore how a decision maker should allocate their attention among different information sources and find that with Poisson signals, speedy decisions are not necessarily more accurate. Auster et al. (2024) generalize the Bayesian analysis of the Wald model to the case where the decision maker is ambiguity-averse.

[3] See also Orlov et al. (2020) for a related dynamic persuasion model where an agent discloses information over time to persuade the principal to delay exercising an option.

manipulation. Moreover, their models assume static learning rather than dynamic learning as in our paper.

3. The Model

We start with an informal description of the model. Consider a decision maker, called Alice, who must choose a binary action (e.g., whether to vote for a candidate, whether to get vaccinated, or whether to support a climate initiative). She holds a prior about an underlying state and would like her action to match the state. She can either act on her prior or she can collect more information sequentially to sharpen her belief. There is an adversary (e.g., a foreign government, an anti-vaccine activist, or a fossil-fuel firm) called Bob, who has a clear agenda and can incur a cost to manipulate the additional information that Alice wishes to collect. Both players are rational and Bayesian.

We cast our problem in a strategic variant of the sequential sampling model of Wald (1945). Alice incurs a constant flow cost to observe a Brownian motion whose drift depends on the unknown state and Bob's hidden action of manipulation. Alice can learn about the state from the Brownian signals over time, but her signals are distorted by Bob's hidden manipulation. Based on observed signals, Alice decides when to stop observing and when she stops she must choose one of the two actions, high or low, to take. Alice would like her action to match the state, but Bob always wants her to take the high action independent of the state. We are interested in how equilibrium manipulation affects decision time, decision quality and players' welfare.

3.1. Basic Setup

Formally, let $(\Omega, \mathcal{F}, P)$ be a probability space supporting a standard Brownian motion $W = (W_t)_{t\in\mathbb{R}_+}$ and two random variables, θ and γ. We assume that θ, W, and γ are independent. We call θ the *state* and its distribution μ the *prior distribution*. The *prior belief* is $p_0 = P(\theta \geq 0) \in (0,1)$.[4] We call γ Bob's lump-sum *manipulation cost* and we assume that its distribution Γ is absolutely continuous with full support on $[0,1]$. The distributions μ and Γ are common knowledge among the players, but the manipulation cost γ is privately observed by Bob and neither player knows the realization of the state θ.

Time $t \in [0,\infty)$ is continuous. At time $t = 0$, Bob privately observes the realization of γ, and chooses an *action* $y \in \{0, \overline{m}\}$ with cost $\gamma y/\overline{m}$. That is, Bob incurs a lump-sum cost γ when choosing $\overline{m}$ and 0 otherwise. From time $t = 0$ on, Alice observes a process $X^y = (X^y_t)_{t\geq 0}$ with

$$dX^y_t = (\theta + y)dt + dW_t$$

We call X^y the *manipulated process*. The (completed) filtration generated by X^y is denoted by

[4] We assume that θ has bounded support, and that $p_0 \in (0,1)$ and $P(\theta = 0) = 0$.

$\mathcal{X}^y = (\mathcal{X}^y_t)_{t\in\mathbb{R}_+}$. It induces the *belief process* $\pi^y = (\pi^y_t)_{t\in\mathbb{R}_+}$ given by

$$\pi^y_t = P\left(\theta \geq 0 \mid \mathcal{X}^y_t\right).$$

We assume $0 < \overline{m} < \operatorname{ess\,inf}|\theta|$ so that θ and $\theta + y(\gamma)$ have the same sign, and hence X^y almost surely reveals the sign of θ in the limit.

Alice faces a Wald problem with a manipulated process. That is, she chooses a stopping time τ for the manipulated process $\mathcal{X}^y$, and upon stopping, she chooses one of two actions $\{h, l\}$, called the high and the low action, optimally based on her belief π^y_τ. Alice wants to match the state: her utility is 1 if she takes the high (low) action in a positive (negative) state, and 0 otherwise. The *observation cost* is $c > 0$ per unit of time. Both players do not discount. Hence, Alice's expected payoff given τ is

$$\mathrm{E}\left[g(\pi^y_\tau) - c\tau\right],$$

where $g(p) = p \vee (1-p)$ for $p \in [0,1]$.

Bob prefers Alice to take the high action independently of the state: his utility is 1 if Alice takes the high action, and 0 otherwise. Hence, if Bob chooses y and Alice chooses the stopping time τ for X^y (and chooses an action optimally upon stopping), then for given γ, Bob's expected payoff is

$$P\left(\pi_\tau \geq \frac{1}{2}\right) - \gamma\frac{y}{\overline{m}}.$$

The players' utility functions and Alice's observation cost c are common knowledge. Table 1 summarizes these and later definitions.

3.2. Strategies and Equilibrium Concept

The two players' strategies are modeled as follows. Bob's *manipulation strategy* is represented a measurable function $y(\gamma)$ with $y\colon \mathbb{R}_+ \to \{0, \overline{m}\}$. The restriction of Bob's strategy to be deterministic is without loss because γ is continuous and has a full support distribution.

For a fixed manipulation strategy for Bob, Alice faces an optimal stopping problem with two types of uncertainties: a non-strategic uncertainty about the state θ and a strategic uncertainty about Bob's manipulation choice $y(\gamma)$. The strategic uncertainty, absent in the standard Wald model, complicates Alice's learning problem. Fortunately, we can translate Alice's problem with strategic uncertainty into a Wald problem with non-strategic uncertainty only. For any given manipulation strategy $y(\gamma)$, we can redefine the "manipulated state" as $\theta + y(\gamma)$ and let μ^y denote the distribution of $\theta + y(\gamma)$, called the "manipulated prior". Then Alice's optimal stopping problem for a given manipulation strategy becomes a standard Wald problem with prior μ^y. The characterization of solutions to the Wald problem with a general prior is given below in Section 4.

Since Alice's stopping problem for given $y(\gamma)$ can be reformulated as a Wald problem, we can restrict Alice's strategies to stopping times induced by two time-dependent boundaries $b_h(t)$

and $b_l(t)$ on the anticipated observed process X^y (see Appendix A for the definition of b_h, b_l). Note that Bob's best response involves reasoning about Alice's belief process if she anticipates y and Bob plays $\tilde{y}$. We characterize this belief process in Appendix B through a stochastic differential equation. We then translate the optimal boundaries β_h, β_l for the belief process π^y to optimal boundaries b_h, b_l on Alice's anticipated observed process X^y and let Alice stop the process $X^{\tilde{y}}$ according to these boundaries b_h, b_l. We refer to b_h, b_l as the boundaries for the observed process X^y induced by boundaries β_h, β_l for the belief process π^y. The formulation based on b_h, b_l for the anticipated observed process X^y is sometimes more convenient, especially in computing Bob's expected deviation gains. The restriction to deterministic boundaries entails no substantial loss since optimal deterministic boundaries b_h, b_l always exist.

We consider Bayes-Nash equilibria of this game.

Definition 1. A triple (b_h, b_l, y) is an *equilibrium* if the following hold.

1. Optimality for Alice: b_h and b_l are optimal boundaries for the observed process X^y with respect to the prior distribution μ^y, where μ^y is the distribution of $\theta + y(\gamma)$.

2. Optimality for Bob:

$$P\left(X^y_\tau = b_h(\tau)\right) - \gamma\frac{y(\gamma)}{\overline{m}} \geq P\left(X^{\tilde{y}}_{\tilde{\tau}} = b_h(\tilde{\tau})\right) - \gamma\frac{\tilde{y}}{\overline{m}}$$

for any measurable function $\tilde{y}\colon [0,1] \to \{0, \overline{m}\}$, where $\tau = \inf\{t \in \mathbb{R}_+ \colon X^y_t \in \{b_h(t), b_l(t)\}\}$ and $\tilde{\tau} = \inf\{t \in \mathbb{R}_+ \colon X^{\tilde{y}}_t \in \{b_h(t), b_l(t)\}\}$.

Recall that Bob's manipulation cost is independent of the state and the Brownian motion W. Hence, the probability of changing Alice's action from low to high by manipulating is independent of the manipulation cost. It follows that Bob's gain from manipulation is strictly decreasing in the manipulation cost, and so a best response for Bob is given by a cutoff such that he manipulates if his manipulation cost is below the cutoff and does not manipulate otherwise. The optimal cutoff is the difference between the probability that Alice takes the high action if Bob plays $\overline{m}$ and the probability that she takes the high action if Bob plays 0.

Definition 2. A manipulation strategy $y\colon [0,1] \to \{0, \overline{m}\}$ is a *cutoff strategy* if there is $\gamma_0 \in [0,1]$ such that $y = \overline{m}\mathbf{1}_{[0,\gamma_0]}$.

The preceding remarks give the following statement.

Proposition 1. *For any strategy for Alice given by boundaries b_h, b_l for the observed process, Bob's unique (up to a set of measure 0) best response is a cutoff strategy.*

We thus take Bob's strategy set to be $[0,1]$ from now on, where $\gamma_0 \in [0,1]$ corresponds to the cutoff strategy $\overline{m}\mathbf{1}_{[0,\gamma_0]}$.

As we mentioned earlier, to solve our strategic Wald problem, we first have to solve the Wald problem with a general prior, which we will do below.

4. The Single-Agent Problem

We begin this subsection by presenting in Section 4.1 the solution of the single-agent Wald problem under a general prior. This part of the analysis follows Ekström and Vaicenavicius (2015). Then, we examine in Section 4.2 how optimal boundaries depend on the prior distribution, which is key to our later analysis of how information manipulation affects equilibrium and welfare outcomes. This part of the analysis is original and may be of independent interest.

4.1. Optimal Stopping with a General Prior

A decision maker wants to choose one of the two actions $\{h, l\}$ to match the state θ: her utility is 1 if she takes the high (low) action in a positive (negative) state, and 0 otherwise. The observed process $X = (X_t)_{t\in\mathbb{R}_+}$ is given by:

$$dX_t = \theta dt + dW_t.$$

The filtration $\mathcal{X} = (\mathcal{X}_t)_{t\in\mathbb{R}_+}$ generated by X induces the belief process $\pi = (\pi_t)_{t\in\mathbb{R}_+}$ given by

$$\pi_t = P\left(\theta \geq 0 \mid \mathcal{X}_t\right).$$

It can be shown that π satisfies

$$d\pi_t = \sigma(t, \pi_t)d\hat{W}_t$$

for $\sigma\colon \mathbb{R}_+ \times (0, 1) \to \mathbb{R}$ continuously differentiable and $\hat{W} = (\hat{W}_t)_{t\in\mathbb{R}_+}$ a standard Brownian motion. We give details in the Appendix A.

The decision maker chooses a stopping time τ for the process X, and upon stopping, she chooses optimally based on her belief π_τ. The decision maker's expected payoff given τ is

$$\mathrm{E}\left[g(\pi_\tau) - c\tau\right],$$

where $g(p) = p \vee (1-p)$ for $p \in [0, 1]$. We say that τ is optimal if it is payoff-maximizing among all $\mathcal{X}$-stopping times.

We denote by $\pi^{t,p}$ the belief process from time t onwards conditional on having belief p at t. That is, for $t \in \mathbb{R}_+$ and $p \in (0, 1)$, let $\pi^{t,p} = (\pi^{t,p}_{t+s})_{s\in\mathbb{R}_+}$ be the unique solution to

$$d\pi^{t,p}_{t+s} = \sigma(t+s, \pi^{t,p}_{t+s})d\hat{W}_{t+s} \quad \text{and} \quad \pi^{t,p}_t = p.$$

The *continuation value* depending on the time and the belief, denoted by $V\colon \mathbb{R}_+ \times (0, 1) \to \mathbb{R}$, is defined as

$$V(t, p) = \sup_\tau \mathrm{E}\left[g(\pi^{t,p}_{t+\tau}) - c\tau\right],$$

where the supremum is taken over all $\mathcal{X}$-stopping times. Heuristically, $V(t, p)$ is the expected payoff (net of the observation cost ct incurred thus far) conditional on having belief p at t.

Ekström and Vaicenavicius (2015) show that an optimal stopping time is given by two time-dependent boundaries on the belief process, one above $\frac{1}{2}$ and weakly decreasing, and one below $\frac{1}{2}$ and weakly increasing, and determined their limit value as t goes to infinity. These results are derived from the fact that $V(t,p)$ is non-increasing in t for each p, and convex in p for each t. The case when the prior distribution μ is supported on two points is special since it is the unique instance where the value function and thus the optimal boundaries are constant in time.

Theorem 1 (Ekström and Vaicenavicius, 2015)**.** *Under the assumptions above, the following hold.*

(i) $V(t,p)$ *is continuous, and it is convex in* p *for all* $t \geq 0$ *and non-increasing in* t *for all* $p \in (0,1)$*. Moreover,* $V(t,p)$ *is constant for all* $p \in (0,1)$ *if and only if* μ *is supported on two points.*

(ii) Let $\beta_h, \beta_l \colon \mathbb{R}_+ \to (0,1)$ *be defined by*

$$\beta_h(t) = \sup\{p \in [\frac{1}{2},1) \colon V(t,p) > g(p)\} \text{ and } \beta_l(t) = \inf\{p \in (0,\frac{1}{2}] \colon V(t,p) > g(p)\}.$$

Then, the following stopping time τ *is optimal:*

$$\tau = \inf\{t \in \mathbb{R}_+ \colon \pi_t \in \{\beta_h(t), \beta_l(t)\}\}.$$

(iii) β_h *is continuous and non-increasing and* β_l *is continuous and non-decreasing, and both are constant if and only if* μ *is supported on two points. Moreover,* $0 < \beta_l(t) < \frac{1}{2} < \beta_h(t) < 1$ *for all* $t \in \mathbb{R}_+$*.*

(iv) Let $z_+ = \inf\{z \geq 0 \colon \mu([z, z+\epsilon)) > 0 \text{ for all } \epsilon > 0\}$ *and* $z_- = \sup\{z \leq 0 \colon \mu((z-\epsilon, z]) > 0 \text{ for all } \epsilon > 0\}$*. Then,* $\lim_{t\to\infty} \beta_h(t) = \lim_{t\to\infty} \beta_l(t) = \frac{1}{2}$ *if* $z_+ = z_- = 0$*, and* $\lim_{t\to\infty} \beta_h(t) = p_h$ *and* $\lim_{t\to\infty} \beta_l(t) = p_l$ *if* $z_+ > z_-$*, where* p_h, p_l *are the constant optimal boundaries for a two-point prior distribution supported on* $\{z_+, z_-\}$*.*

4.2. Comparative Statics of Prior Distributions

We study how the continuation value and, thus, the optimal boundaries depend on the prior distribution. We write V^μ, β_h^μ, π^μ, and so forth to indicate which prior distribution the objects are derived from. When two prior distributions $\mu, \tilde{\mu}$ are considered, we use the shorthands $V, \tilde{V}$, $\beta_h, \tilde{\beta}_h$, $\pi, \tilde{\pi}$, and so forth.

First, we show that the expected payoff, that is, the continuation value at time 0, is convex in the prior distribution. This follows from the fact that for fixed boundaries, the expected payoff is linear in the prior distribution. Applying this fact to the optimal boundaries for the convex combination of two prior distributions, and using that a maximum of linear functions is convex gives the statement.

Lemma 1. *For any prior distribution* μ*, let* $p_0^\mu = \mu([0,\infty))$*. Then,* $V^\mu(0, p_0^\mu)$ *is convex in* μ*.*

Lemma 1 shows that more randomness in the prior distribution decreases the expected payoff. We show that for certain convex combinations of prior distributions—the ones which arise in the strategic sampling problem considered in Section 3—the optimal boundaries become narrower at time 0 as a consequence.

Definition 3. Let θ, ξ be independent random variables with distributions μ, ν such that ν has finite support and $\theta > 0$ if and only if $\theta + \xi > 0$. Then, we say that the convolution $\tilde{\mu} = \mu * \nu$ is a *sign-preserving random shift* of μ.

In other words, $\tilde{\mu}$ is a sign-preserving random shift of μ if it is a convex combination of shifts of μ, each preserving the the probability on positive states.[5] Proposition 2 shows that replacing a prior distribution by a sign-preserving random shift moves the optimal boundaries at time 0 and for any sufficiently late time closer to $\frac{1}{2}$.

Proposition 2. *Let μ be a prior distribution, and let $\tilde{\mu}$ be a sign-preserving random shift of μ. Then, for all $p \in (0,1)$, $V(0,p) \geq \tilde{V}(0,p)$. Moreover, $\beta_h(0) \geq \tilde{\beta}_h(0)$ and $\tilde{\beta}_l(0) \geq \beta_l(0)$, and there is $t_0 \in \mathbb{R}_+$ such that $\beta_h(t) \geq \tilde{\beta}_h(t)$ and $\tilde{\beta}_l(t) \geq \beta_l(t)$ for all $t \geq t_0$.*

When the prior distribution is supported on two points, one can say more about how sign-preserving random shifts change the continuation value, the optimal boundaries, and the accuracy of the decisions. Recall that the continuation value and the optimal boundary are independent of time for two-point prior distributions, and that the optimal boundaries are monotonic for any prior distribution (Theorem 1(iii)). In combination with the fact that sign-preserving random shifts decrease the continuation value and narrow the optimal boundaries at time 0 (Proposition 2), this implies that both effects hold for all times.

Corollary 1. *Let μ be a prior distribution supported on two points, and let $\tilde{\mu}$ be a sign-preserving random shift of μ. Then, for all $t \in \mathbb{R}_+$ and $p \in (0,1)$, $V(t,p) \geq \tilde{V}(t,p)$. Moreover, for all $t \in \mathbb{R}_+$, $\beta_h(t) \geq \tilde{\beta}_h(t)$ and $\tilde{\beta}_l(t) \geq \beta_l(t)$.*

Fudenberg et al. (2018) define the accuracy of a stopping time at some time t as the probability of choosing the correct action conditional on stopping at t. They show that under optimal stopping, the accuracy is non-increasing in the decision time for normal prior distributions. By contrast, we compare the accuracy under optimal stopping across different prior distributions. The expected accuracy of a stopping time is the probability with which it leads to choosing the correct action.

Definition 4. Let τ be an $\mathcal{X}$-stopping time. Then, the *expected accuracy* of τ is

$$\mathrm{Acc}(\tau) = \mathrm{E}\left[g(\pi_\tau)\mathbf{1}\{\tau < \infty\}\right].$$

From the preceding results, we can conclude that sign-preserving random shifts of two-point prior distributions decrease the expected accuracy under optimal stopping. In fact, the

[5] The statements about random shifts extend to distributions ν with infinite support. We restrict to the finite case since it is sufficient for the rest of the paper and avoids technicalities.

same arguments show that the choices for two-point prior distributions are almost surely more accurate—that is, taken with a more extreme belief—than for the random shift. It is open whether Corollary 2 holds if μ is supported on more than two points.

Corollary 2. *Let μ be a prior distribution supported on two points, and let $\tilde{\mu}$ be a sign-preserving random shift of μ. Let $\tau, \tilde{\tau}$ be optimal stopping times for $\mu, \tilde{\mu}$ given by Theorem 1(ii). Then,* $\mathrm{Acc}[\tilde{\tau}] \leq \mathrm{Acc}[\tau]$.

Lastly, we consider the expected observation time. Since the cost of observations is constant, the expected observation time is proportional to the expected observation cost.

Definition 5. Let τ be an $\mathcal{X}$-stopping time. Then, the *expected observation time* of τ is $\mathrm{E}[\tau]$.

We show that sign-preserving random shifts can increase or decrease the expected observation time under optimal stopping. The examples we give use two-point prior distributions and sign-preserving random shifts thereof that are symmetric about 0. Hence, even symmetry of the optimal boundaries and the fact that sign-preserving random shifts of two-point prior distributions narrow the optimal boundaries do not suffice to pin down the change in the expected observation time.[6]

Proposition 3. *There exists a prior distribution μ supported on two points and a sign-preserving random shift $\tilde{\mu}$ of μ such that $\mathrm{E}[\tau] < \mathrm{E}[\tilde{\tau}]$ when $\tau, \tilde{\tau}$ denote optimal stopping times for $\mu, \tilde{\mu}$, respectively. Likewise, there exist such $\mu, \tilde{\mu}$ such that the opposite inequality holds.*

5. Equilibrium Analysis

We study properties of equilibria, compare equilibrium outcomes to the counter-factual case where Bob cannot influence the observed process, and analyze the comparative statics of different prior distributions and different distributions for the manipulation cost.

5.1. Equilibrium Existence

We first show that an equilibrium always exists as follows. Consider the function mapping $q \in [0,1]$ to optimal boundaries b_h^q, b_l^q for the observed process when Bob manipulates with probability q, and the function mapping the pair b_h^q, b_l^q to $\Phi(q) = P(\gamma \leq \gamma_0^q)$, where γ_0^q is the optimal cutoff against b_h^q, b_l^q. Composing these maps gives a continuous function Φ from $[0,1]$ to $[0,1]$, which has a fixed-point by the intermediate value theorem. Since Φ is defined through iterated best responses, any such fixed-point induces an equilibrium.

Proposition 4. *For any prior distribution μ and any distribution Γ of the manipulation cost satisfying the above assumptions, there exists an equilibrium $(\beta_h, \beta_l, \gamma_0)$.*

[6] Note that the expected observation time can increase despite narrower boundaries since sign-preserving random shifts also change the volatility σ.

Next, we show that in any equilibrium, Alice almost surely observes for a positive amount of time and Bob manipulates non-trivially (with probability strictly between 0 and 1), unless the prior belief is extreme and Alice would stop at time 0 for the (non-manipulated) prior distribution.

Proposition 5. *Let μ be any prior distribution, let β_h, β_l be the optimal boundaries for μ as defined in Theorem 1(ii), and assume that $\beta_l(0) < p_0 < \beta_h(0)$. Then, in any equilibrium, (i) Alice almost surely observes past time* 0*, and (ii) Bob manipulates with probability strictly between* 0 *and* 1*.*

5.2. Welfare, Timing, and Accuracy

Consider two scenarios. In the first, Bob cannot influence the process observed by Alice, and Alice solves the optimal stopping problem for the given prior distribution. In the second, Bob can manipulate the process observed by Alice, and we consider an equilibrium where Alice solves the optimal stopping problem for the manipulated prior distribution. We refer to the first scenario as the *no manipulation benchmark*, and to the second as *equilibrium play*. We compare the players' welfare and the accuracy and timing of Alice's decisions between both scenarios.

First, we observe that in any equilibrium, Alice is worse off compared to the no manipulation benchmark, and if the prior distribution is supported on two points, her decisions are less accurate in equilibrium. The first claim follows from Proposition 2 and the fact that Bob manipulates with probability strictly between 0 and 1 in any equilibrium by Proposition 5. The second part is a consequence of Corollary 2.

Corollary 3. *Let μ be any prior distribution. Then, in any equilibrium, Alice's expected payoff is no more than $V^{\mu}(0, p_0)$. If μ is supported on two points, then in any equilibrium, Alice's accuracy is lower than in the no manipulation benchmark.*

By contrast, the comparison for Bob's expected payoff between both scenarios can go either way. This implies that social welfare—the sum of the players' expected payoffs—can be lower in equilibrium than in the no manipulation benchmark. On the other hand, Bob's gain from manipulation can outweigh Alice's loss, and so social welfare can also be higher in equilibrium. Similarly, the comparison for the expected observation time can go either way.

Proposition 6. *There exist a prior distribution μ, a distribution Γ for the manipulation cost, and an equilibrium such that compared to the no manipulation benchmark,*

1. *Bob's expected payoff is higher (lower),*
2. *social welfare is higher (lower), and*
3. *the expected observation time is lower (higher).*

5.3. Comparative Statics of Equilibria

We examine how equilibria depend on the prior distribution and on the distribution of the manipulation cost.

First, consider a two-point prior distribution with a prior belief for which Alice would observe past time 0 under the no manipulation benchmark. By Proposition 5, Bob's cutoff is non-trivial and Alice observes past time 0 in any equilibrium. We show that, however, that Bob's cutoff and Alice's expected observation time both go to 0 as the prior belief approaches the boundary of the observation region at time 0. To see the first part, recall from Proposition 2 and Proposition 5(i) that the boundary at time 0 in any equilibrium is in between the prior belief and the (constant) boundary for the no manipulation benchmark, and so as the prior belief goes to, say, the upper boundary for the no manipulation benchmark, so does the equilibrium boundary at time 0. Moreover, the upper equilibrium boundary is non-increasing. For a high prior belief, a small upwards shock thus suffices for Alice to stop and take the high action. Since on small time scales, the belief process is dominated by the volatility part and the drift part is negligible, Alice takes the high action with probability close to 1 even if Bob does not manipulate. As a result, manipulation is optimal for Bob only if his manipulation cost is very small, and so since Γ is continuous, the probability with which he manipulates is small as well. The fact that Alice's expected observation time goes to 0 follows from similar arguments.

Proposition 7. *Let $\mu = p_0\delta_1 + (1 - p_0)\delta_{-1}$ be a prior distribution supported on two points, and let β_h, β_l be the constant optimal boundaries for μ, and assume that $\beta_l < p_0 < \beta_h$. Let $(\beta_h^*, \beta_l^*, \gamma_0^*)$ be any equilibrium, and denote by τ^* the induced stopping time for the belief process in equilibrium. Then, Bob's cutoff γ_0^* and his manipulation probability $P(\gamma \leq \gamma_0^*)$, and Alice's expected observation time $\mathrm{E}[\tau^*]$ go to 0 as p_0 goes to β_h or β_l.*

Now we compare how the distribution of the manipulation cost affects the probability with which Bob manipulates in equilibrium. It is natural to assume that all else being equal, Bob manipulates less frequently if manipulation is more expensive. To make this precise, denote by F_Γ the cumulative distribution function of a distribution Γ. If $\Gamma, \tilde{\Gamma}$ are two distributions, Γ *stochastically dominates* $\tilde{\Gamma}$ if $F_\Gamma \leq F_{\tilde{\Gamma}}$. We show that if Γ stochastically dominates $\tilde{\Gamma}$, then for any equilibrium under Γ, there is an equilibrium under $\tilde{\Gamma}$ where Bob's manipulation probability is higher.

Proposition 8. *Let μ be any prior distribution, and let $\Gamma, \tilde{\Gamma}$ be two distributions for the manipulation cost such that Γ stochastically dominates $\tilde{\Gamma}$. Assume that under Γ, there is an equilibrium with cutoff γ_0^* where Bob manipulates with probability $q^* = F_\Gamma(\gamma_0^*)$. Then, under $\tilde{\Gamma}$, there is an equilibrium where Bob manipulates with probability at least q^*.*

Acknowledgements

The authors thank Philipp Strack and Matt Weinberg for valuable comments. Florian Brandl acknowledges support by the DFG under the Excellence Strategy EXC-2047. Xianwen Shi thanks financial support under an Insight grant from the Social Sciences and Humanities Research Council of Canada.

APPENDIX

A. Preliminaries

In preparation for the proofs, we make further definitions and summarize some known facts. Recall that $X = (X_t)_{t\in\mathbb{R}_+}$ is given by the equation

$$dX_t = \theta dt + dW_t,$$

and that $\mathcal{X} = (\mathcal{X}_t)_{t\in\mathbb{R}_+}$ denotes the completion of the filtration generated by X. It follows from the innovation theorem of Brownian motion (see, e.g., Harrison, 2013, Theorem 1.12) that

$$dX_t = \mathrm{E}\,[\theta \mid \mathcal{X}_t]dt + d\hat{W}_t,$$

where

$$\hat{W}_t = X_t - \int_0^t \mathrm{E}\,[\theta \mid \mathcal{X}_s]ds$$

is a standard Brownian motion.[7] For each $t \in \mathbb{R}_+$ and each $x \in \mathbb{R}$, let

$$\mu_{t,x}(dz) = \frac{\exp\left(\frac{2zx-z^2t}{2}\right)\mu(dz)}{\int_{\mathbb{R}} \exp\left(\frac{2z'x-z'^2t}{2}\right)\mu(dz')},$$

and

$$p(t,x) = \int_{\mathbb{R}_+} \mu_{t,x}(dz). \tag{1}$$

Heuristically, $\mu_{t,x}$ and $p(t,x)$ are the posterior distribution and the posterior belief upon observing $X_t = x$.[8] Following Ekström and Vaicenavicius (2015, Proposition 3.1, Proposition 3.4), for each $t \in \mathbb{R}_+$, almost surely,

$$\pi_t = p(t, X_t).$$

[7] Clearly, $\hat{W}$ is adapted to the filtration $\mathcal{X}$. In fact, it can be shown that the completion of the filtration generated by $\hat{W}$ equals $\mathcal{X}$ (see, e.g., Bain and Crisan, 2009, p. 35).

[8] To see that $\mu_{t,x}$ and $p(t,x)$ can be extended to $t = 0$, observe that $\mu_{0,x}$ can be seen as updating a belief at time $t = -2\epsilon$ on an observation at $t = 0$ (see also Ekström and Vaicenavicius, 2015, Section 3.3). Let

$$\nu(dz) = \frac{\exp(z^2\epsilon)\mu(dz)}{\int_{\mathbb{R}} \exp(z'^2\epsilon)\mu(dz')},$$

which is well-defined if μ has bounded support. Then,

$$\nu_{2\epsilon,x}(dz) = \frac{\exp\left(\frac{2zx-2z^2\epsilon}{2}\right)\nu(dz)}{\int_{\mathbb{R}} \exp\left(\frac{2z'x-2z'^2\epsilon}{2}\right)\nu(dz')} = \frac{\exp\left(\frac{2zx}{2}\right)\mu(dz)}{\int_{\mathbb{R}} \exp\left(\frac{2z'x}{2}\right)\mu(dz')} = \mu_{0,x}(dz).$$

Hence, if θ has distribution ν and $\tilde{X}_t = \theta(t+2\epsilon) + W_{t+2\epsilon}$ for $t \geq -2\epsilon$, $\tilde{X}_{-2\epsilon} = 0$, then $\mu_{0,\tilde{X}_0}(dz) = P\left(\theta \in dz \mid \tilde{X}_0\right)$.

Moreover, for each $t \in \mathbb{R}_+$, $p(t,\cdot)\colon \mathbb{R} \to (0,1)$ is differentiable in both variables, strictly increasing, and bijective. Hence, its inverse $p(t,\cdot)^{-1} = x(t,\cdot)\colon (0,1) \to \mathbb{R}$ exists and has the same properties. Hence, observing $X_t = x(t,p)$ induces the belief p.

The belief process π satisfies the equation

$$d\pi_t = \sigma(t,\pi_t)d\hat{W}_t,$$

where the volatility $\sigma\colon \mathbb{R}_+ \times (0,1)$ is given by $\sigma(t,p) = \partial_2 p(t,x(t,p))$ for all $t \in \mathbb{R}_+$ and $p \in (0,1)$. A calculation shows that

$$\sigma(t,p) = (1-p)\int_{\mathbb{R}_+} z\mu_{t,x(t,p)}(dz) - p\int_{\mathbb{R}_-} z\mu_{t,x(t,p)}(dz). \tag{2}$$

If μ is supported on two points $z_- < 0 < z_+$, then $\sigma(t,p) = p(1-p)(z_+ - z_-)$. In that case, the constant optimal boundaries are given by β_h, β_l, where β_h is the unique solution to

$$\frac{(z_+ - z_-)^2}{2c} = \frac{\beta_h}{1-\beta_h} - \frac{1-\beta_h}{\beta_h} + 2\log\left(\frac{\beta_h}{1-\beta_h}\right)$$

and $\beta_l = 1 - \beta_h$ (Shiryaev, 2008, Section 4.2, Theorem 5 and the remark thereafter). Since the term on the right-hand side is strictly increasing in β_h, it follows that, conversely, β_h is strictly increasing in $z_+ - z_-$.

B. The Off-Path Belief Process

Checking the equilibrium condition for Bob requires determining the probability that Alice chooses each action for deviations by Bob from his equilibrium strategy. To this end, the following lemma determines Alice's belief process evolves if she anticipates the cutoff strategy γ_0 but Bob chooses $m \in \{0, \overline{m}\}$ irrespective of his manipulation cost γ. Since θ, W, γ are independent, the distribution of this process does is independent of γ.

Lemma 2. *Let μ be any prior distribution, let Γ be any distribution of the manipulation cost, and let $\gamma_0 \in [0,1]$ be a cutoff strategy for Bob. For $m \in \{0, \overline{m}\}$, define $\pi^{\gamma_0,m} = (\pi_t^{\gamma_0,m})_{t\in\mathbb{R}_+}$ by letting $\pi_t^{\gamma_0,m} = p^{\gamma_0}(t, X_t^m)$, called Alice's belief process when anticipating γ_0 and observing X^m. Then, $\pi^{\gamma_0,m}$ satisfies*

$$d\pi_t^{\gamma_0,m} = \left(\mathrm{E}\left[\theta \mid \mathcal{X}_t^m\right] + m - \int_{\mathbb{R}} z\mu^{\gamma_0}_{t,X_t^m}(dz)\right)\sigma^{\gamma_0}(t,\pi_t^{\gamma_0,m})dt + \sigma^{\gamma_0}(t,\pi_t^{\gamma_0,m})d\hat{W}_t^m.$$

Here, the superscripts γ_0 and m refer to the corresponding manipulated prior distributions. For example, X^m is the observed process with drift term $\theta + m$, which has distribution $\mu * \delta_m$.

Proof. We omit the superscript γ_0 since it is fixed throughout the proof. Using (1), we observe that $p(t,x) \in C^\infty(\mathbb{R}_{++} \times \mathbb{R})$ and calculate partial derivatives of $p(t,x)$.

$$\begin{aligned}
\partial_1 p(t,x) &= -\int_{\mathbb{R}_+} \frac{z^2}{2}\mu_{t,x}(dz) + \int_{\mathbb{R}_+} \mu_{t,x}(dz) \int_{\mathbb{R}} \frac{z^2}{2}\mu_{t,x}(dz) \\
\partial_2 p(t,x) &= \int_{\mathbb{R}_+} z\mu_{t,x}(dz) - \int_{\mathbb{R}_+} \mu_{t,x}(dz) \int_{\mathbb{R}} z\mu_{t,x}(dz) \\
\partial_2^2 p(t,x) &= \int_{\mathbb{R}_+} z^2\mu_{t,x}(dz) - 2\int_{\mathbb{R}} z\mu_{t,x}(dz) \int_{\mathbb{R}_+} z\mu_{t,x}(dz) \\
&\quad - \int_{\mathbb{R}_+} \mu_{t,x}(dz) \int_{\mathbb{R}} z^2\mu_{t,x}(dz) + 2\int_{\mathbb{R}_+} \mu_{t,x}(dz) \left(\int_{\mathbb{R}} z\mu_{t,x}(dz)\right)^2.
\end{aligned}$$

Recall from Appendix A that

$$dX_t^m = \mathrm{E}\left[\theta + m \mid \mathcal{X}_t^m\right]dt + d\hat{W}_t^m = \left(\mathrm{E}\left[\theta \mid \mathcal{X}_t^m\right] + m\right)dt + d\hat{W}_t^m,$$

and by definition,

$$\sigma(t, \pi_t^{y,m}) = \partial_2 p(t, x(t, \pi_t^{y,m})) = \partial_2 p(t, x(t, p(t, X_t^m))) = \partial_2 p(t, X_t^m).$$

Applying Ito's formula (see, e.g., Kuo, 2006, Theorem 7.4.3) to $\pi_t^{y,m} = p(t, X_t^m)$ and using the expressions for the partial derivatives above, we get[9]

$$\begin{aligned}
d\pi_t^{y,m} &= \left(\partial_1 p(t, X_t^m) + \partial_2 p(t, X_t^m)\left(\mathrm{E}\left[\theta \mid \mathcal{X}_t^m\right] + m\right) + \frac{1}{2}\partial_2^2 p(t, X_t^m)\right)dt + \partial_2 p(t, X_t^m)d\hat{W}_t^m \\
&= \left(\partial_1 p(t, X_t^m) + \partial_2 p(t, X_t^m)\int_{\mathbb{R}} z\mu_{t,X_t^m}(dz) + \frac{1}{2}\partial_2^2 p(t, X_t^m)\right)dt + \partial_2 p(t, X_t^m)d\hat{W}_t^m \\
&\quad + \left(\mathrm{E}\left[\theta \mid \mathcal{X}_t^m\right] + m - \int_{\mathbb{R}} z\mu_{t,X_t^m}(dz)\right)\partial_2 p(t, X_t^m)dt \\
&= \left(\mathrm{E}\left[\theta \mid \mathcal{X}_t^m\right] + m - \int_{\mathbb{R}} z\mu_{t,X_t^m}(dz)\right)\sigma(t, \pi_t^{y,m})dt + \sigma(t, \pi_t^{y,m})d\hat{W}_t^m.
\end{aligned}$$

□

C. Proofs Omitted From Section 4

Lemma 1. *For any prior distribution μ, let $p_0^\mu = \mu([0,\infty))$. Then, $V^\mu(0, p_0^\mu)$ is convex in μ.*

Proof. Let $b_h, b_l \colon \mathbb{R}_+ \to \mathbb{R}$ be continuous such that $b_l(t) < b_h(t)$ for all $t \in \mathbb{R}_+$. For any prior distribution μ, let

$$\tau^\mu = \inf\{t \in \mathbb{R}_+ \colon X_t^\mu \in \{b_h(t), b_l(t)\}\}$$

denote the first (possibly infinite) time at which X^μ hits one of the boundaries b_h, b_l. Since the distribution of sample paths of X^μ is linear in μ, it follows that $\mathrm{E}\left[\tau^\mu\right]$ is linear in μ. Similarly, the probability of choosing the high action when $\theta^\mu \geq 0$,

$$P\left(\tau^\mu < \infty \wedge X_{\tau^\mu}^\mu = b_h(\tau^\mu) \wedge \theta^\mu \geq 0\right),$$

[9] A useful way of verifying that the first "dt"-term after the second equality cancels is the symbolism $\partial_1 p(t,x) = -\frac{1}{2}z^2\mathbb{R}_+ + (z^0\mathbb{R}_+)(\frac{1}{2}z^2\mathbb{R})$, etc.

and the corresponding expression for b_l and $\theta^\mu < 0$ are linear in μ. Observe that

$$\begin{aligned} V^\mu(0,p_0^\mu) &= \sup_\tau \mathrm{E}\left[g(\pi_\tau) - c\tau\right] \\ &\geq P\left(\tau^\mu < \infty \wedge X^\mu_{\tau^\mu} = b_h(\tau^\mu) \wedge \theta^\mu \geq 0\right) + P\left(\tau^\mu < \infty \wedge X^\mu_{\tau^\mu} = b_l(\tau^\mu) \wedge \theta^\mu < 0\right) - c\mathrm{E}\left[\tau^\mu\right] \\ &= \mathrm{Acc}(\tau^\mu) - c\mathrm{E}\left[\tau^\mu\right], \end{aligned} \tag{3}$$

where equality holds if b_h, b_l are optimal boundaries for the observed process X^μ.

Now fix a prior distribution μ. By Theorem 1(iii), the optimal boundaries β_h^μ, β_l^μ on the belief process are continuous, and x^μ is continuous, and so $b_h(t) := b_h^\mu(t) = x^\mu(t, \beta_h^\mu(t))$ and $b_l(t) := b_l^\mu(t) = x^\mu(t, \beta_l^\mu(t))$ are continuous. If $\mu = \lambda\mu_1 + (1-\lambda)\mu_2$ for two prior distributions μ_1, μ_2 and $\lambda \in [0,1]$, then linearity of the right-hands-side in (3) and equality for μ imply that

$$\begin{aligned} \lambda V^{\mu_1}(0,p_0^{\mu_1}) + (1-\lambda)V^{\mu_2}(0,p_0^{\mu_2}) &\geq \lambda\left(\mathrm{Acc}[\tau^{\mu_1}] - c\mathrm{E}\left[\tau^{\mu_1}\right]\right) + (1-\lambda)\left(\mathrm{Acc}[\tau^{\mu_2}] - c\mathrm{E}\left[\tau^{\mu_2}\right]\right) \\ &= \mathrm{Acc}[\tau^\mu] - c\mathrm{E}\left[\tau^\mu\right] = V^\mu(0, p_0^{\mu_\lambda}). \end{aligned}$$

which proves the claim. □

Proposition 2. *Let μ be a prior distribution, and let $\tilde{\mu}$ be a sign-preserving random shift of μ. Then, for all $p \in (0,1)$, $V(0,p) \geq \tilde{V}(0,p)$. Moreover, $\beta_h(0) \geq \tilde{\beta}_h(0)$ and $\tilde{\beta}_l(0) \geq \beta_l(0)$, and there is $t_0 \in \mathbb{R}_+$ such that $\beta_h(t) \geq \tilde{\beta}_h(t)$ and $\tilde{\beta}_l(t) \geq \beta_l(t)$ for all $t \geq t_0$.*

Proof. For a random variable θ, denote by $M_\theta(a) = \mathrm{E}\left[e^{a\theta}\right]$, $a \in \mathbb{R}$, its moment generating function. If θ has distribution μ, we also write M_μ for its moment generating function. For all $a, x \in \mathbb{R}$, we have

$$M_{\mu_x}(a) = \int_{\mathbb{R}} e^{az}\mu_x(dz) = \frac{\int_{\mathbb{R}} e^{(a+x)z}\mu(dz)}{\int_{\mathbb{R}} e^{xz}\mu(dz)} = \frac{M_\mu(x+a)}{M_\mu(x)},$$

where we write $\mu_x = \mu_{0,x}$.

Now let θ, ξ be independent random variables with distributions μ, ν as in the definition of sign-preserving random shifts. Let $\tilde{\mu} = \mu * \nu$ be the distribution of $\theta + \xi$. We show that $V(0,p) \geq \tilde{V}(0,p)$ for all $p \in [0,1]$. For $p \in \{0,1\}$, we have $V(0,p) = g(p) = \tilde{V}(0,p)$. So assume from now on that $p \in (0,1)$. Let $x = x^\mu(0,p)$ be the observation at time 0 that induces belief p, so that $\mu_x([0,\infty)) = p$. The idea is to show that $\tilde{\mu}_x = \mu_x * \nu_x$, and thus that $\tilde{\mu}_x$ is an sign-preserving random shift μ_x. This reduces the problem to the case $p = p_0$ since for μ_x, p is the belief induced by observing 0 at time 0. Then the statement follows from Lemma 1.

For all $a \in \mathbb{R}$,

$$M_{\mu_x}(a)M_{\nu_x}(a) = \frac{M_\mu(x+a)}{M_\mu(x)}\frac{M_\nu(x+a)}{M_\nu(x)} = \frac{M_{\mu*\nu}(x+a)}{M_{\mu*\nu}(x)} = M_{\tilde{\mu}_x}(a),$$

where the second equality uses that the moment generating function of the convolution of two distributions is the product of their moment generating functions. Since a distribution is uniquely determined by its moment generating function, it follows that $\tilde{\mu}_x = \mu_x * \nu_x$, and so $\tilde{\mu}_x$

is the distribution of the sum of two independent random variables, say θ_x, ξ_x, with distributions μ_x, ν_x. Note that $\theta_x > 0$ if and only if $\theta_x + \xi_x > 0$ since this property holds for θ and ξ and the supports of μ, ν are the same as those of μ_x, ν_x, respectively. This shows that $\tilde{\mu}_x$ is a sign-preserving random shift of μ_x, and reduces the problem to the case $p = p_0$, which follows from Lemma 1.

Using that $V(0,p) \geq \tilde{V}(0,p)$ all $p \in [0,1]$, the definition of the boundaries (cf. Theorem 1(ii)) implies that $\beta_h(0) \geq \tilde{\beta}_h(0)$ and $\tilde{\beta}_l(0) \geq \beta_l(0)$. Let $z_+ = \inf\{z \geq 0\colon \mu([z, z+\epsilon)) > 0 \text{ for all } \epsilon > 0\}$ and $z_- = \sup\{z \leq 0\colon \mu((z-\epsilon, z]) > 0 \text{ for all } \epsilon > 0\}$, and define $\tilde{z}_+, \tilde{z}_-$ analogously. Then, $z_+ \geq \tilde{z}_+$ and $\tilde{z}_- \geq z_-$, and the inequalities are strict if $\mu \neq \tilde{\mu}$. By the remarks at the end of Appendix A, the optimal boundaries for a two-point prior distribution with support $\{\tilde{z}_+, \tilde{z}_-\}$ are closer to $\frac{1}{2}$ than the optimal boundaries for a two-point prior with support $\{z_+, z_-\}$. It thus follows from Theorem 1(iv) that $\lim_{t\to\infty} \beta_h(t) \geq \lim_{t\to\infty} \tilde{\beta}_h(t)$ and $\lim_{t\to\infty} \tilde{\beta}_l(t) \geq \lim_{t\to\infty} \beta_l(t)$, where both equalities are strict if $\mu \neq \tilde{\mu}$. Hence, the last statement follows. □

Corollary 1. *Let μ be a prior distribution supported on two points, and let $\tilde{\mu}$ be a sign-preserving random shift of μ. Then, for all $t \in \mathbb{R}_+$ and $p \in (0,1)$, $V(t,p) \geq \tilde{V}(t,p)$. Moreover, for all $t \in \mathbb{R}_+$, $\beta_h(t) \geq \tilde{\beta}_h(t)$ and $\tilde{\beta}_l(t) \geq \beta_l(t)$.*

Proof. By Proposition 2, $V(0,p) \geq \tilde{V}(0,p)$ for all $p \in [0,1]$, and $\beta_h(0) \geq \tilde{\beta}_h(0)$ and $\tilde{\beta}_l(0) \geq \beta_l(0)$. Theorem 1(i) shows that $V(t,p)$ is constant in t for all p, and $\tilde{V}(t,p)$ is non-increasing in t for all p. Hence, $V(t,p) \geq \tilde{V}(t,p)$ for all $t \geq 0$ and $p \in [0,1]$. Moreover, β_h, β_l are constant, and $\tilde{\beta}_h, \tilde{\beta}_l$ are non-increasing and non-decreasing, respectively, which implies $\beta_h(t) \geq \tilde{\beta}_h(t)$ and $\tilde{\beta}_l(t) \geq \beta_l(t)$ for all $t \in \mathbb{R}_+$. □

Corollary 2. *Let μ be a prior distribution supported on two points, and let $\tilde{\mu}$ be a sign-preserving random shift of μ. Let $\tau, \tilde{\tau}$ be optimal stopping times for $\mu, \tilde{\mu}$ given by Theorem 1(ii). Then,* $\mathrm{Acc}[\tilde{\tau}] \leq \mathrm{Acc}[\tau]$.

Proof. Let β_h, β_l and $\tilde{\beta}_h, \tilde{\beta}_l$ be the optimal boundaries for μ and $\tilde{\mu}$ respectively. By Theorem 1, β_h, β_l are constant, and we also write β_h, β_l for the corresponding constants. Moreover, the optimal boundaries for any two-point prior distribution are symmetric about $\frac{1}{2}$ (Shiryaev, 2008, Section 4.2, Theorem 5 and the remark thereafter), and so $\beta_h = 1 - \beta_l$.[10] Hence, if τ is the optimal stopping time induced by β_h, β_l, then $g(\pi_\tau) = \pi_\tau \vee (1 - \pi_\tau) = \beta_h$ almost surely. By Corollary 1, $\beta_h \geq \tilde{\beta}_h(t) > \frac{1}{2} > \tilde{\beta}_l(t) \geq \beta_l$ for all $t \in \mathbb{R}_+$. Hence, if $\tilde{\tau}$ is the optimal stopping time induced by $\tilde{\beta}_h, \tilde{\beta}_l$, then

$$g(\tilde{\pi}_{\tilde{\tau}}) \leq \tilde{\beta}_h(\tilde{\pi}_{\tilde{\tau}}) \vee (1 - \tilde{\beta}_l(\tilde{\pi}_{\tilde{\tau}})) \leq \beta_h.$$

[10] The symmetry is crucial for the argument. In general, if β_h, β_l and $\tilde{\beta}_h, \beta_l$ are two pairs of boundaries (not necessarily optimal) for two belief processes $\pi, \tilde{\pi}$ such that $\beta_h(t) \geq \tilde{\beta}_h(t) > \frac{1}{2} > \tilde{\beta}_l(t) \geq \beta_l(t)$ for all $t \in \mathbb{R}_+$, it is not true that β_h, β_l induce higher expected accuracy than $\tilde{\beta}_h, \tilde{\beta}_l$, not even if $\pi = \tilde{\pi}$. For example, if β_h is close to 1 and β_l is close to $\frac{1}{2}$, it is very likely that π hits β_l before β_h, and so the high accuracy of decisions at β_h rarely matters. Moving β_h closer to $\frac{1}{2}$ reduces the accuracy of decisions at β_h but increases the probability that π hits β_h before β_l. The second effect can outweigh the first effect.

almost surely. It follows that $g(\pi_\tau) \geq g(\tilde{\pi}_{\tilde{\tau}})$ almost surely. In particular, $\mathrm{Acc}[\tilde{\tau}] \leq \mathrm{Acc}[\tau]$. $\square$

Heuristically, it is clear that if the prior distribution is concentrated on states with high absolute value, then the value is close to 1; likewise, if it is concentrated on states with low absolute value, the value is close to the accuracy achieved by deciding based on the prior belief. The next lemma makes this precise. Via a time change of the underlying Brownian motion, analogous statements hold for small (large) observation cost c.

Lemma 3. *For each $\alpha > 0$, let μ^α be the distribution of $\alpha\theta$. Then, $V^{\mu_\alpha}(0, p_0) \to 1$ for $\alpha \to \infty$ and $V^{\mu_\alpha}(0, p_0) \to g(p_0)$ for $\alpha \to 0$.*

Proof. We write V^α instead of V^{μ_α}, and so forth. First, consider the case $\alpha \to \infty$. Fix $T, x > 0$, and consider the $\mathcal{X}^\alpha$-stopping time τ induced by the first exit from the rectangle $[0, T] \times [-x, x]$. That is,

$$\tau^\alpha = \inf\{t \in \mathbb{R}_+ \colon X_t^\alpha \notin [-x, x]\} \wedge T.$$

Standard estimates for normal distributions show that $\mathrm{Acc}[\tau^\alpha] \to 1$ for $\alpha \to \infty$. Hence, $\liminf_{\alpha\to\infty} V^\alpha(0, p_0) \geq 1 - cT$. Since T was arbitrary, we have $V^\alpha(0, p_0) \to 1$ for $\alpha \to \infty$.

Second, consider the case $\alpha \to 0$. Since μ has bounded support, there is $z \in \mathbb{R}_{++}$ such that the support of μ is contained in $[-z, z]$. Let $\nu = p_0\delta_z + (1 - p_0)\delta_{-z}$. Then, for all $\alpha > 0$, $t \in \mathbb{R}_+$, and $p \in (0, 1)$, $\sigma^{\nu_\alpha}(t, p) \geq \sigma^{\mu_\alpha}(t, p)$ by (2). Hence, for all $\alpha > 0$, $t \in \mathbb{R}_+$, and $p \in (0, 1)$, $V^{\nu_\alpha}(t, p) \geq V^{\mu_\alpha}(t, p)$ (see, e.g., Janson and Tysk, 2003, Theorem 7). But $V^{\nu_\alpha}(0, p_0) \to g(p_0)$ for $\alpha \to 0$ since $\beta_h^{\nu_\alpha}, \beta_l^{\nu_\alpha} \to \frac{1}{2}$, and so the claim follows. $\square$

A straightforward consequence of Lemma 3 is the following. If τ^α is an optimal stopping time for μ^α, then $\mathrm{Acc}[\tau^\alpha] \to 1$ and $\mathrm{E}[\tau^\alpha] \to 0$ for $\alpha \to \infty$.

The next lemma shows that quick decisions necessarily come at the cost of accuracy. More precisely, any stopping time with an expected observation time close to 0 has an expected accuracy close to the accuracy of deciding based on the prior belief.

Lemma 4. *Let μ be any prior distribution with prior belief p_0. Let $\epsilon > 0$. Then, there exists $\delta > 0$ such that for any $\mathcal{X}$-stopping time τ with $\mathrm{E}[\tau] \leq \delta$, it holds that $\mathrm{Acc}(\tau) \leq g(p_0) + \epsilon$.*

Proof. Assume for contradiction that for each $\delta > 0$, there is an $\mathcal{X}$-stopping time τ with $\mathrm{E}[\tau] \leq \delta$ and $\mathrm{Acc}(\tau) > g(p_0) + \epsilon$. For each $\alpha > 0$, define the process $X^\alpha = (X_t^\alpha)_{t\in\mathbb{R}_+}$ by letting $X_t^\alpha = \alpha^{-\frac{1}{2}} X_{\alpha t}$. Then, $X_t^\alpha = \alpha^{\frac{1}{2}}\theta t + \tilde{W}_t$ where $\tilde{W}_t = \alpha^{-\frac{1}{2}} W_{\alpha t}$ is a standard Brownian motion by the scaling property of Brownian motion. But then, $V^\alpha(0, p_0) \to g(p_0)$ for $\alpha \to 0$ by Lemma 3.

Now let $\alpha > 0$ such that $V^\alpha(0, p_0) \leq g(p_0) + \frac{\epsilon}{2}$. By assumption, there is an $\mathcal{X}$-stopping time τ with $\mathrm{E}[\tau] \leq \frac{\alpha\epsilon}{2c}$ and $\mathrm{Acc}(\tau) > g(p_0) + \epsilon$. Since $\mathcal{X}_t^\alpha = \mathcal{X}_{\alpha t}$, $\tau^\alpha = \alpha^{-1}\tau$ is an $\mathcal{X}^\alpha$-stopping time with $\mathrm{E}[\tau^\alpha] \leq \frac{\epsilon}{2c}$ and $\mathrm{Acc}(\tau^\alpha) = \mathrm{Acc}(\tau) > g(p_0) + \epsilon$. But then, $V^\alpha(0, p_0) \geq \mathrm{Acc}(\tau^\alpha) - c\mathrm{E}[\tau^\alpha] > g(p_0) + \frac{\epsilon}{2}$, which contradicts the choice of α. $\square$

Proposition 3. *There exists a prior distribution μ supported on two points and a sign-preserving random shift $\tilde{\mu}$ of μ such that $\mathrm{E}[\tau] < \mathrm{E}[\tilde{\tau}]$ when $\tau, \tilde{\tau}$ denote optimal stopping times for $\mu, \tilde{\mu}$, respectively. Likewise, there exist such $\mu, \tilde{\mu}$ such that the opposite inequality holds.*

Proof. For $z \in \mathbb{R}_{++}$, let $\mu^z = \frac{1}{2}\delta_z + \frac{1}{2}\delta_{-z}$, and so the prior belief $p_0^z = \frac{1}{2}$. Denote by τ^z the corresponding optimal $\mathcal{X}^z$-stopping time defined by Theorem 1(ii).

First, we show that a sign-preserving random shift can increase the expected observation time under optimal stopping. Let $\frac{1}{6} > \epsilon > 0$. Let $z \in \mathbb{R}_{++}$ such that $V^z(0, \frac{1}{2}) \geq 1 - \epsilon$, which exists since $V^z(0, \frac{1}{2}) \to 1$ for $z \to \infty$ by Lemma 3. Let $\delta > 0$ such that for any $\mathcal{X}^z$-stopping time τ with $\mathrm{E}[\tau] \leq \delta$, we have $\mathrm{Acc}(\tau) \leq \frac{1}{2} + \epsilon$, which exists by Lemma 4.

Now let $z' \in \mathbb{R}_{++}$, $z' > z$, such that $\mathrm{E}\left[\tau^{z'}\right] \leq \frac{\delta}{2}$, which exists by (the remarks after) Lemma 3. Let $\tilde{\mu} = \mu^{z'} * \mu^{z'-z}$. More concretely, $\tilde{\mu} = \frac{1}{4}\delta_{z''} + \frac{1}{4}\delta_z + \frac{1}{4}\delta_{-z} + \frac{1}{4}\delta_{-z''}$, where $z'' = z' + (z' - z)$. From (2), it follows that $\tilde{\sigma}(t,p) \geq \sigma^z(t,p)$ for all $t \in \mathbb{R}_+$ and $p \in (0,1)$, and so $\tilde{V}(0, \frac{1}{2}) \geq V^z(0, \frac{1}{2}) \geq 1 - \epsilon$ (see, e.g., Janson and Tysk, 2003, Theorem 7). Denote by $\tilde{\beta}_h, \tilde{\beta}_l$ and $\tilde{\tau}$ the optimal boundaries and the optimal $\tilde{\mathcal{X}}$-stopping time for $\tilde{\mu}$ defined by Theorem 1(ii). Assume for contradiction that $\mathrm{E}[\tilde{\tau}] \leq \frac{\delta}{2}$. The corresponding optimal boundaries $\tilde{b}_h, \tilde{b}_l \colon \mathbb{R}_+ \to \mathbb{R}$ on the observed process $\tilde{X}$ are given by $\tilde{b}_h(t) = \tilde{x}(t, \tilde{\beta}_h(t))$ and $\tilde{b}_l(t) = \tilde{x}(t, \tilde{\beta}_l(t))$. Let $\tau^* = \inf\{t \in \mathbb{R}_+ \colon X_t^z \in \{\tilde{b}_h(t), \tilde{b}_l(t)\}\}$ and $\tau^{**} = \inf\{t \in \mathbb{R}_+ \colon X_t^{z''} \in \{\tilde{b}_h(t), \tilde{b}_l(t)\}\}$ be the $\mathcal{X}^z$ and $\mathcal{X}^{z''}$-stopping times induced by these boundaries. Then, since the expected observation time and the expected accuracy for fixed boundaries on the observed process are linear in the prior distribution (cf. the proof of Lemma 1), and $\tilde{\mu} = \frac{1}{2}\mu^z + \frac{1}{2}\mu^{z''}$, it follows that

$$\mathrm{E}[\tilde{\tau}] = \frac{1}{2}\mathrm{E}[\tau^*] + \frac{1}{2}\mathrm{E}[\tau^{**}] \quad \text{and} \quad \mathrm{Acc}[\tilde{\tau}] = \frac{1}{2}\mathrm{Acc}[\tau^*] + \frac{1}{2}\mathrm{Acc}[\tau^{**}].$$

Hence,

$$\mathrm{E}[\tau^*] \leq 2\mathrm{E}[\tilde{\tau}] \leq \delta \quad \text{and} \quad \mathrm{Acc}[\tau^*] \geq 2\mathrm{Acc}[\tilde{\tau}] - 1 \geq 1 - 2\epsilon,$$

which contradicts the choice of δ. □

D. Omitted Proofs From Section 5

We prove the statements in Section 5 along with some auxiliary statements. The first shows that the value of the optimal stopping problem in Section 4 depends continuously on the prior distribution.

Lemma 5. *For all $t \in \mathbb{R}_+$ and $p \in (0,1)$, $V^\mu(t,p)$ is continuous in μ with respect to the weak topology on the set of prior distributions.*

Proof. Let $t \in \mathbb{R}_+$ and $p \in (0,1)$ be fixed throughout. For $T \in \mathbb{R}_+ \cup \{\infty\}$ and a prior distribution μ, denote by $\mathcal{T}_T^\mu$ the set of $\mathcal{X}^\mu$-stopping times τ such that $\tau \leq T$. Let

$$V_T^\mu(t,p) = \sup_{\tau \in \mathcal{T}_T^\mu} \mathrm{E}\left[g(\pi_{t+\tau}^{\mu,t,p}) - c\tau\right]$$

be the value of the optimal stopping problem in Section 4 restricted to stopping times bounded by T. For $T \in \mathbb{R}_+$, it follows from Coquet and Toldo (2007, Theorem 5) (applied to the time-restricted belief process $(\pi^{\mu,t,p}_{t+s})_{s\in[0,T]}$) that $V^\mu_T(t,p)$ is a continuous function of μ with respect to the weak topology on the set of prior distributions.

If τ is an $\mathcal{X}^\mu$-stopping time, then since g is bounded by 1,

$$\mathrm{E}\left[g(\pi^{\mu,t,p}_{t+\tau}) - c\tau\right] \leq \mathrm{E}\left[g(\pi^{\mu,t,p}_{t+(\tau\wedge T)}) - c(\tau \wedge T)\right] + P\left(\tau \geq T\right) \leq V^\mu_T(t,p) + P\left(\tau \geq T\right).$$

If $\tau \in \mathcal{T}^\mu_\infty$ attains the supremum on the right-hand side above, then $P\left(\tau \geq T\right) \leq \frac{1}{2cT}$ since $V^\mu_\infty(t,p) = V^\mu(t,p) \geq \frac{1}{2}$. It follows that $V^\mu_T(t,p)$ converges to $V^\mu(t,p)$ as $T \to \infty$ uniformly in μ. Hence, $V^\mu(t,p)$ is continuous in μ. □

Proposition 9 (Ekström and Vaicenavicius, 2015, Proposition 4.8)**.** *For each prior distribution μ, the data $V^\mu, \beta^\mu_h, \beta^\mu_l$ is a solution to the following free boundary problem.*

$$\begin{aligned} &\partial_1 V^\mu(t,p) + \frac{\sigma^\mu(t,p)^2}{2}\partial_2^2 V^\mu(t,p) - c = 0 && \beta^\mu_l(t) < p < \beta^\mu_h(t) \\ &V^\mu(t,p) = g(p) && p \notin (\beta^\mu_l(t), \beta^\mu_h(t)) \end{aligned}$$

Moreover, $V^\mu(t,p)$ is C^1 on $(0,1)$ for each $t \in \mathbb{R}_+$.

Proposition 4. *For any prior distribution μ and any distribution Γ of the manipulation cost satisfying the above assumptions, there exists an equilibrium $(\beta_h, \beta_l, \gamma_0)$.*

Proof. For $q \in [0,1]$, denote by $\mu^q = \mu * (q\delta_{\overline{m}} + (1-q)\delta_0)$ be the manipulated prior distribution if Bob manipulates with probability q. Let β^q_h, β^q_l and b^q_h, b^q_l be the optimal boundaries for the belief process and the observed process for the prior distribution μ^q defined by Theorem 1(ii). Similarly, we use throughout a superscript q for objects corresponding to the prior distribution μ_q. The proof has two parts: (i) the optimal boundaries on the observed process depend continuously on Bob's manipulation probability (in a sense made precise below), and (ii) Bob's manipulation probability depends continuously on the boundaries on the observed process. Then, the equilibrium existence follows from a fixed-point argument.

The following four claims establish that for each $T \in \mathbb{R}_+$, $b^q_h|_{[0,T]}$ is a continuous function of q in the topology of uniform convergence for functions on $[0,T]$. The analogous statement holds for $b^q_l|_{[0,T]}$.

Claim 1. The family of functions $(\beta^q_h|_{[0,T]})_{q\in[0,1]}$ is Hölder-$\frac{1}{2}$ equicontinuous.

Proof. We have to show that there is $C > 0$ such that for all $s,t \in [0,T]$ and $q \in [0,1]$, $|\beta^q_h(s) - \beta^q_h(t)| \leq C|s-t|^{\frac{1}{2}}$. Indeed, $\sigma^q(t,p)$ is bounded away from 0 uniformly for $t \in [0,2T]$, $\beta^q_l(t) < p < \beta^q_h(t)$, and $q \in [0,1]$ since $\beta^q_h(0)$ is bounded away from 1 and $\beta^q_l(0)$ is bounded away from 0 uniformly in q (by Proposition 2), and for each q, β^q_h is non-increasing and β^q_l is non-decreasing (by Theorem 1 (iii)). A routine argument shows that $\partial_2^2 V^q(t,p)$ is bounded above uniformly for $t \in [0,T]$, $\beta^q_l(t) < p < \beta^q_h(t)$, and $q \in [0,1]$, say, $\partial_2^2 V^q(t,p) \leq B$ for all such t, p, q.[11] Since the support of μ is bounded by assumption, $\sigma^q(t,p)$ is also bounded above. Then, the upper bound on $\partial_2^2 V^q(t,p)$ and Proposition 9 imply that there is $A \in \mathbb{R}_+$ such that $\partial_1 V^q(t,p) \geq -A$ for all $t \in [0,T]$, $\beta_l^q(t) < p < \beta_h^q(t)$, and $q \in [0,1]$. Second, since $\partial_1 V^q(t,q) \leq 0$ (by Theorem 1 (i)) and $\sigma^q(t,p)$ is bounded above, it follows from Proposition 9 that there is $\tilde{c} \in \mathbb{R}_{++}$ such that $\partial_2^2 V^q(t,p) \geq \tilde{c}$ for $t \in [0,T]$, $\beta_l^q(t) < p < \beta_h^q(t)$, and $q \in [0,1]$.

Then, let $0 \leq s \leq t \leq T$. Using that $\partial_2 V^q(s, \beta_h^q(s)) = g'(\beta_h^q(s)) = -1$ by the last part of Proposition 9 and integrating, we have that

$$\begin{aligned} V^q(s, \beta_h^q(t)) &= V^q(s, \beta_h^q(s)) + \int_{\beta_h^q(s)}^{\beta_h^q(t)} \left(\partial_2 V^q(s, \beta_h^q(s)) + \int_{\beta_h^q(s)}^{p'} \partial_2^2 V^q(s, p'') dp'' \right) dp' \\ &\geq \beta_h^q(t) + \frac{\tilde{c}}{2} (\beta_h^q(s) - \beta_h^q(t))^2. \end{aligned}$$

Moreover,

$$V^q(s, \beta_h^q(t)) \leq V^q(t, \beta_h^q(t)) + (t-s)A = \beta_h^q(t) + (t-s)A.$$

Hence, $\beta_h^q(s) - \beta_h^q(t) \leq (\frac{2A}{\tilde{c}}(t-s))^{\frac{1}{2}}$. ◻

Claim 2. For each $t \in [0,T]$, $\beta_h^q(t)$ is continuous in q.

Proof. Fix $t \in [0,T]$. Let $q_0 \in [0,1]$, and let $(q_n)_{n \in \mathbb{N}} \subset [0,1]$ such that $q_n \to q_0$ as $n \to \infty$. Let $b_n = \beta_h^{q_n}(t)$, and assume for contradiction that $b_n \not\to b_0$. By passing to a subsequence, we may assume that $b_n \to b^* \neq b_0$.

Case 1. Suppose $b^* < b_0$. Then, $V^{q_0}(t, b^*) > b^*$ by definition of b_0, and so by continuity of $V^{q_0}(t,p)$ in p, $V^{q_0}(t, \tilde{b}) > g(\tilde{b})$ for some $\tilde{b} > b^*$. For n large enough, $b_n \leq \tilde{b}$, and so $V^{q_n}(t, \tilde{b}) = g(\tilde{b})$, which contradicts that $V^q(t, \tilde{b})$ is continuous in q as asserted by Lemma 5.

Case 2. Suppose $b^* > b_0$. As in the proof of Claim 1, one shows that there are $\tilde{c} > 0$ and $\delta > 0$ such that $\partial_2^2 V^{q_n}(t,p) \geq \tilde{c}$ for all $p \in [b_0, b_0 + \delta]$ and n large enough. Hence, again as in the proof of Claim 1, we get that

$$V^{q_n}(t, b_0) \geq b_0 + \frac{\tilde{c}}{2}(b_n - b_0)^2 = V^{q_0}(t, b_0) + \frac{\tilde{c}}{2}(b_n - b_0)^2$$

for n large enough. This contradicts that $V^q(t, b_0)$ is continuous in q as asserted by Lemma 5.

◻

Claim 3. $\beta_h^q|_{[0,T]}$ is continuous in q in the topology of uniform convergence on $[0,T]$.

[11] Assume that $\sigma^q(t,p) \geq B'$ for all $t \in [0, 2T]$, $\beta_l^q(t) < p < \beta_h^q(t)$, and $q \in [0,1]$. Fix t, p, q in this region. Since $\partial_2^2 V^q(t,p')$ is continuous in p' (see, e.g., Strulovici and Szydlowski, 2015), there is $\rho > 0$ so that $\partial_2^2 V^q(t,p') \geq \frac{1}{2}\partial_2^2 V^q(t,p)$ for all $p' \in [p-\rho, p+\rho]$. For the process $\pi^{t,p}$, consider the stopping time $\tau = \inf\{s \in \mathbb{R}_+ \colon \pi_{t+s}^{t,p} \in \{p+2\rho, p-2\rho\}\}$. Then, using Theorem 1.1 of Geiß and Manthey (1994) and the fact that the expected time for a standard Brownian motion starting at 0 to hit $\{-2\rho, 2\rho\}$ is $4\rho^2$, it follows that $\mathrm{E}[\tau] \leq \frac{4\rho^2}{B'^2}$. We moreover have that $V^q(t, p+2\rho) \geq V^q(t,p) + 2\rho\partial_2 V^q(t,p) + \frac{\rho^2}{2}\partial_2^2 V^q(t,p)$ and similarly for $V^q(t, p-2\rho)$. Using the stopping time $\tau \wedge T$, we have that $V^q(t,p) \geq V^q(t,p) + \frac{\rho^2}{2}\partial_2^2 V^q(t,p) - \frac{4c\rho^2}{B'^2} + o(\rho^2)$, and so letting ρ go to 0, $\partial_2^2 V^q(t,p) \leq 8cB'^2$. (Here, the $o(\rho^2)$ comes from the fact that $P(\tau \geq T)$ goes to 0 as ρ goes to 0.)

Proof. The family $(\beta_h^q|_{[0,T]})_{q\in[0,1]}$ is Hölder-$\frac{1}{2}$ equicontinuous by Claim 1, and so, in particular, uniformly equicontinuous. Moreover, for each $t \in [0,T]$, $\beta_h^q(t)$ is continuous in q by Claim 2. It is an exercise in basic calculus to show that together, these statements imply continuity in the topology of uniform convergence on the compact set $[0,T]$. □

Claim 4. $b_h^q|_{[0,T]}$ is continuous in q in the topology of uniform convergence on $[0,T]$.

Proof. By definition, $b_h^q(t) = x^q(t, \beta_h^q(t))$. Moreover, $x^q(t,p)$ is continuous in t,p,q jointly, and so for each $\rho > 0$, it is Lipschitz continuous when restricted to $t \in [0,T]$, $p \in [\rho, 1-\rho]$, and $q \in [0,1]$. Choosing ρ such that $\rho \le \beta_l^q(t) < \beta_h^q(t) \le 1-\rho$ for all $t \in [0,T]$ and $q \in [0,1]$, the claim follows from Claim 3. □

For each $q \in [0,1]$, let γ_0^q be Bob's optimal cutoff when Alice's boundaries on the observed process are b_h^q, b_l^q. That is, for $\tau^q = \inf\{t \in \mathbb{R}_+ \colon X_t \in \{b_h^q(t), b_l^q(t)\}\}$ and $\tilde{\tau}^q = \inf\{t \in \mathbb{R}_+ \colon X_t + \overline{m}t \in \{b_h^q(t), b_l^q(t)\}\}$, $\gamma_0^q = P\left(X_{\tilde{\tau}^q} + \overline{m}\tilde{\tau}^q = b_h^q(\tilde{\tau}^q)\right) - P\left(X_{\tau^q} = b_h^q(\tau^q)\right)$ is the probability that Alice takes the high action if Bob manipulates minus that probability if Bob does not manipulate.[12] We claim that γ_0^q is a continuous function of q. We show that $P\left(X_{\tau^q} = b_h^q(\tau^q)\right)$ is continuous in q, and the proof for $P\left(X_{\tilde{\tau}^q} + \overline{m}\tilde{\tau}^q = b_h^q(\tilde{\tau}^q)\right)$ is analogous. Fix $q_0 \in [0,1]$ and $\epsilon > 0$, and let $(q_n)_{n\in\mathbb{N}}$ be a sequence converging to q_0. Let $T \in \mathbb{R}_+$ such that $P\left(\tau^{q_n} \le T\right) \ge 1 - \frac{\epsilon}{3}$ for all $n \in \mathbb{N} \cup \{0\}$, which exists by the same argument as in the proof of Lemma 5. By Claim 4, $b_h^{q_n}|_{[0,T]}$ converges to $b_h^{q_0}|_{[0,T]}$ uniformly on $[0,T]$. Hence, for n large enough, $|P\left(X_{\tau^{q_n}\wedge T} = b_h^{q_n}(\tau^{q_n}\wedge T)\right) - P\left(X_{\tau^{q_0}\wedge T} = b_h^{q_0}(\tau^{q_0}\wedge T)\right)| \le \frac{\epsilon}{3}$. Together, we have that $|P\left(X_{\tau^{q_n}} = b_h^{q_n}(\tau^{q_n})\right) - P\left(X_{\tau^{q_0}} = b_h^{q_0}(\tau^{q_0})\right)| \le \epsilon$ proving the claim.

To conclude, we observe that since γ_0^q is a continuous function of q and Γ is absolutely continuous with full support on $[0,1]$, Bob's manipulation probability $\phi(q) = P\left(\gamma \le \gamma_0^q\right)$ is a continuous function of q, and hence ϕ has a fixed-point $q^* \in [0,1]$ by the intermediate value theorem. By the construction of ϕ, $(\beta_h^{q^*}, \beta_l^{q^*}, \gamma_0^{q^*})$ is an equilibrium, which finishes the proof. □

Proposition 5. *Let μ be any prior distribution, let β_h, β_l be the optimal boundaries for μ as defined in Theorem 1(ii), and assume that $\beta_l(0) < p_0 < \beta_h(0)$. Then, in any equilibrium, (i) Alice almost surely observes past time* 0*, and (ii) Bob manipulates with probability strictly between* 0 *and* 1*.*

Proof. Let $(\beta_h^*, \beta_l^*, \gamma_0^*)$ be an equilibrium for the prior distribution μ.

We prove (ii) first. Since Γ has full support, it suffices to show that $\gamma_0^* \in (0,1)$. Assume for contradiction that $\gamma_0^* = 0$. Then, Alice observes the (non-manipulated) process X, and so $\beta_h^* = \beta_h$ and $\beta_l^* = \beta_l$. Denote by $b_h(t) = x(t, \beta_h(t))$ and $b_l(t) = x(t, \beta_l(t))$ the corresponding optimal boundaries for the observed process X. If Bob were to manipulate, Alice would observe $\tilde{X} = (\tilde{X}_t)_{t\in\mathbb{R}_+}$ with $\tilde{X}_t = X_t + \overline{m}t$. Let

$$\tau = \inf\{t \in \mathbb{R}_+ \colon X_t \in \{b_h(t), b_l(t)\}\} \quad \text{and} \quad \tilde{\tau} = \inf\{t \in \mathbb{R}_+ \colon \tilde{X}_t \in \{b_h(t), b_l(t)\}\}.$$

[12] $\tau^q, \tilde{\tau}^q$ have finite expectations since b_h^q, b_l^q are optimal for the prior distribution μ_q, and so Alice's expected observation time is finite.

It suffices to show that

$$P\left(\tilde{X}_{\tilde{\tau}}=b_h(\tilde{\tau})\right)>P\left(X_\tau=b_h(\tau)\right),$$

since then manipulating is a best response for Bob for $\gamma>0$ small enough. Since $\beta_l(0)<p_0<\beta_h(0)$ by assumption, we have that $\tau\not\equiv 0$ and $P\left(X_\tau=b_h(\tau)\right)\in(0,1)$. Observe that $\{X_\tau=b_h(\tau)\}\subset\{\tilde{X}_{\tilde{\tau}}=b_h(\tilde{\tau})\}=\{b_h(t)-X_t\leq\overline{m}t \text{ for some } t\in\mathbb{R}_+\}$ as events. Then,

$$P\left(\tilde{X}_{\tilde{\tau}}=b_h(\tilde{\tau})\right)-P\left(X_\tau=b_h(\tau)\right)=P\left(\tilde{X}_{\tilde{\tau}}=b_h(\tilde{\tau})\wedge X_\tau=b_l(\tau)\right)>0,$$

where the inequality follows from standard estimates for Brownian motion. This proves the claim.

If $\gamma_0^*=1$, we also have $\beta_h^*=\beta_h$ and $\beta_l^*=\beta_l$. Denoting by π^* Alice's belief process in equilibrium and letting $\tau^*=\inf\{t\in\mathbb{R}_+\colon\pi_t^*\in\{\beta_h(t),\beta_l(t)\}\}$, we have that $P\left(\pi^*_{\tau^*}=\beta_h(\tau^*)\right)\in(0,1)$ since $\beta_l(0)<p_0<\beta_h(0)$. Bob's expected payoff for $\overline{m}$ is bounded from above by $1-\gamma$, and so his expected payoff for $\overline{m}$ is smaller than his expected payoff for 0 for γ close to 1. Hence, $\gamma_0^*=1$ is not a best response, which is a contradiction.

Turning to the proof of (i), observe that it suffices to show that $\beta_l^*(0)<p_0<\beta_h^*(0)$. Assuming for contradiction that $p_0\geq\beta_h^*(0)$, it follows that $X_0=0\geq b_h^*(0)$, where $b_h^*(t)=x^*(t,\beta_h^*(t))$ is the optimal boundary for the observed process in equilibrium.[13] Hence, Alice almost surely stops at time 0 and takes the high action, and so Bob's payoff is $1-\gamma\mathbf{1}_{[0,\gamma_0^*]}(\gamma)$. In particular, his payoff is not non-increasing in γ since $\gamma_0^*\in(0,1)$ by (ii). This contradicts that γ_0^* is a best response. The proof for the case $p_0\leq\beta_l^*(0)$ is analogous. □

Proposition 7. *Let $\mu=p_0\delta_1+(1-p_0)\delta_{-1}$ be a prior distribution supported on two points, and let β_h,β_l be the constant optimal boundaries for μ, and assume that $\beta_l<p_0<\beta_h$. Let $(\beta_h^*,\beta_l^*,\gamma_0^*)$ be any equilibrium, and denote by τ^* the induced stopping time for the belief process in equilibrium. Then, Bob's cutoff γ_0^* and his manipulation probability $P\left(\gamma\leq\gamma_0^*\right)$, and Alice's expected observation time $\mathrm{E}\left[\tau^*\right]$ go to 0 as p_0 goes to β_h or β_l.*

Proof. Let $(\beta_h^*,\beta_l^*,\gamma_0^*)$ be an equilibrium for the prior distribution μ. Observe that the distribution of $\theta+\overline{m}\mathbf{1}_{[0,\gamma_0^*]}(\gamma)$ is a non-trivial (by Proposition 5(ii)) sign-preserving random shift of μ. Hence, by Proposition 2 and Proposition 5(i), we have $\beta_h\geq\beta_h^*(0)>p_0>\beta_l^*(0)\geq\beta_l$. Moreover, by Theorem 1(iii), $\beta_h^*(0)>\frac{1}{2}>\beta_l^*(0)$. Denote by π^* Alice's belief process in equilibrium, and let $\tau^*=\inf\{t\in\mathbb{R}_+\colon\pi_t^*\in\{\beta_h^*(t),\beta_l^*(t)\}\}$. We claim that $P\left(\pi^*_{\tau^*}=\beta_h^*(\tau^*)\mid\gamma>\gamma_0^*\right)$ goes to 1 as p_0 goes to β_h—in words, the probability of Alice taking the high action conditional on Bob not manipulating goes to 1 as the prior belief goes to β_h. The distribution of π^* conditional on the event $\{\gamma>\gamma_0^*\}$ equals the distribution of $\pi^{\gamma_0^*,0}$ in Lemma 2, and by that lemma,

$$d\pi_t^{\gamma_0^*,0}=\xi_t dt+\sigma^{\gamma_0^*}(t,\pi_t^{\gamma_0^*,0})d\hat{W}_t$$

for some process $\xi=(\xi_t)_{t\in\mathbb{R}_+}$ measurable with respect to X such that ξ_t is lower bounded by $(-1)\sigma^{\gamma_0^*}(t,\pi_t^{\gamma_0^*,0})\geq-1$ (where the -1 comes from the fact that the manipulated prior

[13] More precisely, $x^*=x^{\mu^*}$, where μ^* is the distribution of $\theta+\overline{m}\mathbf{1}_{[0,\gamma_0^*]}(\gamma)$.

distribution has support contained in $[-1,\infty)$ and we use (2) to bound $\sigma^{\gamma_0^*}$). Let $\tilde{\pi}=(\tilde{\pi}_t)_{t\in\mathbb{R}_+}$ such that

$$d\tilde{\pi}_t = -dt + \sigma^{\gamma_0^*}(t,\tilde{\pi}_t)d\hat{W}_t.$$

Then, $P\left(\pi_t^{\gamma_0^*,0} \geq \tilde{\pi}_t \text{ for all } t\in\mathbb{R}_+\right) = 1$ (see, e.g., Geiß and Manthey, 1994, Theorem 1.1). Letting $\tilde{\tau}=\inf\{t\in\mathbb{R}_+ : \tilde{\pi}_t \in \{\beta_h^*(0),\frac{1}{2}\}\}$ and using that $\sigma^{\gamma_0^*}(t,p)$ is bounded above and away from 0 uniformly on $\mathbb{R}_+\times[\beta_l,\beta_h]$ and uniformly in $\gamma_0^*\in[0,1]$, we have that as p_0 goes to β_h, $P\left(\tilde{\pi}_{\tilde{\tau}}=\beta_h^*(0)\right)$ goes to 1, and so the probability that $\pi^{\gamma_0^*,0}$ hits β_h^* before β_l^* goes to 1 as p_0 goes to β_h. Hence, for any $\gamma_0\in(0,1]$, $\overline{m}$ is not a best response when $\gamma\geq\gamma_0$ for p_0 close to β_h. It follows that γ_0^* goes to 0 as p_0 goes to β_h, and so since Γ is continuous, it follows that $P\left(\gamma\leq\gamma_0^*\right)$ goes to 0 as p_0 goes to β_h.

The proof for the case that p_0 goes to β_l is similar. □

Proposition 8. *Let μ be any prior distribution, and let $\Gamma,\tilde{\Gamma}$ be two distributions for the manipulation cost such that Γ stochastically dominates $\tilde{\Gamma}$. Assume that under Γ, there is an equilibrium with cutoff γ_0^* where Bob manipulates with probability $q^*=F_\Gamma(\gamma_0^*)$. Then, under $\tilde{\Gamma}$, there is an equilibrium where Bob manipulates with probability at least q^*.*

Proof. Similar to the proof of Proposition 4, define the function $\Phi_\Gamma\colon[0,1]\to[0,1]$ as follows. For $q\in[0,1]$, let β_h^q,β_l^q be Alice's optimal boundaries if Bob manipulates with probability q—that is, the optimal boundaries for the manipulated prior distribution $\mu*((1-q)\delta_0+q\delta_{\overline{m}})$ as defined in Theorem 1(ii). Then, let $\gamma_0^q\in[0,1]$ be Bob's unique best response to β_h^q,β_l^q, and define $\Phi_\Gamma(q)=F_\Gamma(\gamma_0^q)$. Hence, assuming that Alice expects Bob to manipulate with probability q, γ_0^q is the probability that Alice takes the high action if Bob plays $\overline{m}$ minus the probability that Alice takes the high action if Bob plays 0, and $\Phi_\Gamma(q)$ is the that Bob manipulates with cutoff γ_0^q under Γ. We show in the proof of Proposition 4 that Φ_Γ is continuous.

Observe that the maps $q\mapsto(\beta_h^q,\beta_l^q)$ and $(\beta_h^q,\beta_l^q)\mapsto\gamma_0^q$ defined above do not depend on Γ. The map $\gamma_0^q\mapsto\Phi_\Gamma(q)=F_\Gamma(\gamma_0^q)$ is decreasing in the stochastic dominance order. Hence, since Γ stochastically dominates $\tilde{\Gamma}$, $\Phi_\Gamma(q)\leq\Phi_{\tilde{\Gamma}}(q)$ for all $q\in[0,1]$. Since under Γ, there is an equilibrium where Bob manipulates with probability q^*, we have $\Phi_\Gamma(q^*)=q^*$, and so $\Phi_{\tilde{\Gamma}}(q^*)\geq q^*$. By the intermediate value theorem, there is $\tilde{q}^*\in[q^*,1]$ with $\Phi_{\tilde{\Gamma}}(\tilde{q}^*)=\tilde{q}^*$, and so under $\tilde{\Gamma}$, there is an equilibrium $(\tilde{\beta}_h^{\tilde{q}^*},\tilde{\beta}_l^{\tilde{q}^*},\tilde{\gamma}_0^*)$, $\tilde{\gamma}_0^*=F_{\tilde{\Gamma}}^{-1}(\tilde{q}^*)$, where Bob manipulates with probability $\tilde{q}^*\geq q^*$. □

References

K. J. Arrow, D. Blackwell, and M. A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica, Journal of the Econometric Society*, pages 213–244, 1949.

S. Auster, Y.-K. Che, and K. Mierendorff. Prolonged learning and hasty stopping: The Wald problem with ambiguity. *American Economic Review*, 114(2):426–461, 2024.

A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering Fundamentals of Stochastic Filtering Fundamentals of Stochastic Filtering*, volume 60 of *Stochastic Modelling Stochastic Modelling and Applied Probability*. Springer, 2009.

Y. Benkler, R. Faris, and H. Roberts. *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford University Press, 2018.

D. Cetemen and C. Margaria. Dynamic signaling in Wald options. 2023.

J. Chan, A. Lizzeri, W. Suen, and L. Yariv. Deliberating collective decisions. *The Review of Economic Studies*, 85(2):929–963, 2018.

M. Charikar, J. Steinhardt, and G. Valiant. Learning from untrusted data. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 47–60, 2017.

Y.-K. Che and K. Mierendorff. Optimal dynamic allocation of attention. *American Economic Review*, 109(8):2993–3029, 2019.

H. Chernoff. *Sequential analysis and optimal design*. CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, 1972.

F. Coquet and S. Toldo. Convergence of values in optimal stopping and convergence of optimal stopping times. *Electronic Journal of Probability*, 12:207–228, 2007.

B. Daley and B. Green. Waiting for news in the market for lemons. *Econometrica*, 80(4): 1433–1504, 2012.

C. Dellarocas. Strategic manipulation of internet opinion forums: Implications for consumers and firms. *Management science*, 52(10):1577–1593, 2006.

I. Diakonikolas, G. Kamath, D. Kane, J. Li, A. Moitra, and A. Stewart. Robust estimators in high-dimensions without the computational intractability. *SIAM Journal on Computing*, 48 (2):742–864, 2019.

F. Dilmé. Dynamic quality signaling with hidden actions. *Games and Economic Behavior*, 113: 116–136, 2019.

M. Ekmekci and L. Maestri. Wait or act now? Learning dynamics in stopping games. *Journal of Economic Theory*, 205:105541, 2022.

M. Ekmekci, L. Gorno, L. Maestri, J. Sun, and D. Wei. Learning from manipulable signals. *American Economic Review*, 112(12):3995–4040, 2022.

E. Ekström and J. Vaicenavicius. Bayesian sequential testing of the drift of a Brownian motion. *ESAIM: Probability and Statistics*, 19:626–648, 2015.

D. Fudenberg, P. Strack, and T. Strzalecki. Speed, accuracy, and the optimal timing of choices. *American Economic Review*, 108(2):3651–3684, 2018.

C. Geiß and R. Manthey. Comparison theorems for stochastic differential equations in finite and infinite dimensions. *Stochastic Processes and their Applications*, 53(1):23–35, 1994.

J. Glazer, H. Herrera, and M. Perry. Fake reviews. *The Economic Journal*, 131(636):1772–1787, 2021.

S. Gryglewicz and A. Kolb. Dynamic signaling with stochastic stakes. *Theoretical Economics*, 17(2):539–559, 2022.

K. Gwiazdon and D. A. Brown. The climate change disinformation campaign: Attacking the common good, advancing the self, and destroying democracy. In *The Routledge Handbook of Applied Climate Change Ethics*, pages 198–213. Routledge, 2023.

J. M. Harrison. *Brownian Models of Performance and Control.* Cambride University Press, 2013.

E. Henry and M. Ottaviani. Research and the approval process: The organization of persuasion. *American Economic Review*, 109(3):911–955, 2019.

S. Janson and J. Tysk. Volatility time and properties of option prices. *The Annals of Applied Probability*, 13(3):890–913, 2003.

H.-H. Kuo. *Introduction to Stochastic Integration.* Springer, 2006.

K. A. Lai, A. B. Rao, and S. Vempala. Agnostic estimation of mean and covariance. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 665–674. IEEE, 2016.

T. L. Lai. On optimal stopping problems in sequential hypothesis testing. *Statistica Sinica*, 7 (1):33–51, 1997.

A. Liang, X. Mu, and V. Syrgkanis. Dynamically aggregating diverse information. *Econometrica*, 90(1):47–80, 2022.

A. McClellan. Experimentation and approval mechanisms. *Econometrica*, 90(5):2215–2247, 2022.

C. Melchior and M. Oliveira. Health-related fake news on social media platforms: A systematic literature review. *New Media & Society*, 24(6):1500–1522, 2022.

D. Orlov, A. Skrzypacz, and P. Zryumov. Persuading the principal to wait. *Journal of Political Economy*, 128(7):2542–2578, 2020.

A. N. Shiryaev. *Optimal Stopping Rules.* Stochastic Modelling and Applied Probability, volume 8. Springer, 2008.

B. Strulovici and M. Szydlowski. On the smoothness of value functions and the existence of optimal strategies in diffusion models. *Journal of Economic Theory*, 159:1016–1055, 2015.

A. Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16 (2):117–186, 1945.

A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, pages 326–339, 1948.

Symbol	Name	Mathematical object
$(\Omega, \mathcal{F}, P)$		probability space
$W = (W_t)_{t\in\mathbb{R}_+}$		Brownian motion
θ	state	random variable
p_0, $p(t, x)$	prior & posterior belief	elements of $[0, 1]$
μ, $\mu_{t,x}$	prior & posterior distribution	probability distribution on $\mathbb{R}$
$X = (X_t)_{t\in\mathbb{R}_+}$	observed process	stochastic process
$\mathcal{X} = (\mathcal{X}_t)_{t\in\mathbb{R}_+}$	observed information process	filtration of $\mathcal{F}$
$\pi = (\pi_t)_{t\in\mathbb{R}_+}$	belief process	stochastic process
σ	volatility of the belief process	function $\mathbb{R}_+ \times [0, 1] \to \mathbb{R}_+$
$\{h, l\}$	action set	set
g	stopping payoff	function $[0, 1] \to \mathbb{R}$
c	observation cost per unit of time	element of $\mathbb{R}_{++}$
τ	stopping time	stopping time
V	value function	function $\mathbb{R}_+ \times [0, 1] \to \mathbb{R}$
β_h, β_l	optimal boundaries	functions $\mathbb{R}_+ \to [0, 1]$
γ	manipulation cost	random variable on $\mathbb{R}_+$
Γ	distribution of the manipulation cost	probability distribution on $\mathbb{R}_+$

Table 1: Reference table of mathematical objects and their interpretations.