Sequential Sampling under Adversarial Manipulation*

Florian Brandl — University of Bonn

Xianwen Shi — University of Toronto

April 18, 2026

Abstract

We study a continuous-time game of bad-news learning between two symmetrically uninformed players: a decision-maker who samples to learn a binary state and a manipulator who can adaptively delay the arrival of breakdown signals at a cost in the low state. Equilibria with manipulation have a simple structure: along the no-breakdown history, the decision-maker stops deterministically, and the manipulator randomizes between never manipulating and manipulating at full intensity from a deterministic time until the decision-maker stops. These equilibria exist on a nondegenerate interval of priors; for any fixed prior in this interval, manipulation starts at the same time and the decision maker stops at the same time, so multiplicity is confined to the mixing probability. A Wald equilibrium without manipulation also exists on this interval, and the players rank Wald and manipulation equilibria in opposite order. Manipulation can flip the comparison of sampling procedures: in equilibrium, static sampling can outperform sequential sampling, and a hybrid protocol can outperform both.

*The authors thank Jimmy Chan, Kim-Sau Chung, Laura Doval, Johannes Hörner, Michihiro Kandori, Yingkai Li, Wing Suen, Balázs Szentes, Allen Wong, Yuichi Yamamoto, Xin Zhao, and participants at UIBE, PKU, NTU, UTokyo, NUS, SMU, HKBU, HKU, Beijing BEAT conference and SAET conference at Ischia for valuable comments and suggestions. Part of this research was conducted while the second author visited UIBE, NTU, ZJU and HKU; he gratefully acknowledges their hospitality. Florian Brandl acknowledges support from the German Research Foundation (DFG) under the Excellence Strategy EXC-2047. Xianwen Shi acknowledges financial support from the Social Sciences and Humanities Research Council of Canada through an Insight Grant.

Florian Brandl: fbrandl@uni-bonn.de; Xianwen Shi: xianwen.shi@utoronto.ca.

1 Introduction

A foundational insight of Wald (1945) is that sequential sampling can outperform fixed-sample procedures because it allows a decision-maker to stop once the evidence becomes sufficiently informative. This advantage, however, relies on the evidence process being exogenous rather than strategically responsive to the sampling procedure. We ask how this conclusion changes when adverse signals can be strategically obstructed. Such concerns arise, for example, when a firm assesses a manager's trial run of a new business line, a venture capitalist monitors an unproven startup, or a regulator oversees the rollout of a novel product. In these settings, neither side may know whether the underlying venture is viable, yet the evaluated party can take costly actions to suppress early failure signals. When such interference can be adapted to the evolving sampling process, the same flexibility that makes sequential sampling attractive also creates scope for strategic manipulation.

We model learning with adversarial manipulation as a continuous-time game between a decision-maker, Alice, and a manipulator, Bob. The state is binary and learning is through bad news: in the low state, a breakdown arrives at an exogenous rate and reveals the state; in the high state, no breakdown occurs. Alice pays a flow cost while sampling and chooses when to stop and take a binary action, aiming to match her action to the state. Bob always prefers the high action and can suppress breakdowns through costly manipulation by reducing their arrival rate at a linear cost. Both players are symmetrically uninformed ex ante, Alice's stopping decision is public, and Bob's manipulation is concealed and adapted to Alice's sampling process.

Equilibria of this game have a simple structure. Along the no-breakdown history, Alice stops at a deterministic time, so her policy reduces to a deterministic stopping rule, as in the standard sequential-sampling problem without manipulation. Bob randomizes between two extremal pure strategies: never manipulate, or manipulate at full intensity from a fixed time until Alice exits. Intuitively, if Bob manipulates, he delays costly interference as much as possible, because early manipulation may be wasted if a breakdown occurs before Alice stops. Mixing is necessary for equilibrium since a pure strategy that concentrates manipulation entirely in a block just before Alice's exit time would make earlier stopping optimal.

Given this structure, equilibria fall into four classes: instantaneous-deterrence equilibria, deadline-deterrence equilibria, Wald equilibria, and manipulation equilibria. The first two are deterrence outcomes with no on-path manipulation and are mainly useful for a complete taxonomy. Our main focus is on the latter two classes. The Wald equilibrium provides the natural benchmark without on-path manipulation, while the manipulation equilibrium captures how strategic interference reshapes Alice's learning

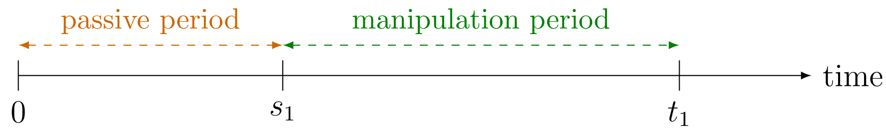

Figure 1: The equilibrium structure of a manipulation equilibrium. Along the no-breakdown history, Alice exits deterministically at t_1. Bob randomizes between never manipulating and manipulating from s_1 to t_1 at full intensity. This splits the sampling region into a passive period and a manipulation period.

dynamics and stopping behavior.

Manipulation equilibria exist on an intermediate range of priors and coexist there with the Wald equilibrium. Their multiplicity is tightly disciplined: for a fixed prior, all manipulation equilibria share the same on-path manipulation interval and differ only in the probability of manipulation. Lower manipulation probabilities are sustained by off-path continuation play that makes continued sampling sufficiently unattractive for Alice.

Manipulation equilibria and Wald equilibria are ranked in opposite ways by the two players. Alice strictly prefers the Wald equilibrium to any manipulation equilibrium, while Bob strictly prefers any manipulation equilibrium to the Wald equilibrium. Equilibrium selection is therefore economically substantive, and becomes a coordination and commitment problem. This equilibrium comparison has a design implication: once the evidence stream is strategically manipulable, sequential sampling need not dominate static sampling in equilibrium. A fixed-sample procedure can deter manipulation and can strictly outperform the worst manipulation equilibrium under sequential sampling. Moreover, a hybrid protocol that combines sequential exploration with a fixed confirmatory phase can outperform both pure procedures.

Finally, we clarify which modeling ingredients are essential. Symmetric uncertainty and adapted manipulation are central. If Bob is informed, the unique equilibrium features immediate exit at time 0; if manipulation effort must be chosen and paid for up front, the unique equilibrium outcome is the no-manipulation Wald outcome. Linear manipulation cost is not the main driver of the results. With slightly convex costs, we construct a nearby manipulation equilibrium with a similar qualitative structure: Alice still stops deterministically along the no-breakdown history, and Bob delays intervention and then increases effort over time.

Taken together, our results yield four lessons. First, in adversarial learning, sequential flexibility can become a liability: the same adaptivity that makes sequential sampling efficient also creates scope for strategic interference. Second, the bad-news learning game delivers a sharp equilibrium structure: along the no-breakdown history, Alice stops at a deterministic time, while Bob mixes once between no manipulation and full-intensity manipulation over a single interval. Third, multiplicity is both limited and economically substantive: manipulation equilibria and Wald equilibria can coexist and are ranked oppositely by the two players, making commitment and coordination central. Fourth, these forces inform procedure choice: when evidence can be strategically obstructed, sequential sampling need not dominate static or hybrid protocols that embed commitment.

1.1 Related Literature

Our paper contributes to several strands of literature at the intersection of sequential analysis, strategic experimentation, and strategic communication.

First, we build on the classical theory of sequential analysis pioneered by Wald (1945), which shows that sequential tests can substantially reduce sample size relative to fixed-sample tests with comparable power.[1] We depart from this literature by allowing a strategic player to interfere with the data-generating process itself. In our setting, the flexibility of sequential sampling may become a source of vulnerability, so that static or hybrid procedures can outperform pure sequential sampling.[2] Our paper also complements recent strategic extensions of the Wald framework, including persuasion (Henry and Ottaviani, 2019), committee deliberation (Chan, Lizzeri, Suen, and Yariv, 2018), and delegation (McClellan, 2022). In contrast to these papers, we study adversarial manipulation of the evidence process itself.

Second, our model is related to the "bad news" learning structure in the strategic experimentation literature, especially Keller and Rady (2015) and Bonatti and Hörner (2017b), where agents learn from the arrival of negative signals.[3] As in Bonatti and Hörner (2017b), private actions can generate equilibrium mixing over extreme choices. Our model shares this feature, but delivers a sharper timing result: the manipulator randomizes only once, between no manipulation and full manipulation over a single interval. More importantly, our environment has asymmetric roles and conflicting objectives, and we use the equilibrium characterization to compare sequential, static, and hybrid sampling procedures.

Third, our paper is related to strategic communication, signaling, and reputation.

[1]The Wald framework is a standard model of dynamic information acquisition; see, for example, Fudenberg, Strack, and Strzalecki (2018) and Che and Mierendorff (2019).

[2]A related insight appears in Adusumilli (2026), who show that in A/B testing the regret-minimizing policy need not adapt to past outcomes.

[3]See also Hörner and Samuelson (2026), who use a related bandit framework to study how disclosing expert information can create perverse incentives.

In cheap-talk models (Crawford and Sobel, 1982; Kartik, 2009), an informed sender influences an uninformed receiver through direct messages. In signaling and reputation models (Spence, 1973; Daley and Green, 2012; Board and Meyer-ter Vehn, 2013; Dilme, 2019; Gryglewicz and Kolb, 2022; Cetemen and Margaria, 2024; Ekmekci and Maestri, 2022; Ekmekci, Gorno, Maestri, Sun, and Wei, 2022; Sun, 2024), the informed party takes costly actions to shape beliefs. Our model differs in that the manipulator is initially uninformed and does not transmit private information; instead, he distorts the learning process itself. In this respect, our paper is also related to career-concerns models (Holmström, 1999; Bonatti and Hörner, 2017a) and to signal-jamming models in industrial organization (Fudenberg and Tirole, 1986). Unlike those literatures, however, the manipulator in our model cares only about the decision-maker's final action, not about beliefs per se, and equilibrium involves mixed manipulation with persistent uncertainty rather than pure signaling or full revelation.[4]

Finally, our analysis is connected to strategic classification (Brückner and Scheffer, 2011; Hardt, Megiddo, Papadimitriou, and Wootters, 2016) and adversarial learning in computer science (Diakonikolas, Kamath, Kane, Li, Moitra, and Stewart, 2019; Lai, Rao, and Vempala, 2016; Charikar, Steinhardt, and Valiant, 2017). Like these papers, we study decision-making under strategic interference. The focus, however, is different. Strategic-classification models study how agents manipulate observable covariates to obtain favorable classifications, while adversarial-learning models typically seek algorithms that are robust to worst-case data corruption without explicitly modeling the adversary's incentives. By contrast, we analyze optimal sampling when manipulation is endogenous, strategic, and dynamically adapted to the learning process.

2 The Model

We study a continuous-time strategic interaction between a decision-maker, Alice, and an adversary, Bob. Alice sequentially samples data from a bad news Poisson process with arrival rate $\lambda > 0$ at a flow cost $c > 0$.[5] She aims to match her action $a \in \{-1, 1\}$ with an unknown binary state $\omega \in \{-1, 1\}$. Bob continually chooses a manipulation rate $\beta(t) \in [0, 1]$, which reduces the arrival rate to $(1 - \beta(t))\lambda$, at flow cost $\gamma\beta(t)$, and he may randomize. Bob prefers Alice to take the high action. We analyze Nash equilibria of this game. We assume throughout that $c < \gamma < \lambda/2$. This assumption identifies a parameter range of interest and simplifies our exposition.[6] A formal description of the

[4] Relatedly, Bardhi (2024) study how correlation across project attributes affects selective sampling, and Di Tillio, Ottaviani, and Sørensen (2021) examine how sample selection can be used persuasively.

[5] See, e.g., Keller and Rady (2015) or Bonatti and Hörner (2017b).

[6] See Proposition 7 and Remark 1 in the Online Appendix for further discussion of this assumption.

model follows.

2.1 Strategies and Beliefs

A breakdown reveals the bad state, making it optimal for Alice to exit immediately. A strategy for Alice thus consists of choosing when to stop absent a breakdown. Her strategy space, denoted by Σ, consists of all simple random variables on $\mathbb{R}_+$.[7] A generic element of Σ is denoted by σ and represents the random time at which Alice chooses to stop sampling absent a breakdown.

A pure strategy for Bob is a measurable function $\beta : \mathbb{R}_+ \to [0,1]$, where $\beta(t)$ is the manipulation rate at time t, conditional on no prior exit by Alice. Let $\mathcal{D}$ denote the set of all pure strategies.[8] For any $\beta \in \mathcal{D}$, define the cumulative manipulation effort up to time t as $B(t) = \int_0^t \beta(s)ds$. We follow Aumann (1964) and define a mixed strategy for Bob as a measurable function $\phi\colon [0,1] \to \mathcal{D}$.[9] The interpretation is that Bob draws r uniformly at random from $[0,1]$ and then follows the pure strategy ϕ_r. Let Φ denote the set of all mixed strategies. For each r, we write $B_r(t)$ to denote the cumulative manipulation effort under ϕ_r.

Fix Bob's strategy ϕ and denote by τ the random time of the first breakdown. If $\omega = 1$, no breakdown ever occurs and $\tau = \infty$. If $\omega = -1$, τ is exponentially distributed with a time-varying hazard rate $\lambda(1 - \phi_r(t))$, and its distribution is

$$\mathbb{P}[\tau \leq t \mid \omega = -1] = \mathbb{E}_{r\sim U[0,1]}\left[\int_0^t \lambda(1-\phi_r(s))e^{-\lambda(s-B_r(s))}ds\right].$$

Alice exits and the game stops at time $\sigma \wedge \tau$.

Denote by $\tilde{p}_t = \mathbb{P}[\omega = -1 \mid \tau > t]$ Alice's belief process. Bob knows his pure strategy ϕ_r, and so his belief process is $p_t = \mathbb{P}[\omega = -1 \mid \tau > t, \phi_r]$. Thus, the players' beliefs will, in general, diverge unless Bob's strategy is pure.

[7] A random variable is simple if it takes only finitely many distinct values. We assume that Alice's strategy is simple to reduce technicalities.

[8] Since any two manipulation strategies that agree outside a Lebesgue null set yield the same expected utilities, we identify any two such strategies.

[9] Here, $\mathcal{D}$ is equipped with the Borel σ-algebra generated by the topology of pointwise convergence.

2.2 Payoffs

Let $U_A(\sigma, \phi)$ and $U_B(\sigma, \phi)$ denote the expected payoffs of Alice and Bob, respectively, under the strategy profile (σ, ϕ). Alice's payoff is

$$\begin{aligned} U_A(\sigma, \phi) = (1 - p_0)\, \mathbb{E}\left[(1_{\tilde{p}_\sigma \leq \frac{1}{2}} - c\sigma)\right] \\ + p_0\, \mathbb{E}\left[\int_0^\sigma \lambda(1 - \phi_r(t))e^{-\lambda(t - B_r(t))}(1 - ct)dt + e^{-\lambda(\sigma - B_r(\sigma))}\left(1_{\tilde{p}_\sigma > \frac{1}{2}} - c\sigma\right)\right]. \end{aligned}$$

The first term is the contribution from the good state ($\omega = 1$), which has probability $1 - p_0$. In this case, no breakdown occurs, and Alice stops at σ, incurring a sampling cost $c\sigma$ and taking action $a = 1$ if $\tilde{p}_\sigma \leq \frac{1}{2}$.[10] The second term is the contribution of the bad state ($\omega = -1$), which has probability p_0. Its first summand captures the event that a breakdown occurs before σ, in which case Alice learns that the state is bad, chooses $a = -1$, and receives payoff $1 - ct$ if the breakdown occurs at time t. Recall that the breakdown time τ is exponentially distributed with time-varying hazard rate $\lambda(1 - \phi_r(t))$, and so has density $\lambda(1 - \phi_r(t))e^{-\lambda(t - B_r(t))}$. Its second summand captures the event that no breakdown occurs before σ, which has probability $e^{-\lambda(\sigma - B_r(\sigma))}$. In this case, Alice stops at σ, incurs cost $c\sigma$, and takes action $a = 1$ if $\tilde{p}_\sigma \leq \frac{1}{2}$.

Similarly, Bob's expected payoff is

$$\begin{aligned} U_B(\sigma, \phi) = (1 - p_0)\, \mathbb{E}\left[1_{\tilde{p}_\sigma \leq \frac{1}{2}} - \gamma B_r(\sigma)\right] \\ + p_0\, \mathbb{E}\left[\int_0^\sigma \lambda(1 - \phi_r(t))e^{-\lambda(t - B_r(t))}(-\gamma B_r(t))dt + e^{-\lambda(\sigma - B_r(\sigma))}\left(1_{\tilde{p}_\sigma \leq \frac{1}{2}} - \gamma B_r(\sigma)\right)\right]. \end{aligned}$$

The structure mirrors Alice's payoff, with two key differences: (i) Bob receives a positive payoff only when $a = 1$ (i.e., $\tilde{p}_\sigma \leq \frac{1}{2}$), and (ii) his cost is proportional to the cumulative manipulation effort $B_r(t)$.

Observe that it is never optimal for Alice to exit at $t > 0$ absent a breakdown if $\tilde{p}_t \geq \frac{1}{2}$. If she did, it would be optimal to choose $a = -1$, which is also optimal at time 0 or after a breakdown. Thus, stopping with such a belief incurs unnecessary sampling costs without improving the accuracy.

Lemma 1 (No early stopping)**.** *If $\sigma \in \Sigma$ is a best response to $\phi \in \Phi$, then* $\mathbb{P}[\sigma > 0 \wedge \tilde{p}_\sigma \geq \frac{1}{2}] = 0$.

The proof follows directly from inspecting U_A and is omitted. Henceforth, we assume that $a = 1$ if Alice exits at $t > 0$ absent a breakdown.

[10] We assume Alice breaks ties in Bob's favor, choosing $a = 1$ when $\tilde{p}_\sigma = \frac{1}{2}$.

3 Equilibrium Analysis

We consider Bayes-Nash equilibria of this game. First, we show that in every equilibrium, Alice adopts a deterministic stopping rule, and Bob may initially stay passive before randomizing between two extremal strategies: no manipulation at all or full-intensity manipulation until Alice stops. Second, we show that equilibria in which Bob manipulates with positive probability exist for an interval of prior beliefs, and that all such equilibria feature the same on-path behavior. Third, we consider equilibria without manipulation and show that they can coexist with manipulation equilibria. These results characterize all equilibria and are used to analyze the welfare effects of manipulation and compare alternative sampling methods (cf. Section 4 and Section 5).

3.1 Equilibrium Structure

We show that every equilibrium has a simple structure: Alice stops at a deterministic time, and Bob randomizes between no manipulation and manipulation at full intensity from a fixed time onward until Alice stops.

First, we determine which pure manipulation strategies can be best responses to any sampling strategy σ. Suppose that σ is supported on a finite set of times $t_1 < \cdots < t_k$, and set $t_0 = 0$. Then any pure best response β must bunch manipulation effort at the end of each interval $[t_{i-1}, t_i]$, so that there is no manipulation on an initial segment of the interval and full-intensity manipulation on the remainder. The intuition is as follows. If Alice does not exit before t_{i-1}, the probability of a breakdown within $[t_{i-1}, t_i]$ depends only on the total manipulation effort throughout that interval. Front-loading manipulation can waste some of the effort because of a breakdown before t_i. By deferring manipulation while holding the total effort fixed, Bob reduces his expected cost without changing the probability of a breakdown. Thus, any best response has a "delay and bunch" structure.

Lemma 2 (Delay and bunch)**.** *Assume that $\beta \in \mathcal{D}$ is a best response to $\sigma \in \Sigma$ with supp$(\sigma) = \{t_1, \ldots, t_k\}$ and $t_1 < \cdots < t_k$. Then, there exist $s_1, \ldots, s_k \in \mathbb{R}_+$ such that $s_i \in [t_{i-1}, t_i]$ for all $i \in [k]$ and for $t \in [0, t_k]$,*

$$\beta(t) = \begin{cases} 1 & \text{if } t \in [s_i, t_i) \text{ for some } i \in [k], \\ 0 & \text{if } t \in [t_{i-1}, s_i) \text{ for some } i \in [k]. \end{cases}$$

Second, we observe that Bob never manipulates with certainty in any equilibrium. Otherwise, Bob manipulates with certainty at full intensity in an interval of the form $[s_i, t_i)$ by the previous lemma. But then Alice does not learn in that interval, and would strictly prefer to exit at s_i instead of t_i since this saves on sampling cost at no loss in accuracy.

Lemma 3 (No blocking)**.** *Let* (σ, ϕ) *be an equilibrium such that* $\text{supp}(\sigma) = \{t_1, \ldots, t_k\}$ *with* $t_1 < \cdots < t_k$*. Then, for almost every* $t \in [0, t_k]$*,* $\mathbb{P}[\phi_r(t) = 1] < 1$*.*

Bob never suppresses a breakdown with certainty on the equilibrium path. Hence, the absence of a breakdown is evidence of the high state, and Alice's belief decreases strictly.

Third, we establish that Alice's strategy is deterministic in any equilibrium. For intuition, suppose that Alice were to randomize between stopping at t_1 and t_2. Since Alice is indifferent between stopping and sampling at t_1, the marginal value of sampling equals its marginal cost. The likelihood that Bob manipulates is weakly increasing on $[t_1, t_2]$ (Lemma 2) and Alice's belief is strictly decreasing (Lemma 3). Hence, the marginal value of sampling is strictly decreasing on $[t_1, t_2]$. Since the marginal sampling cost is constant, Alice thus strictly prefers stopping at any time between t_1 and t_2, violating her equilibrium constraint.

Lemma 4 (Deterministic stopping)**.** *Let* (σ, ϕ) *be an equilibrium. Then,* σ *is deterministic.*

In summary, Alice stops deterministically at some time t_1 and Bob randomizes over delay-and-bunch strategies in any equilibrium. Any such delay and bunch strategy consists of starting to manipulate at full intensity at some time $s_1 \in [0, t_1]$ and continuing until Alice exits. We show that for any such s_1, Bob strictly prefers not to manipulate at any time after s_1 if he started manipulating at s_1. Hence, he randomizes between two pure strategies: manipulating on some interval $[s_1, t_1]$ or not manipulating at all.

Theorem 1 (Equilibrium structure)**.** *Let* (σ, ϕ) *be an equilibrium. Then,* $\sigma = t_1$ *for some* $t_1 \geq 0$*. If* $t_1 > 0$*, there exist* $s_1 \in [0, t_1]$ *and* $\alpha \in [0, 1)$ *such that with probability* α*,* $\phi_r(t) = 1$ *for all* $t \in [s_1, t_1]$ *and* $\phi_r(t) = 0$ *for all* $t \in [0, s_1)$*, and with probability* $1 - \alpha$*,* $\phi_r(t) = 0$ *for all* $t \in [0, t_1]$*.*

Hence, in any equilibrium, either Alice exits immediately ($t_1 = 0$), or Alice samples until some $t_1 > 0$ absent a breakdown and Bob manipulates at full intensity from some $s_1 \geq 0$ to t_1 with some probability $\alpha \in [0, 1)$ and does not manipulate at all with the remaining probability.[11] Thus, the on-path play of any equilibrium is fully described by

[11] Note that this strategy is not equivalent to Bob manipulating at constant rate α from s_1 to t_1 with probability 1. In the former case, Bob's belief process diverges from Alice's once manipulation begins, whereas if Bob manipulates at the constant rate α with probability 1, both players' beliefs remain the same.

a triplet (s_1, α, t_1), where $0 \leq s_1 \leq t_1$ and $\alpha \in [0, 1)$, and four types of equilibria can occur.

Definition 1. Let (σ, ϕ) be an equilibrium with on-path play (s_1, α, t_1). Denote by t_1^w Alice's optimal exit time if Bob does not manipulate. Then, (σ, ϕ) is

(i) a *manipulation equilibrium* if $0 \leq s_1 < t_1$ and $\alpha \in (0, 1)$,

(ii) a *Wald equilibrium* if $t_1 = t_1^w$ and $\alpha = 0$,

(iii) an *instantaneous-deterrence equilibrium* if $0 = t_1 < t_1^w$, and

(iv) a *deadline-deterrence equilibrium* if $0 < s_1 = t_1 < t_1^w$ and $\alpha = 0$.

Instantaneous deterrence occurs if Bob threatens to manipulate at a high intensity from time 0 onward. This makes sampling uninformative and it is optimal for Alice to exit at time 0. In a deadline-deterrence equilibrium, Bob does not manipulate on path, but threatens to initiate manipulation at t_1 if Alice were to continue past t_1; hence Alice exits strictly before the Wald time t_1^w. We analyze manipulation equilibria and Wald equilibria in detail, and discuss deadline-deterrence equilibria as an intermediate case between these two classes.

3.2 Manipulation Equilibria

Manipulation equilibria are a robust occurrence: they exist for an intermediate range of prior beliefs. Moreover, for any prior, all manipulation equilibria feature the same on-path behavior. In all such equilibria, Alice exits at the same time and Bob starts manipulating at the same time. Only the probability of manipulation can vary across equilibria. In deadline-deterrence equilibria, the sampling period is longer than in manipulation equilibria. Hence, deadline deterrence suppresses sampling less than active manipulation and leads to better outcomes for Alice.

To show that manipulation equilibria exist, we characterize the feasible on-path behavior. Consider the on-path play (s_1, α, t_1) of a manipulation equilibrium. To account for Bob's incentives, it suffices to consider deviations to delay-and-bunch strategies since those include a best reply (Lemma 2). The payoff difference between never manipulating and manipulating from some time $t \leq t_1$ onward is the reduction in the probability of a breakdown on $[t, t_1]$ minus the manipulation cost over that interval. Recalling that p_t is Bob's belief, this difference is

$$N(t) = p_t \left(1 - e^{-\lambda(t_1 - t)}\right) - \gamma (t_1 - t). \tag{1}$$

In equilibrium, Bob must be indifferent at s_1, and he must weakly prefer not to start manipulation before s_1 and between s_1 and t_1 in the case where he does not initiate manipulation at s_1. Hence,

$$\begin{cases} N(t) \leq 0 & \text{for all } t \in [0, s_1), \\ N(t) = 0 & \text{for } t = s_1, \\ N(t) \leq 0 & \text{for all } t \in (s_1, t_1]. \end{cases} \tag{2}$$

For Alice, two conditions are clearly necessary: she must weakly prefer sampling at 0 to stopping immediately, and close to t_1, she must weakly prefer sampling for another instant to stopping immediately. The second condition rules out overly aggressive manipulation and implies that

$$\alpha \leq \overline{\alpha} \equiv 1 - \frac{1}{p_{s_1}} \cdot \frac{c}{(\lambda - c)e^{-\lambda(t_1 - s_1)} + c}. \tag{3}$$

Intuitively, $\overline{\alpha}$ is the largest manipulation probability that prevents Alice from exiting before t_1. If manipulation were more frequent, stopping before t_1 would be optimal.[12]

We show that these conditions are not only necessary, but also sufficient for (s_1, α, t_1) to occur on-path of a manipulation equilibrium (cf. Lemma 13 in the Appendix). This characterization of equilibrium behavior by local incentive conditions is useful to prove that manipulation equilibria exist for an interval of intermediate prior beliefs.[13]

Theorem 2 (Existence of manipulation equilibria)**.** *Let $\mathcal{P}$ be the set of priors p_0 for which a manipulation equilibrium exists, let $\underline{p} \equiv \gamma/\lambda$, and let $\overline{p} \equiv \sup \mathcal{P}$. Then $\mathcal{P}$ is nonempty, and no manipulation equilibrium exists for $p_0 \leq \underline{p}$. Moreover, $\mathcal{P}$ is either $(\underline{p}, \overline{p})$ or $(\underline{p}, \overline{p}]$.*

The equilibrium path can take one of two forms: it can begin with a passive period ($s_1 > 0$) or immediate randomization ($s_1 = 0$). When p_0 is above a cutoff p^* (which is defined in Lemma 10 in the Appendix), the equilibrium starts with a period of no manipulation. For a high prior belief, the sampling period is long absent a breakdown, so sustaining manipulation until Alice exits is expensive. Both players then prefer to

[12]If Bob randomizes at s_1 with probability $\overline{\alpha}$ and sticks to his realized actions if sampling continues off-path after t_1, then Alice is indifferent at t_1. The same on-path times (s_1, t_1) can be sustained with any probability $\alpha \in [0, \overline{\alpha})$ if Bob manipulates with sufficiently high probability after t_1. Then Alice's continuation value has a downward kink at t_1, and is thus not differentiable. If Bob switches from manipulating to not manipulating at t_1, sampling at t_1 is more profitable than sampling just before t_1, and so Alice strictly prefers sampling at t_1, which violates the equilibrium constraint.

[13]Because manipulation equilibria require Bob to mix with strictly positive probability, the upper endpoint of the existence region need not be attained. If the maximum manipulation probability of equilibria for priors approaching $\overline{p}$ converges to zero, $\overline{p} = \sup \mathcal{P}$ but $\overline{p} \notin \mathcal{P}$.

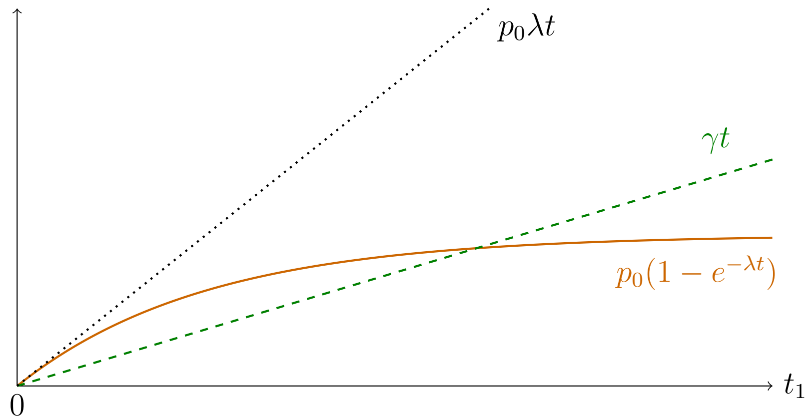

Figure 2: Bob's indifference condition at $t = 0$ if $s_1 = 0$. Alice's stopping time t_1 is determined by the unique strictly positive intersection point of $p_0(1 - e^{-\lambda t})$ and γt.

first let learning proceed. As p_0 falls, the sampling period shortens and so the passive phase also shrinks; at the cutoff p^* Bob begins randomizing immediately.

The lower bound $\underline{p}$ is dictated by Bob's incentives. When sampling begins, Bob must be indifferent between manipulating and staying passive. If the bad state is unlikely, the probability of a breakdown is low, and the cost of manipulation exceeds its gain. To see that no manipulation equilibrium exists for $p_0 \leq \underline{p} = \gamma/\lambda$, consider Bob's indifference condition when there is no passive period: $N(0) = 0$ or equivalently,

$$p_0 \left(1 - e^{-\lambda t_1}\right) = \gamma t_1.$$

The benefit from manipulation on the left is increasing and concave in t_1 with slope $p_0\lambda$ at $t_1 = 0$; the cost of manipulation on the right is linear with slope γ. As illustrated in Figure 2, there is a unique positive solution $t_1 > 0$ if and only if the initial slope of the concave curve exceeds γ, i.e., if $p_0 > \gamma/\lambda$. When $p_0 \leq \gamma/\lambda$, the only solution is $t_1 = 0$, and no manipulation equilibrium exists.

In the proof of Theorem 2, we explicitly construct a manipulation equilibrium. We first show that for priors just above the lower bound $\underline{p}$, there is an equilibrium with on-path play $(0, \alpha, t_1)$: Bob's indifference condition uniquely pins down t_1, and given t_1 we can choose $\alpha > 0$ so that Alice is willing to sample from $t = 0$ to t_1. This guarantees that the equilibrium-existence set $\mathcal{P}$ is nonempty and hence that $\overline{p}$ is well-defined.

It remains to show that $\mathcal{P}$ is an interval. For any prior p_0, we choose (s_1, t_1) so that Bob's incentive condition is satisfied. Since Alice's net gain from starting to sample at $t = 0$ is continuous and strictly decreasing in the manipulation probability α, it suffices to verify that this gain is strictly positive when $\alpha = 0$. This is immediate when

$p_0 \leq 1/2$. When $p_0 > 1/2$, it follows from the fact that Alice's gain under $\alpha = 0$ is strictly decreasing in p_0.

While the mixing probability can vary across manipulation equilibria, the sampling period and manipulation period are uniquely determined. For any fixed prior, Bob starts manipulating and Alice stops sampling at the same time in all manipulation equilibria. Hence, s_1 and t_1 depend only on the prior, and are pinned down by Bob's incentives (2).

Theorem 3 (Uniqueness of manipulation equilibria)**.** *Let $p_0 \in (\underline{p}, \overline{p})$. If $(\hat{\sigma}, \hat{\phi})$ and $(\tilde{\sigma}, \tilde{\phi})$ are two manipulation equilibria with on-path play $(\hat{s}_1, \hat{\alpha}, \hat{t}_1)$ and $(\tilde{s}_1, \tilde{\alpha}, \tilde{t}_1)$, respectively, then $\hat{s}_1 = \tilde{s}_1$ and $\hat{t}_1 = \tilde{t}_1$.*

First, for equilibria without a passive period ($s_1 = 0$), Alice's exit time t_1 is uniquely determined by Bob's indifference at $s_1 = 0$ (see Figure 2). Second, for equilibria that feature a passive period ($s_1 > 0$), Bob's incentive conditions $N(s_1) = N'(s_1) = 0$ (cf. Lemma 13) determine the length of the manipulation window $t_1 - s_1$. The condition $N(s_1) = 0$, together with (1), rules out that the randomization time s_1 differs across equilibria.

In contrast to manipulation equilibria, Bob does not manipulate on-path in a deadline-deterrence equilibrium ($\alpha = 0$), but Alice's exit time $\tilde{t}_1$ is strictly smaller than the Wald time t_1^w because Bob threatens manipulation off-path. Although Bob can, in principle, force Alice to exit earlier by threatening to manipulate earlier, it is never optimal for him to induce exit before the stopping time t_1 in a manipulation equilibrium. Hence, $\tilde{t}_1 \in [t_1, t_1^w)$ (see Proposition 6 in the Online Appendix).

3.3 Wald Equilibria

If Bob is absent, Alice faces the classical sequential sampling problem of Wald (1945). She stops when the marginal cost of sampling equals its marginal benefit. It is well known that there exists an interval of priors $(\underline{p}^w, \overline{p}^w)$ on which sampling for a positive duration is optimal. For any prior in this interval, Alice stops sampling once her belief hits the lower cutoff $\underline{p}^w$. If $t_1^w(p)$ denotes optimal stopping time absent manipulation for the prior p, then $t_1^w(p) = 0$ for $p \in [0, 1] \setminus (\underline{p}^w, \overline{p}^w)$ and $t_1^w(p)$ is continuous and strictly increasing on $[\underline{p}^w, \overline{p}^w)$.[14] Weighing marginal costs against marginal benefits yields $\underline{p}^w = c/\lambda$, and $\overline{p}^w$ is the higher of the two beliefs at which Alice is indifferent between exiting immediately and sampling until her belief hits $\underline{p}^w$.

We now reintroduce Bob. A Wald equilibrium for a prior p_0 is an equilibrium in which Bob never manipulates and Alice exits at $t_1^w(p_0)$. Intuitively, when manipulation

[14] The optimal stopping time is unique for each prior except $\overline{p}^w$. At $\overline{p}^w$, Alice is indifferent between stopping immediately and sampling until the belief reaches c/λ. For concreteness, we define $t_1^w(\overline{p}^w) = 0$.

is cheap, such an equilibrium may fail to exist for some priors (despite $\gamma > c$); when manipulation is more expensive, it exists for all priors. In the rest of this section, we study when the Wald equilibrium exists and how the Wald equilibrium relates to manipulation equilibria.

We first show that, in any manipulation equilibrium, Alice exits before the Wald stopping time. Alice's exit belief under manipulation is strictly higher than her exit belief (c/λ) in the Wald problem without manipulation,[15] which pushes towards earlier stopping; on the other hand, manipulation slows down the belief drift, which delays stopping. While these effects partially cancel each other out, the former effect dominates and Alice stops earlier in manipulation equilibria.

Lemma 5 (Early stopping under manipulation)**.** *Fix any prior $p_0 \in (\underline{p}, \overline{p})$, and let $t_1(p_0)$ be Alice's (unique) exit time in a manipulation equilibrium. Then, $t_1(p_0) < t_1^w(p_0)$.*

Surprisingly, the Wald equilibrium exists whenever a manipulation equilibrium exists, irrespective of the manipulation cost. In a manipulation equilibrium, Bob weakly prefers not manipulating to any deviation. If Alice exits later than in the manipulation equilibrium—specifically at the Wald stopping time—any manipulation would have to be sustained for longer to prevent a breakdown, making manipulation less attractive. Hence, it is optimal not to manipulate if Alice exits at the Wald stopping time. This shows that Wald equilibria exist for priors between $\underline{p}$ and $\overline{p}$. Below $\underline{p}$, the benefit of manipulation is small and no manipulation is optimal irrespective of Alice's stopping time. Wald equilibria with a positive sampling period thus also exist on $(\underline{p}^w, \underline{p}]$.

Theorem 4 ((Co)existence of a Wald equilibrium)**.** *The Wald equilibrium exists for any prior in $(\underline{p}^w, \overline{p}]$.*

Whether the Wald equilibrium exists for priors above $\overline{p}$ depends on the manipulation cost. Existence is guaranteed when manipulation is sufficiently expensive, but can fail otherwise. We provide details in the Online Appendix. Figure 3 illustrates the coexistence of the equilibrium types.

4 Welfare Effects of Manipulation

Manipulation equilibria and Wald equilibria generically coexist. How does the agents' welfare compare across both types of equilibria? Wald equilibria are best-possible for Alice since no information is suppressed through manipulation and she exits at the optimal stopping time. The comparison is less obvious for Bob. On the one hand,

[15] Alice's exit belief $\tilde{p}_{t_1}$ can be derived from (15) by setting $q_{t_1} = c/\lambda$.

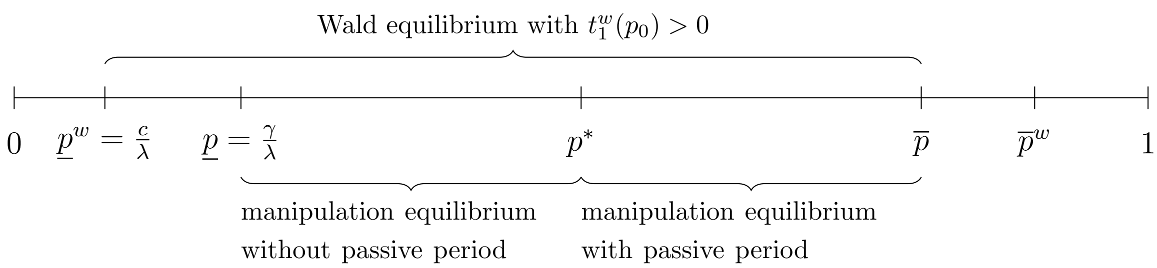

Figure 3: Illustration of the equilibrium types characterized in Theorem 2 and Theorem 4. For $p_0 \in (\underline{p}, \overline{p})$, a manipulation equilibrium exists; it features a passive period when $p_0 > p^*$ and none when $p_0 \leq p^*$. The Wald equilibrium exists for $p_0 \in (\underline{p}^w, \overline{p}]$. For $p_0 \in (\overline{p}, \overline{p}^w]$, the existence of the Wald equilibrium depends on the manipulation cost γ.

manipulation prevents breakdowns and thus prevents Alice from taking the low action. The downside of manipulation is that it is costly. We show that the first effect dominates, so that Bob always prefers manipulation equilibria.

Fix a prior p_0, and recall that t_1^w denotes the Wald stopping time and t_1 is the time at which Alice exits in a manipulation equilibrium absent a breakdown. When there is no manipulation, Alice is worse off stopping at t_1 instead of t_1^w since t_1^w is optimal. Moreover, if Alice exits at t_1, she is worse off in the manipulation equilibrium than without manipulation. Indeed, Alice could recover the manipulation outcome by occasionally ignoring breakdown signals in the latter situation. This shows that Alice always prefers the Wald equilibrium to any manipulation equilibrium. We omit a formal proof of this straightforward fact.

To derive the welfare comparison for Bob, observe that he is indifferent between the two pure strategies in a manipulation equilibrium: never manipulate, or manipulate at full intensity on $[s_1, t_1]$. Hence, his expected payoff does not depend on the manipulation probability and can be computed from the non-manipulation branch: it is the probability that no breakdown occurs until t_1 absent manipulation:

$$p_0 e^{-\lambda t_1} + (1 - p_0).$$

By comparison, his payoff in the Wald equilibrium is

$$p_0 e^{-\lambda t_1^w} + (1 - p_0).$$

Since $t_1 < t_1^w$ (Lemma 5), Bob strictly prefers manipulation equilibria.

This shows that the players have opposing preferences over the two types of equilibria.

Theorem 5 (Welfare effects of manipulation)**.** *For any $p_0 \in (\underline{p}, \overline{p})$, Alice strictly prefers* *the Wald equilibrium to any manipulation equilibrium, while Bob strictly prefers any manipulation equilibrium to the Wald equilibrium.*

For $p_0 \in (\underline{p}, \overline{p})$, both a Wald equilibrium and a continuum of manipulation equilibria exist. Theorem 5 shows Alice and Bob rank them in opposite orders. Selection is therefore a coordination problem: beliefs that coordinate on the Wald profile (exit at t_1^w and no manipulation) sustain it, and beliefs that coordinate on any manipulation outcome (exit at t_1 and manipulate from s_1) sustain that instead. If Alice can publicly commit to stopping at t_1^w, Bob cannot gain by manipulating, so that the Wald outcome is implemented. Absent commitment or additional equilibrium restrictions, standard refinements do not resolve the multiplicity.

5 Comparing Sampling Methods

With *sequential* sampling, Alice can stop at any time. An alternative is *static* sampling: commit at time 0 to a sample size and act only after observing the full sample. In non-adversarial environments, sequential sampling dominates because a breakdown reveals the state and further sampling does not improve the decision (Wald, 1945). We show that, however, Bob—the manipulating adversary—can flip this comparison: the equilibrium in the static sampling game with manipulation can be better for Alice than the worst manipulation equilibrium under sequential sampling. We then introduce a hybrid design that combines both sampling methods and can strictly dominate each pure method.

5.1 Static Sampling vs Sequential Sampling

Consider the following static sampling game with manipulation: at time 0, Alice commits to a sample size t_a. After observing t_a, Bob selects a subsample size $t_b \in [0, t_a]$, which prevents breakdowns in $[0, t_b]$. Alice then observes the full sample $[0, t_a]$ and chooses an action.

Manipulation reduces the probability of breakdown, but is costly. Bob's gain from choosing t_b instead of not manipulating is

$$G(t_b, \gamma) \equiv p_0 \left(e^{-\lambda(t_a - t_b)} - e^{-\lambda t_a}\right) - \gamma t_b.$$

Because $G(\cdot, \gamma)$ is convex, Bob's best response is extremal: $t_b \in \{0, t_a\}$.

Lemma 6 (Static sampling equilibrium)**.** *Let $p_0 \in (c/\lambda, p^*)$ with $p^* \leq \overline{p}$. With static sampling, Bob never manipulates in equilibrium, and Alice chooses*

$$t_a = \begin{cases} t_1 & \text{if} \quad \gamma < \widehat{\gamma} \\ t^s & \text{if} \quad \gamma \geq \widehat{\gamma} \end{cases} \tag{4}$$

where t_1 is Alice's stopping time in the manipulation equilibrium with sequential sampling, $t^s = (\log(\lambda p_0) - \log c)/\lambda$ is the optimal sample size for static sampling without manipulation, and $\widehat{\gamma}$ solves $G(t^s, \widehat{\gamma}) = 0$.

In equilibrium, Alice chooses t_a just large enough to deter manipulation. If manipulation is costly ($\gamma \geq \widehat{\gamma}$), she chooses the benchmark t^s. If manipulation is cheap ($\gamma < \widehat{\gamma}$), she commits to $t_a = t_1$, making Bob exactly indifferent ($G(t_1, \gamma) = 0$).

Absent manipulation, sequential sampling dominates static sampling. With manipulation, however, committing to a fixed sample can deter manipulation (Lemma 6). We show that Alice may strictly prefer static sampling to the *worst* manipulation equilibrium under sequential sampling. Hence, the fact that Alice may face manipulation can flip the comparison of sampling methods in equilibrium.

Proposition 1 (Static vs sequential sampling)**.** *Let $p_0 \in (\underline{p}, 1/2)$. If $\lambda p_0 e^{-\lambda t_1} \geq c$, Alice's expected payoff in the static sampling equilibrium is strictly higher than in the worst manipulation equilibrium under sequential sampling, i.e., when Bob manipulates with probability $\overline{\alpha}$.*

Sequential sampling would still dominate static sampling if the Wald equilibrium (or a manipulation equilibrium with a small α) were expected in the sequential game. Our point is that, with manipulation, the flexibility of sequential sampling also creates a vulnerability that an adversary can exploit; sequential sampling no longer robustly dominates static sampling.[16]

5.2 A Hybrid Approach

Sequential sampling is efficient absent manipulation, whereas static sampling is more robust because it limits Bob's ability to condition on the realized path. This suggests a hybrid protocol: Alice samples sequentially on $[0, s]$ and, conditional on no breakdown by time s, commits to a fixed additional sample on $[s, t]$. We denote this policy by $[0, s] \oplus [s, t]$.

[16] Example 1 in the Online Appendix illustrates how sequential sampling can dominate static sampling for higher priors while the reverse can be true for lower priors. It also shows that static sampling may still dominate even if Alice and Bob move simultaneously, so the reversal of the classical ranking is not driven solely by the assumption that Alice moves first in static sampling.

Proposition 2 (Hybrid sampling)**.** *Let $[s_1, t_1]$ be the manipulation period in a manipulation equilibrium for some prior p_0, and let t^s be the optimal sample size under static sampling without manipulation.*

(i) If $p_0 \geq p^*$ and $\lambda p_0 e^{-\lambda t_1} > c$, Alice strictly prefers $[0,\, t^s - t_1] \oplus [t^s - t_1,\, t^s]$ to static sampling with sample size t^s in equilibrium.

(ii) Suppose a manipulation equilibrium with on-path play $(0, \alpha, t_1^*)$ exists under sequential sampling, and Alice strictly prefers static sampling to this equilibrium.[17] *Then, if $p_0 \geq p^*$, Alice strictly prefers the equilibrium under $[0,\, s_1] \oplus [s_1,\, t_1]$ to the manipulation equilibrium with on-path play (s_1, α, t_1).*

The hybrid protocol uses sequential experimentation only where flexibility is most valuable, and switches to commitment exactly where manipulation would otherwise become attractive. In part (i), Alice keeps the same total no-breakdown sample length as under pure static sampling, but moves the first $t^s - t_1$ units into an exploratory sequential phase. Bob does not manipulate under this hybrid policy, so Alice obtains the same continuation payoff as under static sampling on the no-breakdown path and strictly gains from the possibility of an earlier breakdown during the exploratory phase.

Part (ii) uses the initial passive segment of the manipulation equilibrium. Under $[0, s_1] \oplus [s_1, t_1]$, the history up to s_1 is unchanged, but the commitment at s_1 removes Bob's incentive to manipulate thereafter. After no breakdown by s_1, Alice's posterior equals p^*, so the continuation under the hybrid coincides with the static continuation at p^*; by assumption, this continuation is strictly better for Alice than the manipulation-equilibrium continuation. Thus the hybrid can improve not only on pure static sampling, but also on the manipulation equilibrium under pure sequential sampling.

6 Discussion

Modeling the interaction between a decision-maker and a manipulator presents several challenges. The decision problem must be general enough to capture a range of applications, yet structured enough to reflect the manipulator's capabilities and constraints. At the same time, the model must remain tractable to yield meaningful insights. To address these challenges, we extend the classical Wald sequential sampling framework by incorporating strategic data manipulation. We adopt a simple bad-news learning structure, where Bob's manipulation reduces the arrival rate of breakdowns in the low

[17] We do not have simple sufficient conditions for this ranking. It is easy to find parameter values to satisfy the assumed properties. See Example 3 in the Online Appendix.

state, thereby slowing Alice's learning. This section explores how different modeling choices affect the broader implications of our results.

6.1 Informed Adversary

We assumed that Alice and Bob are symmetrically uninformed ex ante. In some applications, however, it is more natural to assume that Bob knows the true state from the outset. Then the unique equilibrium is trivial: Alice exits immediately.[18]

Proposition 3 (Informed adversary)**.** *Assume that Bob knows the state. Then, there is a unique equilibrium in which Alice exits at time* 0 *and Bob always manipulates in the low state off path.*

This result underscores the importance of symmetric uncertainty in sustaining nontrivial learning dynamics. The intuition is as follows. In the low state, Bob knows that a breakdown arrives at rate λ. Since his manipulation cost is less than λ, manipulation is always optimal. On the other hand, in the high state, Bob knows that no breakdown will occur, so that manipulation is always strictly suboptimal. Thus, in equilibrium, no breakdown ever occurs, Alice gains no information by sampling, and it is optimal to exit at time 0 and decide based on her prior belief.

6.2 Non-adapted manipulation

Bob can adapt manipulation to the realized sampling history, which may not be possible in applications. The manipulator may not observe outcomes in real time, or the available technology may require committing to an intervention plan in advance. To capture this friction, we consider a variant with *non-adapted* manipulation.

At time 0, Bob commits to a (possibly mixed) manipulation strategy, and pays the full manipulation cost up front, independently of Alice's stopping time. That is, if Bob's pure strategy is $\beta\colon \mathbb{R}_+ \to [0,1]$, his manipulation cost is $\gamma \int_0^\infty \beta(t)\, dt$. This makes manipulation more costly compared to the case where Bob only pays for manipulation until Alice exits. Under non-adapted manipulation, only the total amount of manipulation prior to Alice's stopping matters; its timing does not.

The following proposition shows that when Bob has to pay manipulation cost up front, there is a unique equilibrium in which Alice exits at the Wald stopping time, and Bob never manipulates. Therefore, adaptivity is crucial for the existence of a manipulation equilibrium in the baseline model.

[18]The proof of Proposition 3 does not need the assumption that Alice's stopping time σ is purely atomic.

Proposition 4 (Non-adapted manipulation)**.** *Fix $p_0 \in (\underline{p}^w, \overline{p}^w)$ and let $t_1^w(p_0)$ denote the no-manipulation Wald stopping time from prior p_0. Then the game with non-adapted manipulation has a unique equilibrium outcome: Bob does not manipulate ($\beta \equiv 0$), and along the no-breakdown history Alice stops at $t_1^w(p_0)$.*

We sketch the argument and omit the formal proof. Given non-adapted manipulation, Alice's stopping time is deterministic, and Bob never manipulates after she stops. Manipulation weakly raises Alice's posterior along the no-breakdown history, so optimality implies that she stops no earlier than the no-manipulation Wald time. But once Alice waits at least until the Wald time, manipulation is strictly unprofitable for Bob. Hence Bob does not manipulate, and Alice stops exactly at the Wald time.

6.3 Slightly Convex Manipulation Cost

We assume that Bob's cost for manipulating at rate $b \in [0, 1]$ is γb, and thus linear in the manipulation intensity. This delivers a stark equilibrium structure: Alice stops deterministically, and Bob randomizes once between two polar pure strategies: never manipulate, or manipulate at full intensity from a deterministic start time until Alice stops. The qualitative predictions are the same when the cost is *slightly* convex.

Proposition 5 (Manipulation equilibria under slightly convex costs)**.** *Let (σ^0, ϕ^0) be a manipulation equilibrium with on-path play (s_1^0, α^0, t_1^0). Then, there exists $\bar{\varepsilon} > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon})$, the game with manipulation cost function $k_\varepsilon(b) = \gamma b + \varepsilon b^2$ admits an equilibrium $(\sigma^\varepsilon, \phi^\varepsilon)$ such that along the no-breakdown history, Alice stops at t_1^ε, and Bob mixes between no manipulation (probability $1 - \alpha^\varepsilon$) and a pure strategy β^ε (probability α^ε), where $\beta^\varepsilon(t) = 0$ for $t < s_1^\varepsilon$ and β^ε is weakly increasing on $[s_1^\varepsilon, t_1^\varepsilon]$.*

Proposition 5 shows that the linear-cost manipulation equilibrium persists under a small quadratic perturbation of manipulation costs. In the constructed equilibrium, Alice continues to stop at a deterministic time and Bob mixes between not manipulating and a pure strategy that includes manipulation. The difference is on the manipulation branch: with linear costs, Bob first remains passive and then switches immediately to full manipulation (a bang–bang "bunch"), whereas the quadratic term penalizes concentrating effort, so Bob instead increases the manipulation intensity over time and maintains full manipulation when he reaches the cap. This replaces the jump with a (short) ramp and leaves the comparative statics intact.

6.4 Alternative Learning Technology

Our analysis focuses on *bad-news* learning, in which evidence arrives only in the low state. This framework is natural for environments where the decision-maker screens for adverse events—failures, breaches, defects, or side effects—and an interested party benefits from delaying their detection. It isolates a simple strategic tension: Alice values timely bad-news detection, while Bob benefits from suppressing or postponing bad news.

At the same time, our sharp equilibrium characterization is tied to this information technology. Under *good-news* learning, where breakthroughs arrive only in the high state (e.g., Keller, Rady, and Cripps, 2005; Bonatti and Hörner, 2011), a manipulator would instead want to accelerate favorable signals. The strategic forces, and therefore the equilibrium structure, can be quite different. We leave a systematic analysis of good-news manipulation to future work.

We also assume conclusive learning: a single breakdown fully reveals the low state. This assumption helps deliver a tractable closed-form analysis. With richer signal structures, beliefs would evolve more gradually, and the equilibrium need not retain the same simple form. Our companion paper, Brandl and Shi (2024), studies a continuous-learning environment with Brownian signals in which Bob chooses hidden effort that affects the drift. There too manipulation worsens Alice's payoff and decision quality, but a sharp closed-form equilibrium characterization is much harder to obtain.

APPENDIX

A Omitted Proofs from Section 3.1

Lemma 2 (Delay and bunch)**.** *Assume that $\beta \in \mathcal{D}$ is a best response to $\sigma \in \Sigma$ with $supp(\sigma) = \{t_1, \ldots, t_k\}$ and $t_1 < \cdots < t_k$. Then, there exist $s_1, \ldots, s_k \in \mathbb{R}_+$ such that $s_i \in [t_{i-1}, t_i]$ for all $i \in [k]$ and for $t \in [0, t_k]$,*

$$\beta(t) = \begin{cases} 1 & \text{if } t \in [s_i, t_i) \text{ for some } i \in [k], \\ 0 & \text{if } t \in [t_{i-1}, s_i) \text{ for some } i \in [k]. \end{cases}$$

Proof. Suppose β does not have the stated form. Then there exists an interval $[t_{i-1}, t_i]$ and a point $s_i \in [t_{i-1}, t_i)$ such that we can define an alternative strategy $\tilde{\beta} \neq \beta$ by shifting all manipulation effort to the end of the interval:

$$\tilde{\beta}(t) = \begin{cases} 1 & \text{if } t \in [s_i, t_i) \\ 0 & \text{if } t \in [t_{i-1}, s_i) \\ \beta(t) & \text{otherwise} \end{cases}$$

That is, when following $\tilde{\beta}$, Bob does not manipulate from t_{i-1} to s_i and manipulates with intensity 1 from s_i to t_i, where s_i is chosen such that the overall manipulation effort in $[t_{i-1}, t_i]$ remains unchanged: $\int_{t_{i-1}}^{t_i} \beta(t)dt = \int_{t_{i-1}}^{t_i} \tilde{\beta}(t)dt = t_i - s_i$.

We argue below that $U_B(\sigma, \tilde{\beta}) > U_B(\sigma, \beta)$, and thus β cannot be a best response. Recall that $B(t) = \int_0^t \beta(s)ds$ and define $\tilde{B}(t) = \int_0^t \tilde{\beta}(s)ds$. Clearly, $B(t) \geq \tilde{B}(t)$ for all $t \in \mathbb{R}_+$ with equality for all $t \notin [t_{i-1}, t_i]$ and a strict inequality on a positive measure subset of $[t_{i-1}, t_i]$. Since Alice always chooses $a = 1$ if she exits absent a breakdown (cf. Lemma 1),

$$U_B(\sigma, \beta) - U_B(\sigma, \tilde{\beta}) = p_0 \, \mathbb{P}[\sigma \geq t_i] \begin{pmatrix} \int_{t_{i-1}}^{t_i} \lambda(1-\beta(t))e^{-\lambda(t-B(t))}(-\gamma B(t))dt \\ -\int_{t_{i-1}}^{t_i} \lambda(1-\tilde{\beta}(t))e^{-\lambda(t-\tilde{B}(t))}(-\gamma \tilde{B}(t))dt \end{pmatrix}. \tag{5}$$

By definition of $\tilde{\beta}$,

$$\begin{aligned} \int_{t_{i-1}}^{t_i} \lambda(1-\tilde{\beta}(t))e^{-\lambda(t-\tilde{B}(t))}\tilde{B}(t)dt &= \int_{t_{i-1}}^{s_i} \lambda e^{-\lambda(t-B(t_{i-1}))}B(t_{i-1})dt \\ &= B(t_{i-1})\left(e^{-\lambda(t_{i-1}-B(t_{i-1}))} - e^{-\lambda(s_i - B(t_{i-1}))}\right) \\ &= B(t_{i-1})\left(e^{-\lambda(t_{i-1}-B(t_{i-1}))} - e^{-\lambda(t_i - B(t_i))}\right) \end{aligned}$$

where the last equality follows from the fact that $B(t_i) - B(t_{i-1}) = t_i - s_i$ by the choice of s_i. Meanwhile,

$$\begin{aligned} \int_{t_{i-1}}^{t_i} \lambda(1-\beta(t))e^{-\lambda(t-B(t))}B(t)dt &> \int_{t_{i-1}}^{t_i} \lambda(1-\beta(t))e^{-\lambda(t-B(t))}B(t_{i-1})dt \\ &= B(t_{i-1})\left(e^{-\lambda(t_{i-1}-B(t_{i-1}))} - e^{-\lambda(t_i - B(t_i))}\right) \end{aligned}$$

where the inequality follows since β is not of the form claimed in the lemma. Substituting the two previous (in)equalities into (5) gives $U_B(\sigma, \beta) - U_B(\sigma, \tilde{\beta}) < 0$. □

Lemma 3 (No blocking)**.** *Let* (σ, ϕ) *be an equilibrium such that* $\text{supp}(\sigma) = \{t_1, \ldots, t_k\}$ *with* $t_1 < \cdots < t_k$. *Then, for almost every* $t \in [0, t_k]$, $\mathbb{P}[\phi_r(t) = 1] < 1$.

Proof. We begin with the second part. Suppose, for contradiction, that there exists

$i \in [k]$ and $t \in (t_{i-1}, t_i)$ such that $\mathbb{P}[\phi_r(t) = 1] = 1$. Since ϕ is a best response to σ, ϕ_r is a best response for almost every $r \in [0,1]$. Hence, by Lemma 2, $\phi_r(s) = 1$ for all $s \in [t, t_i)$ and almost every $r \in [0,1]$. This implies that Alice's belief $\tilde{p}_s$ remains constant on $[t, t_i)$, contradicting the optimality of waiting until t_i rather than exiting at t. Hence, $\mathbb{P}[\phi_r(t) = 1] < 1$. For the first part, observe that

$$\tilde{p}_t = \mathbb{P}[\omega = -1 \mid \tau \geq t] = \frac{p_0\, \mathbb{P}[\tau \geq t \mid \omega = -1]}{1 - p_0 + p_0\, \mathbb{P}[\tau \geq t \mid \omega = -1]}$$

Since $\mathbb{P}[\tau \geq t \mid \omega = -1]$ is strictly decreasing in t on $[0, t_k]$ by the second part, it follows that $\tilde{p}_t$ is strictly decreasing as well. □

Lemma 4 (Deterministic stopping)**.** *Let (σ, ϕ) be an equilibrium. Then, σ is deterministic.*

Proof. Assume, for contradiction, that σ is not deterministic, and let $\text{supp}(\sigma) = \{t_1, \ldots, t_k\}$ with $t_1 < \cdots < t_k$ and $k \geq 2$. Since ϕ is a best response to σ, it follows that ϕ_r is a best response for almost every $r \in [0,1]$. Hence, for almost all r, Bob's deterministic strategy is as described in Lemma 2.

First, we show that, within the interval $[t_{k-1}, t_k]$, the probability of a breakdown in the next ϵ units of time is decreasing in t, and the expected time until a breakdown is increasing in t.

Claim 1. Let $\epsilon > 0$. Then $\mathbb{P}[\tau \leq t + \epsilon \mid \tau > t]$ is decreasing in t on $[t_{k-1}, t_k - \epsilon]$, and $\mathbb{E}[\epsilon \wedge (\tau - t) \mid \tau > t]$ is increasing in t on $[t_{k-1}, t_k - \epsilon]$.

Proof of Claim 1. Let $r \sim U[0,1]$ denote the random variable determining Bob's deterministic strategy. Denote by $\tilde{q}_t$ the probability that the state is low and Bob does not manipulate at time t, conditional on no breakdown before t:[19]

$$\tilde{q}_t = \mathbb{P}[\omega = -1, \phi_r(t) = 0 \mid \tau \geq t].$$

By Lemma 2, $\phi_r(t) \in \{0,1\}$ almost surely for all $t \in [t_{k-1}, t_k]$. Hence, the instantaneous arrival rate of breakdowns at $t \in [t_{k-1}, t_k]$ is determined by $\tilde{q}_t$. To prove the first part, it suffices, by definition of τ (cf. Section 2.1), to show that $\tilde{q}_t$ is decreasing in t.

Let $R_t = \{r \in [0,1] \mid \phi_r(t) = 1\}$. By Lemma 2, the set of r for which $\phi_r(t) = 1$ is non-decreasing in t and hence the sequence $(R_t)_{t \in [t_{k-1}, t_k)}$ is non-decreasing in the sense

[19] Note that, although ω and r are independent, they are not independent conditional on $\tau \geq t$. Intuitively, those values of r for which Bob manipulates a lot up to t become more likely relative to values of r for which Bob manipulates less up to t since manipulation suppresses a breakdown.

of set inclusion and $\phi_r(t) = 0$ for $r \notin R_t$. Observe that

$$\tilde{q}_t = \mathbb{P}[r \notin R_t \mid \omega = -1, \tau \geq t]\, \mathbb{P}[\omega = -1 \mid \tau \geq t] = \mathbb{P}[r \notin R_t \mid \omega = -1, \tau \geq t]\tilde{p}_t.$$

Let $s, t \in [t_{k-1}, t_k)$ with $s < t$. Then,

$$\begin{aligned}\mathbb{P}[r \in R_s \mid \omega = -1, \tau \geq s] &= \frac{\mathbb{P}[\tau \geq s, r \in R_s \mid \omega = -1]}{\mathbb{P}[\tau \geq s \mid \omega = -1]} = \frac{\mathbb{P}[\tau \geq t, r \in R_s \mid \omega = -1]}{\mathbb{P}[\tau \geq s \mid \omega = -1]} \\ &\leq \frac{\mathbb{P}[\tau \geq t, r \in R_s \mid \omega = -1]}{\mathbb{P}[\tau \geq t \mid \omega = -1]} = \mathbb{P}[r \in R_s \mid \omega = -1, \tau \geq t]\end{aligned}$$

where the second step uses that no breakdown occurs in $[s, t]$ if Bob manipulates in that interval. Using that $R_s \subset R_t$ gives

$$\mathbb{P}[r \notin R_s \mid \omega = -1, \tau \geq s] \geq \mathbb{P}[r \notin R_s \mid \omega = -1, \tau \geq t] \geq \mathbb{P}[r \notin R_t \mid \omega = -1, \tau \geq t].$$

By Lemma 3, $\tilde{p}_t$ is strictly decreasing on $[0, t_k]$ and $\mathbb{P}[r \notin R_t \mid \omega = -1, \tau \geq t] > 0$. Hence, $\tilde{q}_s > \tilde{q}_t$, proving the first part of the claim.

The second part follows from the first: if the expected arrival rate of breakdowns is decreasing, then the conditional distribution of $\tau - t$ given $\tau > t$ becomes stochastically larger as t increases. Formally, the first part (applied for each $\epsilon' \in (0, \epsilon]$) implies that for $s, t \in [t_{k-1}, t_k - \epsilon]$ with $s < t$, the distribution $\mathbb{P}[\epsilon \wedge (\tau - s) \in \cdot \mid \tau > s]$ is stochastically dominated by $\mathbb{P}[\epsilon \wedge (\tau - t) \in \cdot \mid \tau > t]$. This implies that the expected time to breakdown increases. □

Second, we decompose the conditional probability of a breakdown in the interval $[s, t]$ into two components: the change in Alice's belief conditional on survival, $\tilde{p}_s - \tilde{p}_t$, and the belief-weighted probability of a breakdown, $\mathbb{P}[\tau \leq t \mid \tau \geq s] \cdot \tilde{p}_t$.

Claim 2. For $s, t \in [0, t_k]$ with $s < t$,

$$\mathbb{P}[\tau \leq t \mid \tau \geq s] = \tilde{p}_s - \tilde{p}_t + \mathbb{P}[\tau \leq t \mid \tau \geq s] \cdot \tilde{p}_t.$$

Proof. First, observe that

$$\begin{aligned}\tilde{p}_t &= \frac{\mathbb{P}[\tau \geq t, \omega = -1]}{\mathbb{P}[\tau \geq t]} = \frac{\mathbb{P}[\tau \geq t, \omega = -1, \tau \geq s]}{\mathbb{P}[\tau \geq t]} \\ &= \frac{\mathbb{P}[\tau \geq t \mid \omega = -1, \tau \geq s]\, \mathbb{P}[\omega = -1, \tau \geq s]}{\mathbb{P}[\tau \geq t \mid \tau \geq s]\, \mathbb{P}[\tau \geq s]} = \frac{\mathbb{P}[\tau \geq t \mid \omega = -1, \tau \geq s]}{\mathbb{P}[\tau \geq t \mid \tau \geq s]}\tilde{p}_s\end{aligned} \tag{6}$$

The claim now follows from the following series of equations.

$$\begin{aligned}
&\tilde{p}_s - \tilde{p}_t + \mathbb{P}[\tau \le t \mid \tau \ge s] \cdot \tilde{p}_t = \tilde{p}_s - \tilde{p}_t(1 - \mathbb{P}[\tau \le t \mid \tau \ge s]) \\
&= \tilde{p}_s - \tilde{p}_t\, \mathbb{P}[\tau \ge t \mid \tau \ge s] = \tilde{p}_s\,(1 - \mathbb{P}[\tau \ge t \mid \omega = -1, \tau \ge s]) \\
&= \tilde{p}_s\, \mathbb{P}[\tau \le t \mid \omega = -1, \tau \ge s] = \frac{\mathbb{P}[\omega = -1, \tau \ge s]}{\mathbb{P}[\tau \ge s]}\, \mathbb{P}[\tau \le t \mid \omega = -1, \tau \ge s] \\
&= \frac{\mathbb{P}[\tau \le t, \omega = -1, \tau \ge s]}{\mathbb{P}[\tau \ge s]} = \frac{\mathbb{P}[\tau \le t, \tau \ge s]}{\mathbb{P}[\tau \ge s]} = \mathbb{P}[\tau \le t \mid \tau \ge s]
\end{aligned}$$

where the third step uses (6) and the second to last step uses that $\tau = \infty$ if $\omega = 1$. □

Returning to the proof of Lemma 4, fix $\epsilon \in (0, t_k - t_{k-1})$. Denote by $V(t)$ Alice's continuation value at $t \ge 0$ conditional on no breakdown before t. Since $t_{k-1} \in \operatorname{supp}(\sigma)$, Alice must be indifferent between exiting and continuing at t_{k-1}. Hence, we have

$$\begin{aligned}
1 - \tilde{p}_{t_{k-1}} &= -\, c\,\mathbb{E}[\epsilon \wedge (\tau - t_{k-1}) \mid \tau > t_{k-1}] + \mathbb{P}[\tau \le t_{k-1} + \epsilon \mid \tau > t_{k-1}] \\
&\quad + (1 - \mathbb{P}[\tau \le t_{k-1} + \epsilon \mid \tau > t_{k-1}])V(t_{k-1} + \epsilon) \\
&\ge -\, c\,\mathbb{E}[\epsilon \wedge (\tau - t_{k-1}) \mid \tau > t_{k-1}] + \mathbb{P}[\tau \le t_{k-1} + \epsilon \mid \tau > t_{k-1}] \\
&\quad + (1 - \mathbb{P}[\tau \le t_{k-1} + \epsilon \mid \tau > t_{k-1}])\left(1 - \tilde{p}_{t_{k-1}+\epsilon}\right)
\end{aligned} \tag{7}$$

where the inequality follows since the continuation value is never less than the value of exiting immediately. Similarly, Alice weakly prefers continuing at $t_k - \epsilon$, and so

$$\begin{aligned}
1 - \tilde{p}_{t_k - \epsilon} &\le -\, c\,\mathbb{E}[\epsilon \wedge (\tau - t_k + \epsilon) \mid \tau > t_k - \epsilon] + \mathbb{P}[\tau \le t_k \mid \tau > t_k - \epsilon] \\
&\quad + (1 - \mathbb{P}[\tau \le t_k \mid \tau > t_k - \epsilon])(1 - \tilde{p}_{t_k})
\end{aligned} \tag{8}$$

where we use that $V(t_k) = 1 - \tilde{p}_{t_k}$ since Alice is indifferent between exiting and continuing at t_k. Rewriting (7) and (8) and using Claim 2 yield

$$\begin{aligned}
c\,\mathbb{E}[\epsilon \wedge (\tau - t_{k-1}) \mid \tau > t_{k-1}] &\ge \tilde{p}_{t_{k-1}} - \tilde{p}_{t_{k-1}+\epsilon} + \mathbb{P}[\tau \le t_{k-1} + \epsilon \mid \tau > t_{k-1}]\tilde{p}_{t_{k-1}+\epsilon} \\
&= \mathbb{P}[\tau \le t_{k-1} + \epsilon \mid \tau \ge t_{k-1}]
\end{aligned} \tag{9}$$

and

$$\begin{aligned}
c\,\mathbb{E}[\epsilon \wedge (\tau - t_k + \epsilon) \mid \tau > t_k - \epsilon] &\le \tilde{p}_{t_k - \epsilon} - \tilde{p}_{t_k} + \mathbb{P}[\tau \le t_k \mid \tau > t_k - \epsilon]\tilde{p}_{t_k} \\
&= \mathbb{P}[\tau \le t_k \mid \tau \ge t_k - \epsilon]
\end{aligned} \tag{10}$$

But (9) and (10) contradict Claim 1, which implies

$$\mathbb{P}[\tau \leq t_{k-1} + \epsilon \mid \tau \geq t_{k-1}] > \mathbb{P}[\tau \leq t_k \mid \tau \geq t_k - \epsilon]$$
$$\mathbb{E}[\epsilon \wedge (\tau - t_{k-1}) \mid \tau > t_{k-1}] < \mathbb{E}[\epsilon \wedge (\tau - t_k + \epsilon) \mid \tau > t_k - \epsilon].$$

Therefore, our assumption that σ is not deterministic must be false. □

Theorem 1 (Equilibrium structure)**.** *Let (σ, ϕ) be an equilibrium. Then, $\sigma = t_1$ for some $t_1 \geq 0$. If $t_1 > 0$, there exist $s_1 \in [0, t_1]$ and $\alpha \in [0, 1)$ such that with probability α, $\phi_r(t) = 1$ for all $t \in [s_1, t_1]$ and $\phi_r(t) = 0$ for all $t \in [0, s_1)$, and with probability $1 - \alpha$, $\phi_r(t) = 0$ for all $t \in [0, t_1]$.*

Proof. The first part follows from Lemma 4. For the second part, since ϕ is a best response to σ, ϕ_r is a best response to σ for almost all $r \in [0, 1]$. Let $\psi \colon \mathbb{R}_+ \to [0, 1]$ with $\psi(t) = \int_0^1 \phi_r(t) dr$ be the probability that Bob manipulates at $t \in \mathbb{R}_+$. We first claim that $\psi(t) < 1$ for all $t \in [0, t_1)$. To see this, note that by Lemma 2, ψ is non-decreasing on $[0, t_1)$, and thus if $\psi(t) = 1$ for some $t \in [0, t_1)$, then $\psi(s) = 1$ for all $s \in [t, t_1)$. But then, $\tilde{p}_s$ is constant on $[t, t_1)$, which implies that Alice does not learn past t and prefers to exit at t. This contradicts that $\sigma = t_1$ is a best response to ϕ.

Define Bob's utility difference function:

$$N(t) = p_t(1 - e^{-\lambda(t_1 - t)}) - \gamma(t_1 - t),$$

which captures the gain from manipulating at full intensity on $[t, t_1]$ versus not manipulating at all over that interval. For any $t \in [0, t_1)$, it follows from $\psi(t) < 1$ and Lemma 2 that there exists a best response $\beta \in \mathcal{D}$ such that $\beta(s) = 0$ for all $s \in [0, t]$. Hence, never manipulating is a best response by continuity, and so $N(t) \leq 0$ for all $t \in [0, t_1]$.

Suppose now that $\beta = 1_{[s_1, t_1]}$ is a best response for some $s_1 \in [0, t_1)$. Then $N(s_1) = 0$. If $s_1 > 0$, optimality of s_1 implies $N'(s_1) = 0$. If $s_1 = 0$, optimality at the boundary implies instead that $N'(0) \leq 0$. By Lemma 7, the function $N(\cdot)$ is strictly concave on $[0, \hat{t})$ and strictly convex on $(\hat{t}, t_1]$ for some $\hat{t} \in [0, t_1]$. It follows that $N(t) < 0$ for all $t \in (s_1, t_1)$: if $s_1 > 0$, this follows from $N(s_1) = N'(s_1) = N(t_1) = 0$; if $s_1 = 0$, it follows from $N(0) = N(t_1) = 0$ and $N'(0) \leq 0$.

This implies that if $\psi(t) > 0$ for some $t \in [0, t_1]$, then ψ must be constant on $[t, t_1]$. Therefore, ψ must be a step function: there exist $s_1 \in [0, t_1)$ and $\alpha \in [0, 1)$ such that

$$\psi(t) = \begin{cases} 0 & \text{for } t \in [0, s_1), \\ \alpha & \text{for } t \in [s_1, t_1]. \end{cases}$$

Define the set $R = \{r \in [0,1] \mid \phi_r = 1_{[s_1,t_1]} \text{ on } [0,t_1]\}$. Since $\psi(t) < 1$ for all $t \in [0,t_1)$, it follows that $|R| = \alpha < 1$, completing the proof. □

B Omitted Proofs from Section 3.2

We begin with several auxiliary lemmas. They characterize Bob's gain from manipulation (Lemma 7), his incentive conditions (Lemma 8), Alice's expected payoff from sampling (Lemma 9), monotonicity of Alice's payoff in the prior (Lemma 11), and the critical belief cutoff $p^* > 1/2$ (Lemma 10). Their proofs are relegated to the Online Appendix.

We then derive necessary and sufficient conditions for a triple (s_1, α, t_1) to arise as the on-path play of a manipulation equilibrium (Lemma 13), identify the critical belief p^*, and establish monotonicity of the on-path play (Lemma 14). These results together yield the proofs of Theorems 2 and 3.

B.1 Auxiliary Results

Lemma 7 (Shape of Bob's manipulation gain)**.** *For all $t \in [0,t_1]$, $N(t)$ is strictly convex at t if $p_t < 1/2$, strictly concave at t if $p_t > 1/2$, and has vanishing second derivative at t if $p_t = 1/2$. Formally:*

$$\begin{cases} N''(t) > 0 & \textit{if } p_t < 1/2, \\ N''(t) = 0 & \textit{if } p_t = 1/2, \\ N''(t) < 0 & \textit{if } p_t > 1/2. \end{cases}$$

Lemma 8 (Auxiliary function Γ and Bob's incentives)**.** *Define $\Gamma : [0,\infty) \to \mathbb{R}$ as*

$$\Gamma(t) = 1 - e^{-\lambda t} - \lambda t(1 - \gamma t). \tag{11}$$

Then $\Gamma(t)$ is initially concave and then convex, and $\Gamma(t) = 0$ admits a unique positive solution. Moreover, if a triple (s_1, α, t_1) with $t_1 > s_1 \geq 0$ satisfies $N(s_1) = 0$, then

$$N'(s_1) = 0 \iff \Gamma(t_1 - s_1) = 0,$$
$$N'(s_1) < 0 \iff \Gamma(t_1 - s_1) < 0.$$

Lemma 9 (Alice's expected payoffs from sampling up to t_1)**.** *Fix $p_0 \in (0,1)$ and an on-path play (s_1, α, t_1) with $0 \leq s_1 < t_1$ and $\alpha \in (0,1)$. Alice samples continuously and* *stops at $\tau \wedge t_1$, where τ is the first breakdown time (if any). Bob does not manipulate on $[0, s_1)$ and, at s_1, randomizes between full manipulation on $[s_1, t_1]$ with probability α and no manipulation on $[s_1, t_1]$ with probability $1-\alpha$. Conditional on reaching s_1 without a breakdown, Alice's continuation payoff at s_1 from sampling on $[s_1, t_1]$ is*

$$\hat{V}_{s_1}(s_1, \alpha, t_1) = (1-\tilde{p}_{s_1})+\tilde{p}_{s_1}(1-\alpha)\left(1-e^{-\lambda(t_1-s_1)}\right)\left(1-\frac{c}{\lambda}\right)-c(t_1-s_1)\left(1-\tilde{p}_{s_1}+\tilde{p}_{s_1}\alpha\right). \quad (12)$$

Alice's ex-ante expected payoff at $t=0$ from sampling until t_1 is

$$\begin{aligned} \hat{V}_0(s_1, \alpha, t_1) = (1-p_0)(1-ct_1) + p_0(1-\alpha)\left(1-e^{-\lambda t_1}\right)\left(1-\frac{c}{\lambda}\right) \\ + p_0\alpha\left[\left(1-e^{-\lambda s_1}\right)\left(1-\frac{c}{\lambda}\right) - c(t_1-s_1)e^{-\lambda s_1}\right]. \end{aligned} \quad (13)$$

Lemma 10 (Critical belief cutoff)**.** *Let $t_1^* > 0$ be the unique positive solution to $\Gamma(t) = 0$, where Γ is defined in (11), and define*

$$p^* \equiv \frac{\gamma t_1^*}{1-e^{-\lambda t_1^*}}. \quad (14)$$

Then $p^* > 1/2$.

Lemma 11 (Alice's sampling gain decreases in the prior)**.** *For each $p_0 \in (1/2, p^*]$, let $t_1 = t_1^{(p_0)}$ be the unique positive solution to Bob's indifference condition at $t = 0$,*

$$p_0(1-e^{-\lambda t_1}) = \gamma t_1.$$

Suppose Bob does not manipulate (i.e., $\alpha = 0$), and define Alice's gain from sampling on $[0, t_1]$, relative to stopping immediately, by

$$\Delta_0(p_0) \equiv \hat{V}_0(0, 0, t_1) - p_0 = (1-2p_0) + t_1\left(\left(1-\frac{c}{\lambda}\right)\gamma - (1-p_0)c\right).$$

Then $\Delta_0(p_0)$ is strictly decreasing in p_0 on $(1/2, p^*]$.

B.2 Proofs of Theorem 2 and Theorem 3

Before presenting the proofs of Theorems 2 and 3, we first prove three lemmas. They derive Alice's belief process (Lemma 12), characterize the necessary and sufficient conditions for the existence of a manipulation equilibrium with on-path play (s_1, α, t_1) (Lemma 13), and prove monotonicity of the equilibrium path (Lemma 14).

Lemma 12 (Alice's belief process)**.** *Let $p_0 \in (0,1)$, $s_1 \geq 0$, and $\alpha \in [0,1]$. Suppose that $\mathbb{P}[\phi_r = 1_{[s_1,\infty)}] = \alpha$ and $\mathbb{P}[\phi_r = 0] = 1-\alpha$. Let $(\tilde{p}_t)_{t\geq 0}$ be Alice's belief process given ϕ, where $\tilde{p}_t = \mathbb{P}[\omega = -1 \mid \tau > t]$.*

$$\dot{\tilde{p}}_t = \begin{cases} -\lambda \tilde{p}_t(1-\tilde{p}_t) & \text{if } t \in (0, s_1), \\ -\lambda \tilde{p}_t(1-\tilde{p}_t) + \lambda \frac{p_{s_1}\alpha}{1-p_{s_1}}(1-\tilde{p}_t)^2 & \text{if } t \in (s_1, \infty). \end{cases}$$

Proof. First, consider the case $s_1 = 0$. Denote by $q_t = \mathbb{P}[\omega = -1 \wedge \phi_r(t) = 0 \mid \tau > t]$ the probability that the state is low and Bob does not manipulate—the only case in which a breakdown can occur—given Alice's information. Then,

$$\dot{q}_t = -\lambda q_t(1-q_t) \quad \text{and} \quad q_0 = p_0(1-\alpha).$$

At $t = 0$, Bob's manipulation is independent of the state, but the two become correlated for $t > 0$ since a breakdown before t can only occur in the absence of manipulation. Conditional on no prior breakdown, Alice's belief that the state is low and Bob does not manipulate decreases over time. The likelihood ratios of the other three events (i.e., $\{\omega = -1 \wedge \phi_r(t) = 1\}$, $\{\omega = 1 \wedge \phi_r(t) = 0\}$, and $\{\omega = 1 \wedge \phi_r(t) = 1\}$) remain constant over time since no breakdown occurs if the state is high or Bob manipulates. Thus, since $p_0\alpha$ is the prior probability that the state is low and Bob manipulates, Alice's belief process $\tilde{p}_t = \mathbb{P}[\omega = -1 \mid \tau > t]$ is

$$\tilde{p}_t = q_t + (1-q_t) \cdot \frac{p_0\alpha}{1-p_0(1-\alpha)}. \tag{15}$$

After differentiating and simplifying, we obtain the belief dynamics.

$$\dot{\tilde{p}}_t = -\lambda \tilde{p}_t(1-\tilde{p}_t) + \lambda \frac{p_0\alpha}{1-p_0}(1-\tilde{p}_t)^2. \tag{16}$$

Second, consider the general case $s_1 \geq 0$. A standard derivation gives the law of motion for $\tilde{p}_t$ for $t \in (0, s_1)$. From s_1 onward, the process evolves as in the first case with prior p_{s_1}. Hence, replacing p_0 by p_{s_1} in (16) gives the statement. □

Lemma 13 (Incentive conditions for manipulation equilibria)**.** *Let $p_0 \in (0,1)$, $0 \leq s_1 < t_1$, and $\alpha \in (0,1)$. Then, a manipulation equilibrium with on-path play (s_1, α, t_1) exists for the prior p_0 if and only if*

1. *Alice has a weak incentive to initiate sampling at $t = 0$ and $\alpha \leq \overline{\alpha}$, where $\overline{\alpha}$ is defined in (3).*

2. Bob is indifferent at s_1, $N(s_1) = 0$, and moreover,

$$\begin{cases} N'(s_1) \le 0 & \textit{if} \quad s_1 = 0, \\ N'(s_1) = 0 & \textit{if} \quad s_1 > 0. \end{cases}$$

Proof. **Necessity.** Suppose (s_1, α, t_1) with $0 \le s_1 < t_1$ and $\alpha \in (0,1)$ is the on-path play of a manipulation equilibrium (σ, ϕ). Then Bob's necessary incentive condition (2) must hold, which immediately implies the second statement. Since $t_1 > 0$, Alice must weakly prefer initiating sampling at $t = 0$. It remains to show that $\alpha \le \overline{\alpha}$ is necessary for Alice to stop at $t_1 > 0$.

Consider first the case $s_1 = 0$, so that ϕ prescribes full-intensity manipulation on $[0, t_1]$ with probability α and no manipulation on $[0, t_1]$ with probability $1 - \alpha$. By Lemma 1, Alice's stopping belief satisfies $\tilde{p}_{t_1} \le 1/2$. Then Alice's expected payoff of stopping at $t \in [0, t_1]$ with $\tilde{p}_t \le 1/2$ is

$$(1 - p_0)(1 - ct) + p_0 \alpha\,(-ct) + p_0\,(1 - \alpha) \left[\int_0^t \lambda e^{-\lambda s}(1 - cs) ds + e^{-\lambda t}\,(-ct) \right].$$

The (left-)derivative of this function with respect to t must be weakly positive at the equilibrium stopping time t_1, since Alice weakly prefers sampling to stopping just before t_1. This gives

$$\alpha \le 1 - \frac{1}{p_0} \cdot \frac{c}{(\lambda - c)e^{-\lambda t_1} + c} \equiv \overline{\alpha}^{(p_0)}. \tag{17}$$

If $s_1 > 0$, the subgame starting at s_1 is identical to the game that begins at s_1 with prior p_{s_1}. In that case, (17) becomes $\alpha \le \overline{\alpha}^{(p_{s_1})}$. Therefore, a necessary condition for Alice to stop at t_1 is $\alpha \le \overline{\alpha}$.

Sufficiency. Assume the two conditions in the statement hold. Let $(s_1^{(p_0)}, \alpha^{(p_0)}, t_1^{(p_0)})$ be the triplet that satisfies these two conditions for a given p_0. To simplify notation, we omit the dependence of $(s_1^{(p_0)}, \alpha^{(p_0)}, t_1^{(p_0)})$ on p_0. Let $\sigma \equiv t_1$, and let ϕ prescribe: no manipulation on $[0, s_1)$; at time s_1, full manipulation on $[s_1, t_1]$ with probability $\alpha \in (0, \overline{\alpha}]$ and no manipulation on $[s_1, t_1]$ with probability $1 - \alpha$; and full manipulation on (t_1, ∞) if Alice has not stopped by t_1. We verify that (σ, ϕ) is an equilibrium.

Bob's incentives. By Lemma 2, it suffices to consider deviations to delay-and-bunch strategies, indexed by a switching time $\hat{s} \in [0, t_1]$: Bob does not manipulate on $[0, \hat{s})$ and manipulates at full intensity on $[\hat{s}, t_1]$. By definition, $N(\hat{s})$ is Bob's payoff gain from such a deviation relative to no manipulation on $[0, t_1]$. Hence Bob has no profitable deviation iff $N(\hat{s}) \le 0$ for all $\hat{s} \in [0, t_1]$ with equality at $\hat{s} = s_1$. Thus it remains to show that the conditions in the lemma imply these inequalities.

- If $s_1 = 0$, then $N(0) = 0$. If $p_0 \le 1/2$, Lemma 7 implies that N is convex on $[0, t_1]$; since $N(0) = N(t_1) = 0$, it follows that $N(t) \le 0$ for all $t \in [0, t_1]$. If instead $p_0 > 1/2$, let $\hat{t}$ satisfy $p_{\hat{t}} = 1/2$. Then N is concave on $[0, \hat{t}]$ and convex on $[\hat{t}, t_1]$. Because $N(0) = 0$ and $N'(0) \le 0$, concavity implies $N(t) \le 0$ on $[0, \hat{t}]$; then $N(\hat{t}) \le 0$, and convexity together with $N(t_1) = 0$ implies $N(t) \le 0$ on $[\hat{t}, t_1]$.

- If $s_1 > 0$, we first claim that $p_{s_1} > 1/2$. Otherwise, Lemma 7 implies that N is strictly convex on (s_1, t_1), and since $N(s_1) = N'(s_1) = 0$, this would imply $N(t_1) > 0$, a contradiction. Hence $p_{s_1} > 1/2$, so there exists $\hat{t} \in (s_1, t_1)$ such that $p_{\hat{t}} = 1/2$. Again, N is concave on $[0, \hat{t}]$ and convex on $[\hat{t}, t_1]$. Since $N(s_1) = N'(s_1) = 0$ and $s_1 < \hat{t}$, concavity implies $N(t) \le 0$ on both $[0, s_1]$ and $[s_1, \hat{t}]$; then convexity and $N(t_1) = 0$ imply $N(t) \le 0$ on $[\hat{t}, t_1]$.

Therefore, $N(\hat{s}) \le 0$ for all $\hat{s} \in [0, t_1]$ and Bob indeed has no incentive to deviate.

Alice's incentives when $s_1 = 0$. Since Bob manipulates with probability 1 after t_1, Alice stops immediately after t_1. It remains to show that Alice weakly prefers sampling for all $t \in [0, t_1)$.

First, we verify that Alice weakly prefers sampling just before t_1. As in the proof of Lemma 12, let $q_t = \mathbb{P}[\omega = -1, \phi_r(t) = 0 \mid \tau > t]$ denote the probability that the state is low and Bob does not manipulate, conditional on no breakdown before t. By Lemma 1, Alice takes the high action $a = 1$ for any $t \in [t_1 - \epsilon, t_1]$ for sufficiently small $\epsilon > 0$. Thus, Alice weakly prefers sampling for an additional dt unit at $t \in [t_1 - \epsilon, t_1)$ iff

$$1 - \tilde{p}_t \le -cdt + q_t \lambda dt + (1 - q_t \lambda dt)(1 - \tilde{p}_{t+dt}).$$

Simplifying and ignoring second-order terms yields the following condition:

$$c \le q_t \lambda \tilde{p}_t - \dot{\tilde{p}}_t. \tag{18}$$

The left-hand side is the marginal cost of sampling; the right-hand side is the marginal benefit, reflecting both the probability of a breakdown and the improvement in belief. By Bayes' rule

$$q_t = \frac{p_0(1-\alpha)e^{-\lambda t}}{1 - p_0(1-\alpha) + p_0(1-\alpha)e^{-\lambda t}},$$

which, together with Alice's belief expressions (15) and (16), yields

$$q_t \lambda \tilde{p}_t - \dot{\tilde{p}}_t = -\frac{\lambda p_0 \alpha}{1 - p_0} + \frac{p_0 \alpha + 1 - p_0}{1 - p_0} \lambda \tilde{p}_t.$$

Evaluating (18) at $t \uparrow t_1$ and rearranging gives $\alpha \le \overline{\alpha}$. Therefore, given $\alpha \le \overline{\alpha}$, (18) holds and Alice weakly prefers sampling just before t_1.

Next, we verify that Alice's gain from continuing to sample is strictly decreasing over time. We first rewrite (15) as

$$q_t = \left(1 + \frac{p_0\alpha}{1-p_0}\right)\tilde{p}_t - \frac{p_0\alpha}{1-p_0}.$$

We then use the above expression and Alice's belief motion (16) to express Alice's marginal benefit from continuing at time t (i.e., the right-hand side of (18)) as

$$\begin{aligned} q_t\lambda\tilde{p}_t - \dot{\tilde{p}}_t &= \left(\left(1 + \frac{p_0\alpha}{1-p_0}\right)\tilde{p}_t - \frac{p_0\alpha}{1-p_0}\right)\lambda\tilde{p}_t + \lambda\tilde{p}_t\left(1-\tilde{p}_t\right) - \frac{\lambda p_0\alpha}{1-p_0}\left(1-\tilde{p}_t\right)^2 \\ &= -\frac{\lambda p_0\alpha}{1-p_0} + \frac{p_0\alpha + 1 - p_0}{1-p_0}\lambda\tilde{p}_t, \end{aligned}$$

which is strictly decreasing in t, since $\tilde{p}_t$ is. Therefore, Alice prefers sampling to stopping for all $t \in [0, t_1)$, and hence stopping at t_1 is a best response for Alice.

Alice's incentives when $s_1 > 0$. The subgame starting at s_1 is equivalent to the game for the prior p_{s_1} and the on-path play $(0, \alpha, t_1 - s_1)$. By the previous case, Alice strictly prefers sampling in $[s_1, t_1)$, and strictly prefers stopping after t_1. Hence, it remains to show that Alice weakly prefers sampling in $[0, s_1)$.

Let $V(p)$ denote her continuation payoff on $[0, s_1]$ when the current belief is p and no breakdown has yet occurred, with boundary value $V(p_{s_1}) = \hat{V}_{s_1}(s_1, \alpha, t_1)$, where $\hat{V}_{s_1}(s_1, \alpha, t_1)$ is given by (12) in Lemma 9. Since Bob does not manipulate before s_1, standard dynamic programming gives

$$0 = -c + \lambda p\big(1 - V(p)\big) - \lambda p(1-p)V'(p).$$

From the Bob part we already know that $p_{s_1} > 1/2$, so stopping on $[0, s_1]$ yields payoff p. Let $D(p) \equiv V(p) - p$ denote Alice's gain from sampling at belief $p \in [p_{s_1}, p_0]$. Then

$$(1-p)D'(p) + D(p) = -\frac{c}{\lambda p},$$

or equivalently,

$$\frac{d}{dp}\left(\frac{D(p)}{1-p}\right) = -\frac{c}{\lambda p(1-p)^2} < 0.$$

Thus $D(p)/(1-p)$ is strictly decreasing in p. By assumption, Alice weakly prefers initiating sampling at $t = 0$, so $D(p_0) = V(p_0) - p_0 \geq 0$. Along the no-breakdown path before s_1, beliefs lie in $[p_{s_1}, p_0]$. Hence for every $p \in [p_{s_1}, p_0]$,

$$\frac{D(p)}{1-p} \geq \frac{D(p_0)}{1-p_0} \geq 0,$$

and therefore $D(p) \geq 0$. So Alice weakly prefers continuing to sample at every $t \in [0, s_1)$.

We have shown that both players are best responding. Hence (σ, ϕ) is a manipulation equilibrium with on-path play (s_1, α, t_1). □

Lemma 14 (Equilibrium-path monotonicity)**.** *Fix a prior $p \in (0,1)$ and suppose there exists a manipulation equilibrium with on-path play (s_1, α, t_1).*

(i) *If $s_1 = 0$, then $p \in (\gamma/\lambda, p^*]$, $t_1 = t_1^{(p)}$, where $t_1^{(p)} > 0$ solves $p(1 - e^{-\lambda t}) = \gamma t$, and $\alpha \in (0, \overline{\alpha}^{(p)}]$ where $\overline{\alpha}^{(p)}$ is given by (3). Moreover, $t_1^{(p)}$ is strictly increasing in p, and whenever $\overline{\alpha}^{(p)} > 0$, $\overline{\alpha}^{(p)}$ is strictly decreasing in p.*

(ii) *If $s_1 > 0$, then $t_1 - s_1 = t_1^*$, $p_{s_1} = p^* < p$, and $\alpha \leq \overline{\alpha}^{(p^*)}$.*

Proof. Part (i). If $s_1 = 0$, then by Lemma 13, t_1 must satisfy $N(0) = 0$ and $N'(0) \leq 0$. But $N(0) = 0$ admits a unique solution $t_1^{(p)} > 0$ iff $p > \gamma/\lambda$ (see Figure 2), so $t_1 = t_1^{(p)}$ and $p > \gamma/\lambda$. By definition of $t_1^{(p)}$, $p(1 - e^{-\lambda t_1^{(p)}}) = \gamma t_1^{(p)}$. Differentiating implicitly gives

$$\frac{dt_1^{(p)}}{dp} = \frac{1 - e^{-\lambda t_1}}{\gamma - \lambda p e^{-\lambda t_1}} = \frac{t_1(1 - e^{-\lambda t_1})}{p(1 - e^{-\lambda t_1} - \lambda t_1 e^{-\lambda t_1})} > 0, \tag{19}$$

because $1 - e^{-\lambda t} - \lambda t e^{-\lambda t}$ equals 0 at $t = 0$ and is strictly increasing in t for $t > 0$. By Lemma 8, $N'(0) \leq 0$ is equivalent to $\Gamma(t_1^{(p)}) \leq 0$. Since $t_1^{(p)}$ is strictly increasing in p and $t_1^* = t_1^{(p^*)}$ is the unique positive root of Γ, we must have $p \leq p^*$. Therefore, if $s_1 = 0$ then $p \in (\gamma/\lambda, p^*]$.

By Lemma 13, $\alpha \in (0, \overline{\alpha}^{(p)}]$ where $\overline{\alpha}^{(p)}$ is given by (3). It remains to show that, if $\overline{\alpha}^{(p)} > 0$, $p \mapsto \overline{\alpha}^{(p)}$ is strictly decreasing. Write $t = t_1^{(p)}$ and $D = (\lambda - c)e^{-\lambda t} + c$. Then $1 - \overline{\alpha}^{(p)} = c/(pD)$, and hence

$$\frac{d}{dp} \log(1 - \overline{\alpha}^{(p)}) = -\frac{1}{p} + \frac{(\lambda - c)\lambda e^{-\lambda t}}{D} \frac{dt}{dp}.$$

Now $\overline{\alpha}^{(p)} > 0$ is equivalent to

$$\frac{c}{\lambda - c} < \frac{p}{1-p} e^{-\lambda t} \iff D < \frac{(\lambda - c)e^{-\lambda t}}{1-p}.$$

Therefore,

$$\frac{d}{dp}\log(1-\overline{\alpha}^{(p)}) > -\frac{1}{p} + (1-p)\lambda\frac{dt}{dp}.$$

Substituting the expression (19) for dt/dp yields

$$-\frac{1}{p} + (1-p)\lambda\frac{dt}{dp} = -\frac{1}{p} + \frac{(1-p)\lambda t(1-e^{-\lambda t})}{p(1-e^{-\lambda t}-\lambda t e^{-\lambda t})} = \frac{-\Gamma(t)}{p(1-e^{-\lambda t}-\lambda t e^{-\lambda t})},$$

where the last equality uses $p(1-e^{-\lambda t}) = \gamma t$. Since $t = t_1^{(p)} \leq t_1^*$, we have $\Gamma(t) \leq 0$, and since $1 - e^{-\lambda t} - \lambda t e^{-\lambda t} > 0$ for $t > 0$, it follows that

$$\frac{d}{dp}\overline{\alpha}^{(p)} = -(1-\overline{\alpha}^{(p)})\frac{d}{dp}\log(1-\overline{\alpha}^{(p)}) < 0.$$

Thus $\overline{\alpha}^{(p)}$ is strictly decreasing wherever it is positive.

Part (ii). If $s_1 > 0$, it follows from Lemma 13 that $N(s_1) = N'(s_1) = 0$. Moreover, by Lemma 8, $N'(s_1) = 0$ iff $\Gamma(t_1 - s_1) = 0$. By definition, $\Gamma(t_1^*) = 0$ and $\Gamma(t) = 0$ admits a unique positive solution, so $t_1 - s_1 = t_1^*$. Since $N(s_1) = 0$, $p_{s_1} = p^*$. By Lemma 13 and (3), $\alpha \leq \overline{\alpha}^{(p^*)}$. □

Theorem 2 (Existence of manipulation equilibria)**.** *Let $\mathcal{P}$ be the set of priors p_0 for which a manipulation equilibrium exists, let $\underline{p} \equiv \gamma/\lambda$, and let $\overline{p} \equiv \sup \mathcal{P}$. Then $\mathcal{P}$ is nonempty, and no manipulation equilibrium exists for $p_0 \leq \underline{p}$. Moreover, $\mathcal{P}$ is either $(\underline{p}, \overline{p})$ or $(\underline{p}, \overline{p}]$.*

Proof. We proceed in four steps. The first step shows that there is a manipulation equilibrium for priors just above $\underline{p} \equiv \gamma/\lambda$ and hence $\mathcal{P}$ is non-empty. Moreover, there is no manipulation equilibrium for priors weakly below $\underline{p}$. The next three steps successively construct a manipulation equilibrium for priors above $\underline{p}$ but below $1/2$, above $1/2$ but below p^* (as defined in (14)), and above p^*.

Step 1. We first show that there is no manipulation equilibrium for priors $p_0 \leq \underline{p}$. This is true because, for a given prior p_0, Bob's indifference condition at $t = 0$,

$$p_0(1-e^{-\lambda t_1}) - \gamma t_1 = 0, \tag{20}$$

admits a unique strictly positive solution if and only if $p_0 > \underline{p}$. By Lemma 13, no manipulation equilibrium exists for priors $p_0 \leq \underline{p}$.

Next, we construct a manipulation equilibrium for every prior $p_0 > \underline{p}$ sufficiently close to $\underline{p}$, so that $\mathcal{P}$ is non-empty and $\overline{p}$ is well-defined. Fix a prior $p_0 > \underline{p}$ and let

$t_1^{(p_0)} > 0$ denote the unique positive solution to (20). We verify that, for sufficiently small $\alpha > 0$, the triple $(0, \alpha, t_1^{(p_0)})$ satisfies both conditions of Lemma 13.

The second condition of Lemma 13 requires $N(0) = 0$ and $N'(0) \leq 0$, where

$$N(t) = p_t(1 - e^{-\lambda(t_1 - t)}) - \gamma(t_1 - t).$$

The condition $N(0) = 0$ holds by definition of $t_1^{(p_0)}$. For p_0 sufficiently close to $\underline{p} = \gamma/\lambda$, we have $p_0 < 1/2$ because $\gamma < \lambda/2$. By Lemma 14, $\Gamma(t_1^*) = 0$, $t_1^* = t_1^{(p^*)}$ and $p^* > 1/2$. Therefore, $p_0 < p^*$. Since $t_1^{(p_0)}$ is increasing in p_0, $t_1^{(p_0)} < t_1^*$ and $\Gamma(t_1^{(p_0)}) < 0$. By Lemma 8, $\Gamma(t_1^{(p_0)}) < 0$ is equivalent to $N'(0) < 0$. Therefore, the condition $N'(0) < 0$ is also satisfied for $p_0 > \underline{p}$ and sufficiently close to $\underline{p}$.

It remains to verify that the first condition of Lemma 13 holds for small enough $\alpha \in (0, \overline{\alpha}^{(p_0)}]$. First, we verify that $\overline{\alpha}^{(p_0)} > 0$:

$$\overline{\alpha}^{(p_0)} = 1 - \frac{1}{p_0} \cdot \frac{c/\lambda}{(1 - c/\lambda)e^{-\lambda t_1^{(p_0)}} + c/\lambda} \to 1 - c/\gamma > 0$$

because $t_1^{(p_0)} \to 0$ as $p_0 \to \gamma/\lambda$ from above.

Next write $t_1 = t_1^{(p_0)}$ and let $\hat{V}_0(0, \alpha, t_1)$ denote Alice's expected payoff from sampling from $t = 0$ to t_1 when Bob mixes according to α at $s_1 = 0$. Since $p_0 \leq 1/2$, Alice's stopping payoff at p_0 is $1 - p_0$. Therefore, by Lemma 9, Alice's gain from sampling at p_0, relative to immediate stopping, is given by

$$\Delta_\alpha(p_0) \equiv \hat{V}_0(0, \alpha, t_1) - (1 - p_0) = t_1 \left((1 - \alpha) \left(1 - \frac{c}{\lambda} \right) \gamma - (1 - p_0 + p_0\alpha)c \right), \quad (21)$$

which is continuous and strictly decreasing in α. Evaluating it at $\alpha = \overline{\alpha}^{(p_0)}$ yields

$$\begin{aligned} \frac{\Delta_{\overline{\alpha}^{(p_0)}}(p_0)}{t_1} &= \frac{1}{p_0} \frac{c}{(\lambda - c)e^{-\lambda t_1} + c} \left(1 - \frac{c}{\lambda} \right) \gamma - \left(1 - \frac{c}{(\lambda - c)e^{-\lambda t_1} + c} \right) c \\ &= \frac{1}{p_0} \frac{c}{(\lambda - c)e^{-\lambda t_1} + c} \frac{\lambda - c}{\lambda} \left(\gamma - \lambda p_0 e^{-\lambda t_1} \right) \\ &= \frac{1}{p_0} \frac{c\gamma}{(\lambda - c)e^{-\lambda t_1} + c} \frac{\lambda - c}{\lambda} \frac{1 - e^{-\lambda t_1} - \lambda t_1 e^{-\lambda t_1}}{1 - e^{-\lambda t_1}} \\ &> 0 \end{aligned} \quad (22)$$

where the first equality uses (17), the last equality uses (20), and the inequality follows because $1 - e^{-\lambda t} - \lambda t e^{-\lambda t}$ equals 0 at $t = 0$ and is strictly increasing in t for $t > 0$. Therefore, the first condition of Lemma 13 is satisfied for every $\alpha \in (0, \overline{\alpha}^{(p_0)}]$. Thus, for each prior p_0 close to $\underline{p}$, there is a manipulation equilibrium with on-path play $(0, \alpha, t_1)$.

Step 2. We show that $\mathcal{P} \supseteq [p^*, \overline{p})$. Let $p_0 \in [p^*, \overline{p})$ and $\hat{p} > p_0$ with $\hat{p} \in \mathcal{P}$. Let $(\hat{s}_1, \hat{\alpha}, \hat{t}_1)$ be the on-path play of a manipulation equilibrium at $\hat{p}$. By Theorem 1 and Lemma 14, we have $\hat{s}_1 > 0$ and $p_{\hat{s}_1} = p^*$. For any $t \in [0, \hat{s}_1]$, the continuation game from t is the original game with prior p_t. Hence the shifted profile $(\hat{s}_1 - t, \hat{\alpha}, \hat{t}_1 - t)$ is a manipulation equilibrium for prior p_t. Since p_t ranges over $[p^*, \hat{p}]$ as t ranges over $[0, \hat{s}_1]$, we obtain existence for all priors in $[p^*, \hat{p}]$, and in particular for p_0.

Step 3. We show that $\mathcal{P} \supseteq (1/2, \hat{p})$ with $\hat{p} = \min\{p^*, \overline{p}\}$. Let $p_0 \in [1/2, \hat{p})$ and consider the candidate on-path play $(0, \alpha, t_1)$ with $t_1 = t_1^{(p_0)}$. Bob's incentive conditions, $N(0) = 0$ and $N'(0) \leq 0$, are satisfied by construction as in Step 1. For Alice's incentives, recall that by Lemma 14, $\overline{\alpha}^{(p)}$ is strictly monotonically decreasing in p, so that $\overline{\alpha}^{(p_0)} > 0$. Since $p_0 \geq 1/2$, Alice's stopping payoff at $t = 0$ is p_0. Therefore, by Lemma 9, Alice's gain from sampling from $t = 0$ to t_1 with $\alpha > 0$, relative to immediate stopping, is given by

$$\begin{aligned}\Delta_\alpha(p_0) &\equiv \hat{V}_0(0, \alpha, t_1^{(p_0)}) - p_0 \\ &= (1 - 2p_0) + t_1 \left((1 - \alpha) \left(1 - \frac{c}{\lambda} \right) \gamma - (1 - p_0 + p_0 \alpha) c \right). \end{aligned} \tag{23}$$

It is easy to see that $\Delta_\alpha(p_0)$ is continuous and strictly decreasing in α and $\Delta_1(p_0) < 0$.

Choose a sequence $p_n \uparrow \hat{p}$ with $p_n \in \mathcal{P}$, which exists by Step 2 if $\hat{p} = p^*$ and by definition of $\overline{p}$ if $\hat{p} = \overline{p}$. In any manipulation equilibrium at p_n, $\alpha_n > 0$ and $\Delta_{\alpha_n}(p_n) \geq 0$. Since $\Delta_\alpha(p_n)$ is strictly decreasing in α, we have $\Delta_0(p_n) > 0$. By continuity of $p \mapsto \Delta_0(p)$, we get $\Delta_0(\hat{p}) \geq 0$. Moreover, by Lemma 11, $\Delta_0'(p) < 0$ for all $p \in [1/2, \hat{p})$, and thus $\Delta_0(p_0) > 0$.

Since $\alpha \mapsto \Delta_\alpha(p_0)$ is continuous and strictly decreasing, $\Delta_1(p_0) < 0$, and $\Delta_0(p_0) > 0$, there exists $\hat{\alpha}^{(p_0)} > 0$ such that $\Delta_{\hat{\alpha}^{(p_0)}}(p_0) = 0$. Therefore, there is $\alpha \in \big(0, \min\{\hat{\alpha}^{(p_0)}, \overline{\alpha}^{(p_0)}\}\big]$ satisfying the first condition of Lemma 13.

Step 4. We show that $\mathcal{P} \supseteq (\underline{p}, \hat{p})$ with $\hat{p} = \min\{1/2, \overline{p}\}$. Let $p_0 \in (\underline{p}, \hat{p})$ and consider the candidate on-path play $(0, \alpha, t_1)$ with $t_1 = t_1^{(p_0)}$. Since $\overline{\alpha}^{(p)}$ is strictly decreasing in p by Lemma 14, we have $\overline{\alpha}^{(p_0)} > 0$. Since $p_0 \leq 1/2$, Alice's gain from starting sampling at $t = 0$, relative to immediate stopping, is given by (21) and inequality (22) remains valid. Therefore, there is $\alpha \in (0, \overline{\alpha}^{(p_0)}]$ satisfying the first condition of Lemma 13. The second condition of Lemma 13, $N(0) = 0$ and $N'(0) \leq 0$, follows from the same argument as in Step 1.

□

Theorem 3 (Uniqueness of manipulation equilibria). *Let $p_0 \in (\underline{p}, \overline{p})$. If $(\hat{\sigma}, \hat{\phi})$ and $(\tilde{\sigma}, \tilde{\phi})$*

are two manipulation equilibria with on-path play $(\hat{s}_1, \hat{\alpha}, \hat{t}_1)$ *and* $(\tilde{s}_1, \tilde{\alpha}, \tilde{t}_1)$*, respectively, then* $\hat{s}_1 = \tilde{s}_1$ *and* $\hat{t}_1 = \tilde{t}_1$.

Proof. We start with two claims.

Claim 1. If $\hat{s}_1 = \tilde{s}_1 = 0$, then $\hat{t}_1 = \tilde{t}_1$.

Proof. Any equilibrium (σ, ϕ) with on-path play $(0, \alpha, t_1)$ with $\alpha \in (0, 1)$ and $t_1 > 0$ satisfies $N(0) = 0$. If a positive solution $t_1 > 0$ to $N(0) = 0$ exists, it is unique because the function $t \mapsto p(1 - e^{-\lambda t}) - \gamma t$ is strictly concave on $(0, \infty)$ and equals 0 at $t = 0$, so it can cross zero at most once for $t > 0$. Hence, $\hat{t}_1 = \tilde{t}_1$. □

Claim 2. If $\hat{s}_1 \geq \tilde{s}_1 > 0$, then $\hat{s}_1 = \tilde{s}_1$ and $\hat{t}_1 = \tilde{t}_1$.

Proof. If (s_1, α, t_1) with $t_1 > s_1 > 0$ is the on-path play of a manipulation equilibrium, then it follows from Bob's incentive conditions that $N(s_1) = N'(s_1) = 0$. By Lemma 8, $N'(s_1) = 0$ is equivalent to $\Gamma(t_1 - s_1) = 0$, which has a unique positive solution $t_1 - s_1$. Hence, $\hat{t}_1 - \hat{s}_1 = \tilde{t}_1 - \tilde{s}_1$. Bob's indifference conditions imply that

$$p_{\hat{s}_1}(1 - e^{-\lambda(\hat{t}_1 - \hat{s}_1)}) = \gamma(\hat{t}_1 - \hat{s}_1) = \gamma(\tilde{t}_1 - \tilde{s}_1) = p_{\tilde{s}_1}(1 - e^{-\lambda(\tilde{t}_1 - \tilde{s}_1)}) = p_{\tilde{s}_1}(1 - e^{-\lambda(\hat{t}_1 - \hat{s}_1)}).$$

Hence, we have $p_{\hat{s}_1} = p_{\tilde{s}_1}$. Therefore, $\hat{s}_1 = \tilde{s}_1$ and $\hat{t}_1 = \tilde{t}_1$. □

It remains to rule out $\hat{s}_1 > \tilde{s}_1 = 0$. Assume for contradiction that $\hat{s}_1 > \tilde{s}_1 = 0$. By Lemma 8, we have $\Gamma(\hat{t}_1 - \hat{s}_1) = 0$ and $\Gamma(\tilde{t}_1) \leq 0$. It follows from the shape of the function Γ that $\hat{t}_1 - \hat{s}_1 \geq \tilde{t}_1$. Note that for all $t > 0$,

$$\frac{d}{dt}\left(\frac{\gamma t}{1 - e^{-\lambda t}}\right) = \frac{\gamma(1 - e^{-\lambda t} - \lambda t e^{-\lambda t})}{(1 - e^{-\lambda t})^2} > 0,$$

because $1 - e^{-\lambda t} - \lambda t e^{-\lambda t}$ equals 0 at $t = 0$ and is strictly increasing in t for $t > 0$. Thus, Bob's indifference conditions at $\hat{s}_1 > 0$ and at $\tilde{s}_1 = 0$ imply that

$$p_0 = \frac{\gamma \tilde{t}_1}{1 - e^{-\lambda \tilde{t}_1}} \leq \frac{\gamma(\hat{t}_1 - \hat{s}_1)}{1 - e^{-\lambda(\hat{t}_1 - \hat{s}_1)}} = p_{\hat{s}_1},$$

which is a contradiction. This completes the proof. □

C Omitted Proofs From Section 3.3

Lemma 5 (Early stopping under manipulation)**.** *Fix any prior* $p_0 \in (\underline{p}, \overline{p})$*, and let* $t_1(p_0)$ *be Alice's (unique) exit time in a manipulation equilibrium. Then,* $t_1(p_0) < t_1^w(p_0)$.

Proof. For the Wald problem without manipulation, Alice's belief process $(p_t)_t$ satisfies

$$p_t = \frac{p_0 e^{-\lambda t}}{1 - p_0 + p_0 e^{-\lambda t}},$$

and she stops at $\underline{p}^w = c/\lambda$. Hence, the stopping time $t_1^w = t_1^w(p_0)$ satisfies

$$\frac{p_0 e^{-\lambda t_1^w}}{1 - p_0 + p_0 e^{-\lambda t_1^w}} = \frac{c}{\lambda} \implies t_1^w = \frac{1}{\lambda} \log\left(\frac{p_0}{1-p_0} \frac{\lambda - c}{c}\right).$$

Let $s_1 = s_1^{(p_0)}$ and $t_1 = t_1^{(p_0)}$ be the boundaries of the manipulation interval in a manipulation equilibrium. Suppose, for contradiction, that $t_1^w \leq t_1$, so that $e^{-\lambda t_1^w} \geq e^{-\lambda t_1}$. If $s_1 = 0$, it follows from (3) that

$$\overline{\alpha} = 1 - \frac{1}{p} \cdot \frac{c}{(\lambda - c)e^{-\lambda t_1} + c} \leq 1 - \frac{1}{p} \cdot \frac{c}{(\lambda - c)e^{-\lambda t_1^w} + c} = 0,$$

contradicting $\overline{\alpha} > 0$. Thus $t_1^w > t_1$. The case $s_1 > 0$ follows by replacing p_0 by p_{s_1}. □

Theorem 4 ((Co)existence of a Wald equilibrium)**.** *The Wald equilibrium exists for any prior in* $(\underline{p}^w, \overline{p}]$.

Proof. Let $t_1^w = t_1^w(p_0)$, $t_1 = t_1(p_0)$, and $s_1 = s_1(p_0)$ be the boundaries of the manipulation period in a manipulation equilibrium. First consider the case $p_0 \in (\underline{p}, \overline{p})$. We show that no manipulation is optimal if Alice stops at t_1^w. If $s_1 = 0$, Bob's indifference at $t = 0$ in the manipulation equilibrium is equivalent to $N(0) = p_0(1 - e^{-\lambda t_1}) - \gamma t_1 = 0$. Since $p_0(1-e^{-\lambda t})$ is strictly concave in t while γt is linear, $p_0(1-e^{-\lambda t}) - \gamma t$ is strictly decreasing for $t > t_1$ (see Figure 2). Hence, if Alice stops at $t_1^w > t_1$, Bob strictly prefers not to manipulate at $t = 0$, and, by the shape of $N(t)$, at any $t \in (0, t_1^w)$. The proof is similar if $s_1 > 0$. Therefore, a Wald equilibrium exists whenever a manipulation equilibrium exists.

It remains to show that a Wald equilibrium exists for each $p_0 \in (\underline{p}^w, \underline{p}]$. For these priors, Bob has no incentive to manipulate for any stopping time $t > 0$ (see Figure 2), so Alice can sample as if Bob is absent. Hence, the Wald equilibrium exists. □

D Omitted Proofs from Section 5

Lemma 6 (Static sampling equilibrium)**.** *Let $p_0 \in (c/\lambda, p^*)$ with $p^* \leq \overline{p}$. With static sampling, Bob never manipulates in equilibrium, and Alice chooses*

$$t_a = \begin{cases} t_1 & \text{if} \quad \gamma < \widehat{\gamma} \\ t^s & \text{if} \quad \gamma \geq \widehat{\gamma} \end{cases} \tag{4}$$

where t_1 is Alice's stopping time in the manipulation equilibrium with sequential sampling, $t^s = (\log(\lambda p_0) - \log c)/\lambda$ is the optimal sample size for static sampling without manipulation, and $\widehat{\gamma}$ solves $G(t^s, \widehat{\gamma}) = 0$.

Proof. Consider first the static benchmark setting where Bob cannot manipulate and Alice chooses t^s to maximize her expected payoff:

$$p_0(1 - e^{-\lambda t^s}) + (1 - p_0) - ct^s.$$

The objective is concave, and the first-order condition yields $t^s = (\log(\lambda p_0) - \log c)/\lambda$.

Next, from $G(t^s, \widehat{\gamma}) = 0$, we compute the cutoff $\widehat{\gamma}$ as

$$\widehat{\gamma}(p_0, \lambda, c) = \frac{\lambda p_0 - c}{\log(\lambda p_0) - \log c}.$$

Therefore, if $\gamma \geq \widehat{\gamma}$ and $t_a \geq t^s$, Bob has no incentive to manipulate in static sampling.

Finally, as noted in the main text, Bob's best response is extremal with $t_b \in \{0, t_a\}$. Hence, Alice will choose $t_a > 0$ only if she can deter Bob from manipulating, i.e., induce $t_b = 0$. Therefore, Bob must not manipulate in equilibrium, whenever Alice chooses $t_a > 0$. Alice's static sampling problem can be written as

$$\max_{t_a \geq 0} \; p_0(1 - e^{-\lambda t_a}) - ct_a \quad \text{subject to} \quad p_0(1 - e^{-\lambda t_a}) \leq \gamma t_a.$$

Let $\mu \geq 0$ be the Lagrange multiplier. The first-order condition is:

$$\lambda p_0 e^{-\lambda t_a} - c + \mu(\gamma - \lambda p_0 e^{-\lambda t_a}) = 0.$$

If $\mu = 0$, the unconstrained optimum is $\lambda p_0 e^{-\lambda t_a} = c$, which implies $t_a = t^s$. The constraint is then reduced to $\gamma \geq \widehat{\gamma}$. If $\mu > 0$ (or equivalently, $\gamma < \widehat{\gamma}$), the constraint binds, $\gamma t_a = p_0(1 - e^{-\lambda t_a})$, and $t_a = t_1$. This completes the proof. □

Proposition 1 (Static vs sequential sampling)**.** *Let $p_0 \in (\underline{p}, 1/2)$. If $\lambda p_0 e^{-\lambda t_1} \geq c$, Alice's expected payoff in the static sampling equilibrium is strictly higher than in the* *worst manipulation equilibrium under sequential sampling, i.e., when Bob manipulates with probability $\overline{\alpha}$.*

Proof. Consider first sequential sampling. In Step 1 and Step 2 of the proof for Theorem 2, we have shown that, for $p_0 \leq 1/2$, the worst manipulation equilibrium for Alice has an on-path play $(0, \overline{\alpha}, t_1)$, where t_1 and $\overline{\alpha}$ must satisfy:

$$p_0(1 - e^{-\lambda t_1}) = \gamma t_1 \quad \text{and} \quad \overline{\alpha} = 1 - \frac{1}{p_0} \cdot \frac{c}{(\lambda - c)e^{-\lambda t_1} + c}. \tag{24}$$

Let $\hat{V}_0(0, \overline{\alpha}, t_1)$ be Alice's (worst) expected payoff from sequential sampling. Starting from (13) (with $s_1 = 0$), we have

$$\begin{aligned}
\hat{V}_0(\overline{\alpha}, t_1) &= (1 - p_0)(1 - ct_1) + p_0(1 - \overline{\alpha})(1 - e^{-\lambda t_1})\Big(1 - \frac{c}{\lambda}\Big) - p_0 \overline{\alpha} c t_1 \\
&= (1 - p_0) - ct_1\big(1 - p_0(1 - \overline{\alpha})\big) + p_0(1 - \overline{\alpha})(1 - e^{-\lambda t_1})\Big(1 - \frac{c}{\lambda}\Big) \\
&= (1 - p_0) - ct_1\left(\frac{(\lambda - c)e^{-\lambda t_1}}{(\lambda - c)e^{-\lambda t_1} + c}\right) + \frac{c(1 - e^{-\lambda t_1})}{(\lambda - c)e^{-\lambda t_1} + c}\Big(1 - \frac{c}{\lambda}\Big) \\
&= 1 - p_0 + \frac{c}{\lambda} \cdot \frac{(\lambda - c)\big(1 - e^{-\lambda t_1} - \lambda t_1 e^{-\lambda t_1}\big)}{(\lambda - c)e^{-\lambda t_1} + c},
\end{aligned}$$

where in the third line we use (24).

Now consider static sampling. Since $\lambda p_0 e^{-\lambda t_1} \geq c$, we have $t_1 \leq t^s$, and therefore Bob will not manipulate if Alice chooses t^s. Therefore, Alice's expected payoff from static sampling is maximized at $t_a = t^s$, which is weakly higher than what she receives by setting $t_a = t_1$. Alice's payoff from static sampling with sample size t_1 is

$$1 - p_0 + p_0(1 - e^{-\lambda t_1}) - ct_1 = 1 - p_0 + (\gamma - c)t_1.$$

where the equality follows from (24). Therefore, Alice strictly prefers static sampling if

$$(\gamma - c)t_1 > \frac{c}{\lambda} \cdot \frac{(\lambda - c)(1 - e^{-\lambda t_1} - \lambda t_1 e^{-\lambda t_1})}{(\lambda - c)e^{-\lambda t_1} + c}.$$

Multiplying through and rewriting again using (24) yields

$$\gamma(\lambda - c)(\lambda p_0 e^{-\lambda t_1} - c) + \lambda p_0(\gamma - c)c > 0,$$

which holds under the stated assumptions. □

References

Adusumilli, K. (2026): "How to Sample and When to Stop Sampling: The Generalised Wald Problem and Minimax Policies," *Review of Economic Studies*, 93(1), 1–34.

Aumann, R. J. (1964): "Mixed and behaviour policies in infinite extensive games," in *Advances in Game Theory*, ed. by M. Dresher, L. Shapley, and A. Tucker, vol. 52 of *Annals of Mathematics Studies*, chap. 28, pp. 627–650. Princeton University Press.

Bardhi, A. (2024): "Attributes: Selective learning and influence," *Econometrica*, 92(2), 311–353.

Board, S., and M. Meyer-ter Vehn (2013): "Reputation for Quality," *Econometrica*, 81(6), 2381–2462.

Bonatti, A., and J. Hörner (2011): "Collaborating," *American Economic Review*, 101(2), 632–663.

——— (2017a): "Career concerns with exponential learning," *Theoretical Economics*, 12(1), 425–475.

——— (2017b): "Learning to Disagree in a Game of Experimentation," *Journal of Economic Theory*, 169, 234–269.

Brandl, F., and X. Shi (2024): "Disinformation in the Wald Model," Working paper, University of Bonn and University of Toronto.

Brückner, M., and T. Scheffer (2011): "Stackelberg games for adversarial prediction problems," in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 547–555.

Cetemen, D., and C. Margaria (2024): "Dynamic signaling in Wald options," Working paper, Boston University.

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

Charikar, M., J. Steinhardt, and G. Valiant (2017): "Learning from untrusted data," in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 47–60.

Che, Y.-K., and K. Mierendorff (2019): "Optimal Dynamic Allocation of Attention," *American Economic Review*, 109(8), 2293–3029.

Crawford, V. P., and J. Sobel (1982): "Strategic information transmission," *Econometrica*, 50(6), 1431–1451.

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

Di Tillio, A., M. Ottaviani, and P. N. Sørensen (2021): "Strategic sample selection," *Econometrica*, 89(2), 911–953.

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

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

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

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

Fudenberg, D., P. Strack, and T. Strzalecki (2018): "Speed, Accuracy, and the Optimal Timing of Choices," *American Economic Review*, 108(12), 2993—3029.

Fudenberg, D., and J. Tirole (1986): "A "signal-jamming" theory of predation," *The RAND Journal of Economics*, pp. 366–376.

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

Hardt, M., N. Megiddo, C. Papadimitriou, and M. Wootters (2016): "Strategic classification," in *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*, pp. 111–122.

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

Holmström, B. (1999): "Managerial incentive problems: A dynamic perspective," *Review of Economic studies*, 66(1), 169–182.

Hörner, J., and L. Samuelson (2026): "What You Don't Know May Be Good For You," *American Economic Review*, 116(3), 1097–1147.

Kartik, N. (2009): "Strategic communication with lying costs," *Review of Economic Studies*, 76(4), 1359–1395.

Keller, G., and S. Rady (2015): "Strategic experimentation with Poisson bandits," *Theoretical Economics*, 5, 275–311.

Keller, G., S. Rady, and M. Cripps (2005): "Strategic experimentation with exponential bandits," *Econometrica*, 73(1), 39–68.

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

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

Spence, M. (1973): "Job Market Signaling," *The Quarterly Journal of Economics*, 87(3), 355–374.

Sun, Y. (2024): "A dynamic model of censorship," *Theoretical Economics*, 19(1), 29–60.

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

O ONLINE APPENDIX

O1 Proof of Auxiliary Results

Lemma 7 (Shape of Bob's manipulation gain)**.** *For all $t \in [0, t_1]$, $N(t)$ is strictly convex at t if $p_t < 1/2$, strictly concave at t if $p_t > 1/2$, and has vanishing second derivative at t if $p_t = 1/2$. Formally:*

$$\begin{cases} N''(t) > 0 & \text{if } p_t < 1/2, \\ N''(t) = 0 & \text{if } p_t = 1/2, \\ N''(t) < 0 & \text{if } p_t > 1/2. \end{cases}$$

Proof. Recall from (1) that

$$N(t) = p_t \left(1 - e^{-\lambda(t_1 - t)}\right) - \gamma(t_1 - t).$$

Differentiating, we obtain:

$$\begin{aligned} N'(t) &= -p_t(1-p_t)\lambda\left(1 - e^{-\lambda(t_1-t)}\right) - \lambda p_t e^{-\lambda(t_1-t)} + \gamma \\ &= -\lambda p_t + \lambda p_t^2 \left(1 - e^{-\lambda(t_1-t)}\right) + \gamma. \end{aligned}$$

On any history with no manipulation up to time t, Bob's belief follows the no-manipulation dynamics

$$p_t = \frac{p_0 e^{-\lambda t}}{1 - p_0 + p_0 e^{-\lambda t}},$$

so we can rewrite $N'(t)$ as:

$$N'(t) = \gamma - \lambda p_t \cdot \frac{1 - p_0 + p_0 e^{-\lambda t_1}}{1 - p_0 + p_0 e^{-\lambda t}}.$$

Differentiating again, and using $\dot{p}_t = -\lambda p_t(1-p_t)$, we obtain:

$$\begin{aligned} N''(t) &= -\lambda \dot{p}_t \cdot \frac{1 - p_0 + p_0 e^{-\lambda t_1}}{1 - p_0 + p_0 e^{-\lambda t}} + \lambda p_t \cdot \frac{1 - p_0 + p_0 e^{-\lambda t_1}}{(1 - p_0 + p_0 e^{-\lambda t})^2} \cdot (-\lambda p_0 e^{-\lambda t}) \\ &= \lambda^2 p_t(1-p_t) \cdot \frac{1 - p_0 + p_0 e^{-\lambda t_1}}{1 - p_0 + p_0 e^{-\lambda t}} - \lambda^2 p_t^2 \cdot \frac{1 - p_0 + p_0 e^{-\lambda t_1}}{1 - p_0 + p_0 e^{-\lambda t}} \\ &= \lambda^2 p_t(1-2p_t) \cdot \frac{1 - p_0 + p_0 e^{-\lambda t_1}}{1 - p_0 + p_0 e^{-\lambda t}}. \end{aligned}$$

The sign of $N''(t)$ is thus determined by the sign of $1 - 2p_t$, proving the claim. □

Lemma 8 (Auxiliary function Γ and Bob's incentives). *Define* $\Gamma : [0, \infty) \to \mathbb{R}$ *as*

$$\Gamma(t) = 1 - e^{-\lambda t} - \lambda t(1 - \gamma t). \tag{11}$$

Then $\Gamma(t)$ *is initially concave and then convex, and* $\Gamma(t) = 0$ *admits a unique positive solution. Moreover, if a triple* (s_1, α, t_1) *with* $t_1 > s_1 \geq 0$ *satisfies* $N(s_1) = 0$*, then*

$$N'(s_1) = 0 \iff \Gamma(t_1 - s_1) = 0,$$
$$N'(s_1) < 0 \iff \Gamma(t_1 - s_1) < 0.$$

Proof. It is straightforward to verify that $\Gamma(t)$ satisfies: (i) $\Gamma(0) = 0$, $\Gamma'(0) = 0$, and $\lim_{t\to\infty} \Gamma(t) = \infty$; (ii) $\Gamma(t)$ is initially concave and then convex (since $\gamma < \lambda/2$). Hence, the equation $\Gamma(t) = 0$ has a unique positive solution.

From the proof of Lemma 7, the conditions $N(s_1) = 0$ and $N'(s_1) < 0$ can be rewritten as

$$p_{s_1}\big(1 - e^{-\lambda(t_1 - s_1)}\big) = \gamma(t_1 - s_1),$$

and

$$-\lambda p_{s_1} + \lambda p_{s_1}^2\big(1 - e^{-\lambda(t_1 - s_1)}\big) + \gamma < 0.$$

Using the expression for p_{s_1} from the first equation, the condition $N'(s_1) < 0$ becomes

$$-\frac{\lambda\gamma(t_1 - s_1)}{1 - e^{-\lambda(t_1 - s_1)}}\big(1 - \gamma(t_1 - s_1)\big) + \gamma < 0,$$

which is equivalent to $\Gamma(t_1 - s_1) < 0$. The equivalence between $N'(s_1) = 0$ and $\Gamma(t_1 - s_1) = 0$ follows analogously. ☐

Lemma 9 (Alice's expected payoffs from sampling up to t_1). *Fix* $p_0 \in (0, 1)$ *and an on-path play* (s_1, α, t_1) *with* $0 \leq s_1 < t_1$ *and* $\alpha \in (0, 1)$*. Alice samples continuously and stops at* $\tau \wedge t_1$*, where* τ *is the first breakdown time (if any). Bob does not manipulate on* $[0, s_1)$ *and, at* s_1*, randomizes between full manipulation on* $[s_1, t_1]$ *with probability* α *and no manipulation on* $[s_1, t_1]$ *with probability* $1 - \alpha$*. Conditional on reaching* s_1 *without a breakdown, Alice's continuation payoff at* s_1 *from sampling on* $[s_1, t_1]$ *is*

$$\hat{V}_{s_1}(s_1, \alpha, t_1) = (1 - \tilde{p}_{s_1}) + \tilde{p}_{s_1}(1 - \alpha)\left(1 - e^{-\lambda(t_1 - s_1)}\right)\left(1 - \frac{c}{\lambda}\right) - c(t_1 - s_1)\big(1 - \tilde{p}_{s_1} + \tilde{p}_{s_1}\alpha\big). \tag{12}$$

Alice's ex-ante expected payoff at $t = 0$ from sampling until t_1 is

$$\hat{V}_0(s_1, \alpha, t_1) = (1 - p_0)(1 - ct_1) + p_0(1 - \alpha)\left(1 - e^{-\lambda t_1}\right)\left(1 - \frac{c}{\lambda}\right) + p_0\alpha\left[\left(1 - e^{-\lambda s_1}\right)\left(1 - \frac{c}{\lambda}\right) - c(t_1 - s_1)e^{-\lambda s_1}\right]. \tag{13}$$

Proof. Let $\tau \sim \text{Exp}(\lambda)$ be the breakdown time in the low state under no manipulation. Suppose s_1 is reached and consider Alice's continuation payoff at s_1.

- If the state is high (with probability $1 - \tilde{p}_{s_1}$), no breakdown occurs, and Alice samples for $(t_1 - s_1)$ units and receives payoff 1 at t_1, so her payoff is $1 - c(t_1 - s_1)$.
- If the state is low (with probability $\tilde{p}_{s_1}$) and Bob does not manipulate, the residual time to breakdown is distributed according to $\text{Exp}(\lambda)$. Alice obtains payoff 1 iff a breakdown occurs within $(t_1 - s_1)$, which has probability $1 - e^{-\lambda(t_1 - s_1)}$. Her expected sampling time is $\mathbb{E}[\min\{\tau, t_1 - s_1\}] = \int_0^{t_1 - s_1} e^{-\lambda t} dt = (1 - e^{-\lambda(t_1 - s_1)})/\lambda$. Hence her expected payoff in this event is $\left(1 - e^{-\lambda(t_1 - s_1)}\right)(1 - c/\lambda)$.
- If the state is low and Bob manipulates at full intensity on $[s_1, t_1]$, no breakdown occurs on that interval. Alice samples for $(t_1 - s_1)$ units and then chooses action 1, yielding payoff 0. Hence, her payoff in this event is therefore $-c(t_1 - s_1)$.

Weighting these payoffs by the event probabilities yields (12).

Next consider Alice's ex-ante expected payoff at $t = 0$ from sampling until t_1.

- If the state is high (with probability $1 - p_0$), no breakdown occurs. Alice samples till t_1 and obtains payoff $1 - ct_1$.
- If the state is low (with probability p_0) and Bob does not manipulate, Alice obtains payoff 1 iff a breakdown occurs by t_1, which has probability $1 - e^{-\lambda t_1}$, and her expected sampling time is $(1 - e^{-\lambda t_1})/\lambda$. Thus her expected payoff is $\left(1 - e^{-\lambda t_1}\right)(1 - c/\lambda)$.
- If the state is low and Bob manipulates only on $[s_1, t_1]$, then a breakdown can occur only before s_1. Alice obtains payoff 1 iff $\tau \le s_1$, which has probability $1 - e^{-\lambda s_1}$. Her expected sampling time is $(\mathbb{E}[\min\{\tau, s_1\}] + (t_1 - s_1)\mathbb{P}(\tau > s_1))$. Hence, her expected payoff is

$$\left(1 - e^{-\lambda s_1}\right) - c\left(\frac{1 - e^{-\lambda s_1}}{\lambda} + (t_1 - s_1)e^{-\lambda s_1}\right) = \left(1 - e^{-\lambda s_1}\right)\left(1 - \frac{c}{\lambda}\right) - c(t_1 - s_1)e^{-\lambda s_1}.$$

Weighting these payoffs by the event probabilities yields (13). □

Lemma 10 (Critical belief cutoff)**.** *Let* $t_1^* > 0$ *be the unique positive solution to* $\Gamma(t) = 0$*, where* Γ *is defined in* (11)*, and define*

$$p^* \equiv \frac{\gamma t_1^*}{1 - e^{-\lambda t_1^*}}. \tag{14}$$

Then $p^* > 1/2$*.*

Proof. Define $a \equiv \gamma/\lambda \in (0, 1/2)$ and $x \equiv \lambda t$. Then

$$\Gamma(t) = 0 \iff g(x) \equiv 1 - e^{-x} - x(1 - ax) = 0,$$

and $x^* \equiv \lambda t_1^*$ is the unique positive root of g. Moreover,

$$p^* = \frac{\gamma\, t_1^*}{1 - e^{-\lambda t_1^*}} = \frac{ax^*}{1 - e^{-x^*}} = \frac{a}{1 - ax^*}.$$

Therefore, to show $p^* > 1/2$, it is sufficient to show $x^* > x_0 \equiv 1/a - 2$. Since $g(0) = 0$ and x^* is the unique positive root of g, it is sufficient to show $g(x_0) < 0$. Note that

$$g(x_0) = 4a - 1 - e^{-x_0}.$$

If $a \leq 1/4$, then $g(x_0) < 0$. If $a \in (1/4, 1/2)$, consider

$$Q(a) \equiv 2 - 1/a - \log(4a - 1).$$

Then $Q(1/2) = 0$, $\lim_{a \downarrow 1/4} Q(a) = \infty$, and $Q'(a) \leq 0$. Hence $Q(a) > 0$ for all $a \in (1/4, 1/2)$. It follows that $-x_0 > \log(4a-1)$ and thus $g(x_0) < 0$. Therefore, $p^* > 1/2$. □

Lemma 11 (Alice's sampling gain decreases in the prior)**.** *For each* $p_0 \in (1/2, p^*]$*, let* $t_1 = t_1^{(p_0)}$ *be the unique positive solution to Bob's indifference condition at* $t = 0$*,*

$$p_0(1 - e^{-\lambda t_1}) = \gamma t_1.$$

Suppose Bob does not manipulate (i.e., $\alpha = 0$*), and define Alice's gain from sampling on* $[0, t_1]$*, relative to stopping immediately, by*

$$\Delta_0(p_0) \equiv \hat{V}_0(0, 0, t_1) - p_0 = (1 - 2p_0) + t_1\Big(\Big(1 - \frac{c}{\lambda}\Big)\gamma - (1 - p_0)c\Big).$$

Then $\Delta_0(p_0)$ *is strictly decreasing in* p_0 *on* $(1/2, p^*]$*.*

Proof. Let $x \equiv \lambda t_1$. Then Bob's indifference condition becomes

$$p_0(1-e^{-x}) = \frac{\gamma}{\lambda}x, \qquad \text{so} \qquad p_0 = \frac{\gamma}{\lambda}\cdot\frac{x}{1-e^{-x}}.$$

Since

$$\frac{dp_0}{dx} = \frac{\gamma}{\lambda}\cdot\frac{1-(1+x)e^{-x}}{(1-e^{-x})^2} > 0,$$

it is enough to show that

$$\Phi(x) \equiv \Delta_0(p_0(x))$$

is strictly decreasing in x.

We use $x \equiv \lambda t_1$ and the definition of Δ_0 to obtain

$$\Phi(x) = 1 - \frac{2\gamma x}{\lambda(1-e^{-x})} + \frac{x}{\lambda}\left[\left(1-\frac{c}{\lambda}\right)\gamma - c + \frac{c\gamma x}{\lambda(1-e^{-x})}\right].$$

Differentiating and simplifying gives

$$\begin{aligned}(1-e^{-x})^2\Phi'(x) = &-\frac{\gamma}{\lambda}\big(1-2xe^{-x}-e^{-2x}\big)\\ &+\frac{c}{\lambda}\left[\frac{\gamma}{\lambda}\big(2x(1-e^{-x})-x^2e^{-x}-(1-e^{-x})^2\big)-(1-e^{-x})^2\right]. \quad (25)\end{aligned}$$

We can further rewrite (25) as

$$\begin{aligned}(1-e^{-x})^2\Phi'(x) = &-\frac{c}{\lambda}\Big[1-2xe^{-x}-e^{-2x}\\ &\quad -\frac{\gamma}{\lambda}\big(2x(1-e^{-x})-x^2e^{-x}-(1-e^{-x})^2\big)+(1-e^{-x})^2\Big]\\ &-\frac{\gamma-c}{\lambda}\big(1-2xe^{-x}-e^{-2x}\big).\end{aligned}$$

We first note that

$$e^{2x}\big(1-2xe^{-x}-e^{-2x}\big) = e^{2x}-2xe^{x}-1,$$

and the derivative of the right-hand side is

$$2e^{x}(e^{x}-x-1) > 0$$

for every $x > 0$, while its value at $x = 0$ is 0. Hence,

$$e^{2x}\big(1-2xe^{-x}-e^{-2x}\big) > 0 \implies 1-2xe^{-x}-e^{-2x} > 0 \qquad \text{for all } x > 0.$$

Therefore, $\Phi'(x) < 0$ if we can show that

$$1 - 2xe^{-x} - e^{-2x} - \frac{\gamma}{\lambda}\big(2x(1-e^{-x}) - x^2e^{-x} - (1-e^{-x})^2\big) + (1-e^{-x})^2 > 0. \tag{26}$$

Now consider the left-hand side of (26):

$$\begin{aligned} &1 - 2xe^{-x} - e^{-2x} - \frac{\gamma}{\lambda}\big(2x(1-e^{-x}) - x^2e^{-x} - (1-e^{-x})^2\big) + (1-e^{-x})^2 \\ &= 2\big(1-(1+x)e^{-x}\big) + \frac{\gamma}{\lambda}\big(x^2e^{-x} - 2x(1-e^{-x}) + (1-e^{-x})^2\big) \\ &> 2\big(1-(1+x)e^{-x}\big) + \frac{\gamma}{\lambda}\big(-2x\big(1-(1+x)e^{-x}\big)\big) \\ &= 2\big(1-(1+x)e^{-x}\big)\left(1 - \frac{\gamma}{\lambda}x\right), \end{aligned} \tag{27}$$

where the inequality follows because

$$x^2e^{-x} - 2x(1-e^{-x}) + (1-e^{-x})^2 + 2x(1-(1+x)e^{-x}) = e^{-x}\big(e^x + e^{-x} - 2 - x^2\big) > 0.$$

The first factor in the lower bound (27) is strictly positive for every $x > 0$, since $(1 + x)e^{-x} < 1$. Moreover, since $p_0 \le p^*$ and $t_1(p_0)$ is strictly increasing in p_0, we have $x \equiv \lambda t_1 \le \lambda t_1^*$. The inequality $x \le \lambda t_1^*$ and the defining equation of t_1^*,

$$1 - e^{-\lambda t_1^*} = \lambda t_1^*(1 - \gamma t_1^*),$$

imply that

$$1 - \frac{\gamma}{\lambda}x \ge 1 - \gamma t_1^* = \frac{1 - e^{-\lambda t_1^*}}{\lambda t_1^*} > 0.$$

Thus, the lower bound (27) is strictly positive, so (26) holds. By (25), $\Phi'(x) < 0$ for all x corresponding to $p_0 \in (1/2, p^*]$, and therefore $\Delta_0'(p_0) < 0$ for all $p_0 \in (1/2, p^*]$. □

O2 Manipulation-Threat Equilibrium

Proposition 6 (Deadline-deterrence equilibrium)**.** *Fix $p_0 \in (\underline{p}, \overline{p})$. Let t_1^w denote the Wald stopping time and let t_1 be the stopping time in a manipulation equilibrium (σ, ϕ). Let $(\tilde{\sigma}, \tilde{\phi})$ be a deadline-deterrence equilibrium with on-path play $(\tilde{s}_1, \tilde{\alpha}, \tilde{t}_1)$, where $\tilde{s}_1 = \tilde{t}_1 > 0$ and $\tilde{\alpha} = 0$. Then $\tilde{t}_1 \in [t_1, t_1^w)$.*

To see the intuition for the longer sampling period, consider a manipulation equilibrium without a passive period. Bob is indifferent at $s_1 = 0$ if sampling lasts until t_1. The function $t \mapsto p_0(1 - e^{-\lambda t}) - \gamma t$ comparing the benefit and the cost of manipulation is positive for all $t \in (0, t_1)$ (cf. Figure 2). Hence, in a deadline-deterrence equilibrium with exit before t_1, Bob would strictly gain by manipulating immediately at $s_1 = 0$, contradicting the premise that he refrains from manipulation on path. The argument for manipulation equilibria with a passive period is similar.

Proof. It is clear that $\tilde{t}_1 \leq t_1^w$, because otherwise Alice would exit before $\tilde{t}_1$. Assume for contradiction that $\tilde{t}_1 < t_1$. Let

$$\tilde{N}(t) = \tilde{p}_t(1 - e^{-\lambda(\tilde{t}_1 - t)}) - \gamma(\tilde{t}_1 - t),$$

where $\tilde{p}_t = p_t$ along the on-path history because there is no manipulation on path in a deadline-deterrence equilibrium. Let (s_1, α, t_1) be the on-path play for (σ, ϕ). Since (σ, ϕ) is an equilibrium, $p_{s_1}(1 - e^{-\lambda(t_1 - t)}) - \gamma(t_1 - t) \geq 0$ for all $t \in [s_1, t_1]$ (since sustaining manipulation is optimal when initiating manipulation at s_1). This inequality is strict for $t \in (s_1, t_1)$ since the left-hand side is strictly concave in the remaining time $t_1 - t$ (equivalently, strictly concave in t). Hence, for $s_1^* = \max\{0, \tilde{t}_1 - (t_1 - s_1)\}$, $0 \leq s_1^* \leq s_1$, and

$$\begin{aligned}
\tilde{N}(s_1^*) &= p_{s_1^*}(1 - e^{-\lambda(\tilde{t}_1 - s_1^*)}) - \gamma(\tilde{t}_1 - s_1^*) \\
&\geq p_{s_1}(1 - e^{-\lambda(\tilde{t}_1 - s_1^*)}) - \gamma(\tilde{t}_1 - s_1^*) \\
&= p_{s_1}(1 - e^{-\lambda(t_1 - (t_1 - \tilde{t}_1 + s_1^*))}) - \gamma(t_1 - (t_1 - \tilde{t}_1 + s_1^*)) \\
&\geq 0.
\end{aligned}$$

The first inequality is strict if $s_1^* = \tilde{t}_1 - (t_1 - s_1)$ (and so $s_1^* < s_1$), and the second inequality is strict if $s_1^* = 0 > \tilde{t}_1 - (t_1 - s_1)$ (and so $t_1 - \tilde{t}_1 + s_1^* \in (s_1, t_1)$). Hence, $\tilde{N}(s_1^*) > 0$. This contradicts that $(\tilde{\sigma}, \tilde{\phi})$ is an equilibrium since Bob strictly prefers manipulating at $s_1^* < \tilde{t}_1$. □

O3 Hybrid Approach Dominates both Pure Static and Pure Sequential Sampling

Proposition 2 (Hybrid sampling). *Let $[s_1, t_1]$ be the manipulation period in a manipulation equilibrium for some prior p_0, and let t^s be the optimal sample size under static sampling without manipulation.*

(i) If $p_0 \geq p^*$ and $\lambda p_0 e^{-\lambda t_1} > c$, Alice strictly prefers $[0,\, t^s - t_1] \oplus [t^s - t_1,\, t^s]$ to static sampling with sample size t^s in equilibrium.

(ii) Suppose a manipulation equilibrium with on-path play $(0, \alpha, t_1^*)$ exists under se-

quential sampling, and Alice strictly prefers static sampling to this equilibrium.[20] *Then, if* $p_0 \geq p^*$, *Alice strictly prefers the equilibrium under* $[0,\, s_1] \oplus [s_1,\, t_1]$ *to the manipulation equilibrium with on-path play* (s_1, α, t_1).

Proof. First we prove part (i). We begin by showing that Bob does not manipulate under the hybrid strategy. By definition of t_1, Bob weakly prefers not to manipulate at $t = 0$ in sequential sampling:

$$p_0(1 - e^{-\lambda t_1}) - \gamma t_1 \leq 0.$$

In the confirmatory phase $[t^s - t_1, t^s]$, Bob's gain from full manipulation is

$$p_{t^s - t_1}(1 - e^{-\lambda t_1}) - \gamma t_1 < p_0(1 - e^{-\lambda t_1}) - \gamma t_1 \leq 0.$$

In the exploratory phase $[0, t^s - t_1]$, Bob would optimally backload any manipulation. Thus, we consider his gain from manipulating the interval $[t, t^s - t_1]$:

$$\begin{aligned}
\hat{N}(t) &= p_t(1 - e^{-\lambda(t^s - t_1 - t)})e^{-\lambda t_1} - \gamma(t^s - t_1 - t) \\
&= p_t(e^{-\lambda t_1} - e^{-\lambda(t^s - t)}) - \gamma(t^s - t_1 - t) \\
&< p_0(e^{-\lambda t_1} - e^{-\lambda(t^s - t)}) - \gamma(t^s - t_1 - t) \\
&= p_0(1 - e^{-\lambda(t^s - t)}) - \gamma(t^s - t) - \left[p_0(1 - e^{-\lambda t_1}) - \gamma t_1\right],
\end{aligned}$$

where the inequality uses $p_t < p_0$. Since $t_1 < t^s$, we have

$$\lambda p_0 e^{-\lambda t_1} - \gamma < \lambda p_0 e^{-\lambda t^s} - \gamma < \lambda p_0 e^{-\lambda t^s} - c = 0,$$

which implies for $x \in [t_1, t^s]$,

$$\frac{\partial}{\partial x}\left[p_0(1 - e^{-\lambda x}) - \gamma x\right] = \lambda p_0 e^{-\lambda x} - \gamma < 0.$$

Since $t^s - t \geq t_1$ for all $t \in [0, t^s - t_1]$, it follows that

$$p_0(1 - e^{-\lambda(t^s - t)}) - \gamma(t^s - t) \leq p_0(1 - e^{-\lambda t_1}) - \gamma t_1,$$

and hence $\hat{N}(t) \leq 0$ for all $t \in [0, t^s - t_1]$. Therefore, Bob does not manipulate in either phase.

Since Alice's belief at any time depends only on cumulative experimentation time,

[20] We do not have simple sufficient conditions for this ranking. It is easy to find parameter values to satisfy the assumed properties. See Example 3 in the Online Appendix.

which is identical across both strategies conditional on no breakdown, the hybrid strategy strictly dominates static sampling by allowing earlier termination if a breakdown occurs during the initial phase $[0, t^s - t_1]$.

Second, we prove part (ii). We again show that Bob has no incentive to manipulate under the hybrid strategy. For $t < s_1$, Bob's gain from full manipulation over the interval $[t, s_1]$, relative to no manipulation, is given by[21]

$$\begin{aligned}
\tilde{N}(t) &= p_t \left(1 - e^{-\lambda(s_1 - t)}\right) e^{-\lambda(t_1 - s_1)} - \gamma(s_1 - t) \\
&= p_t \left(e^{-\lambda(t_1 - s_1)} - e^{-\lambda(t_1 - t)}\right) - \gamma(s_1 - t) \\
&= N(t) - \left[p_t(1 - e^{-\lambda(t_1 - s_1)}) - \gamma(t_1 - s_1)\right] \\
&< N(t) - \left[p_{s_1}(1 - e^{-\lambda(t_1 - s_1)}) - \gamma(t_1 - s_1)\right] = N(t),
\end{aligned}$$

where the inequality follows from the fact that $p_t > p_{s_1}$ for all $t < s_1$. Since $N(t) \leq 0$ for all $t \in [0, t_1]$, it follows that $\tilde{N}(t) < 0$ for all $t \in [0, s_1]$. In the confirmatory phase $[s_1, t_1]$, Bob's gain from full manipulation is

$$N(s_1) = p_{s_1}(1 - e^{-\lambda(t_1 - s_1)}) - \gamma(t_1 - s_1) = 0.$$

Thus, Bob does not manipulate in either phase.

Because Bob is passive up to time s_1 under both the hybrid protocol and the manipulation equilibrium $(s_1, \alpha, t_1^* + s_1)$, Alice's payoff on $[0, s_1]$ is the same in both. On the no-breakdown path at time s_1, the belief equals the cutoff p^*, so the continuation payoff under the hybrid protocol coincides with the payoff from static sampling of length t_1^* at prior p^*. In contrast, the continuation under the manipulation equilibrium coincides with the manipulation equilibrium continuation $(0, \alpha, t_1^*)$ at prior p^*, shifted to start at s_1. By assumption, Alice strictly prefers static sampling to that manipulation equilibrium at p^*. Therefore the hybrid protocol yields Alice a strictly higher payoff than the corresponding manipulation equilibrium for every $p_0 \geq p^*$. □

O4 Manipulation Equilibrium with Slightly Convex Cost

Proposition 5 (Manipulation equilibria under slightly convex costs)**.** *Let (σ^0, ϕ^0) be a manipulation equilibrium with on-path play (s_1^0, α^0, t_1^0). Then, there exists $\bar{\varepsilon} > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon})$, the game with manipulation cost function $k_\varepsilon(b) = \gamma b + \varepsilon b^2$ admits an equilibrium $(\sigma^\varepsilon, \phi^\varepsilon)$ such that along the no-breakdown history, Alice stops at t_1^ε, and Bob mixes between no manipulation (probability $1 - \alpha^\varepsilon$) and a pure strategy β^ε*

[21] Again, since Bob has an incentive to delay, it is sufficient to show that Bob cannot gain from full manipulation relative to no manipulation at $t \in [0, s_1]$.

(probability α^ε), where $\beta^\varepsilon(t) = 0$ for $t < s_1^\varepsilon$ and β^ε is weakly increasing on $[s_1^\varepsilon, t_1^\varepsilon]$.

Proof. Fix parameters (c, γ, λ) and a prior p_0 such that under linear manipulation costs ($\varepsilon = 0$) there exists a manipulation equilibrium (σ^0, ϕ^0) with on-path play (s_1^0, α^0, t_1^0). Fix $\varepsilon > 0$ and let the flow cost be $k_\varepsilon(b) = \gamma b + \varepsilon b^2$. We construct a manipulation equilibrium $(\sigma^\varepsilon, \phi^\varepsilon)$ for all sufficiently small ε.

Step 0 (No early stopping). Lemma 5 is unaffected by the perturbation because only Bob's payoff changes. In particular, if σ is a best response to ϕ, then

$$\mathbb{P}\big[\sigma > 0 \ \wedge \ \tilde{p}_\sigma > \tfrac{1}{2}\big] = 0.$$

Hence, whenever Alice exits at some $t > 0$ absent a breakdown, she takes action $a = 1$.

Step 1 (Backloading in Bob's best responses on a local deadline interval). Fix any sampling strategy σ whose support is a finite set $\{t_1, \ldots, t_k\}$ with $0 < t_1 < \cdots < t_k$ and set $t_0 \equiv 0$. Let β be any pure best response to σ under the convex cost k_ε.

Fix an index $i \in \{1, \ldots, k\}$ and consider the interval $[t_{i-1}, t_i]$. Conditional on the event $\{\sigma \geq t_{i-1}\}$ (i.e., Alice has not exited before t_{i-1}), the date t_i is a *local deadline* in the sense that Alice does not exit on (t_{i-1}, t_i). Holding fixed Bob's continuation play after t_i, we claim that β can be chosen to be weakly increasing on $[t_{i-1}, t_i]$ and strictly increasing wherever it is interior.

Step 1a: benefit depends only on the cumulative manipulation over $[t_{i-1}, t_i]$. Conditional on $\{\sigma \geq t_{i-1}\}$ and Bob's information at t_{i-1}, in state $\omega = -1$ the survival probability to t_i under β is

$$\mathbb{P}(\tau > t_i \mid I^B_{t_{i-1}}, \beta, \omega = -1) = \exp\Big(- \lambda(t_i - t_{i-1}) + \lambda \int_{t_{i-1}}^{t_i} \beta(u)\, du \Big),$$

which depends on β only through $\int_{t_{i-1}}^{t_i} \beta$. In state $\omega = +1$ we have $\mathbb{P}(\tau > t_i \mid \cdot, \omega = +1) = 1$. Therefore, holding $\int_{t_{i-1}}^{t_i} \beta$ fixed leaves (i) the probability of reaching t_i and (ii) Bob's posterior at t_i conditional on no breakdown unchanged, and hence leaves Bob's continuation value from t_i onward unchanged (since continuation play after t_i is held fixed).

Step 1b: exchange argument $\Rightarrow$ increasing intensity. Bob's expected cost on $[t_{i-1}, t_i]$ is incurred only while no breakdown has occurred:

$$\mathbb{E}\left[\int_{t_{i-1}}^{t_i \wedge \tau} k_\varepsilon(\beta(t))\,dt \;\middle|\; I^B_{t_{i-1}}\right] = \int_{t_{i-1}}^{t_i} k_\varepsilon(\beta(t))\,\mathbb{P}(\tau > t \mid I^B_{t_{i-1}}, \beta)\,dt,$$

where the weight $t \mapsto \mathbb{P}(\tau > t \mid I^B_{t_{i-1}}, \beta)$ is weakly decreasing. Take any feasible β on $[t_{i-1}, t_i]$ and suppose there exist $t < t'$ in $[t_{i-1}, t_i]$ and $\delta > 0$ such that $\beta(t) \geq \beta(t') + 2\delta$ and both values are interior (so perturbations remain in $[0, 1]$). Define β^δ by shifting intensity from t to t':

$$\beta^\delta(t) = \beta(t) - \delta, \qquad \beta^\delta(t') = \beta(t') + \delta, \qquad \beta^\delta(u) = \beta(u) \text{ otherwise.}$$

Then $\int_{t_{i-1}}^{t_i} \beta^\delta = \int_{t_{i-1}}^{t_i} \beta$, so by Step 1a the benefit terms are unchanged. But strict convexity of k_ε and the fact that survival weights are weakly larger at t than at t' imply that the expected cost strictly decreases. Hence such an interior decrease cannot occur in a best response. Therefore, on $[t_{i-1}, t_i]$ any pure best response can be chosen *weakly increasing*, and it is *strictly increasing* on the region where it is interior (i.e. wherever $0 < \beta(t) < 1$).

In particular, letting $s_i \equiv \inf\{t \in [t_{i-1}, t_i] : \beta(t) > 0\}$ (with $s_i = t_i$ if the set is empty), we have $\beta(t) = 0$ for $t \in [t_{i-1}, s_i)$ and β is weakly increasing on $[s_i, t_i]$. We will use only this monotonicity (and the implied cap persistence: if β ever reaches 1 on $[t_{i-1}, t_i]$, then it stays at 1 until t_i).

Step 2 (No full blocking and strict decline of Alice's posterior). Let (σ, ϕ) be an equilibrium with $\text{supp}(\sigma) = \{t_1, \ldots, t_k\}$ and $t_1 < \cdots < t_k$. Since ϕ is a best response to σ, ϕ_r is a pure best response to σ for almost every $r \in [0, 1]$.

Step 2a: no blocking. Suppose, toward a contradiction, that there exist $i \in \{1, \ldots, k\}$ and $t \in (t_{i-1}, t_i)$ such that $\mathbb{P}[\phi_r(t) = 1] = 1$. Then for almost every r, ϕ_r is a pure best response and satisfies $\phi_r(t) = 1$. By Step 1 applied to the local deadline interval $[t_{i-1}, t_i]$, each such $\phi_r(\cdot)$ is weakly increasing on $[t_{i-1}, t_i]$ and hence

$$\phi_r(s) = 1 \qquad \forall\, s \in [t, t_i).$$

Thus, in state $\omega = -1$ the breakdown intensity on $[t, t_i)$ is $\lambda(1 - \phi_r(s)) = 0$, so "no breakdown" on $[t, t_i)$ is uninformative and Alice's posterior is constant on that interval. Since sampling costs $c > 0$ per unit time, Alice strictly prefers stopping at t to waiting until t_i, contradicting that σ assigns positive probability to stopping at t_i conditional on reaching t. Therefore, for almost every $t \in [0, t_k)$ we have $\mathbb{P}[\phi_r(t) = 1] < 1$.

Step 2b: $\tilde{p}_t$ strictly decreases along the no-breakdown history. Because $\mathbb{P}[\phi_r(t) = 1] < 1$ for almost every $t < t_k$, in the low state the conditional hazard of a breakdown is strictly positive at almost every time prior to t_k. Hence $\mathbb{P}(\tau \geq t \mid \omega = -1)$ is strictly decreasing in t on $[0, t_k]$. Bayes' rule then implies that Alice's posterior $\tilde{p}_t = \mathbb{P}(\omega = -1 \mid \tau \geq t)$ is strictly decreasing on $[0, t_k]$.

Step 3 (Alice's stopping time is deterministic). Let (σ, ϕ) be an equilibrium. Suppose, toward a contradiction, that σ is not deterministic and has finite support $\text{supp}(\sigma) = \{t_1, \ldots, t_k\}$ with $k \geq 2$ and $t_1 < \cdots < t_k$. Consider the last interval (t_{k-1}, t_k). Since Alice places positive probability on stopping at t_{k-1} and at t_k, she is indifferent between stopping and continuing at t_{k-1} along the no-breakdown history.

For any $t \in (t_{k-1}, t_k)$, Step 1 implies that along the local deadline interval $[t_{k-1}, t_k]$ every pure best response $\phi_r(\cdot)$ is weakly increasing, so (relative to t_{k-1}) Bob's manipulation intensity is weakly higher at t on every realization r. Thus the (subjective) breakdown hazard faced by Alice at time t is weakly lower than at t_{k-1}. Meanwhile, by Step 2, Alice's posterior that the state is low strictly decreases over time on $[t_{k-1}, t_k]$. Both effects strictly reduce the marginal value of additional waiting as t increases, while the marginal sampling cost remains constant at c. Hence, if Alice is indifferent between stopping and continuing at t_{k-1}, she must strictly prefer stopping at any later time in (t_{k-1}, t_k), contradicting that she weakly prefers to continue sampling until t_k on the equilibrium path. Therefore σ must be deterministic.

Step 4 (Given deterministic exit t_1, Bob's best response is delay–and–ramp). Fix a deterministic exit time $t_1 > 0$ for Alice along the no-breakdown history. By Step 0, if Alice exits at t_1 absent a breakdown, she takes action $a = 1$. Let p_t denote Bob's private posterior conditional on no breakdown and his private information. If Bob chooses intensity $b_t \in [0, 1]$, then conditional on no breakdown

$$\dot{p}_t = -\lambda(1 - b_t)p_t(1 - p_t). \tag{28}$$

Step 4a: HJB and best-response map. Define $J_\varepsilon(t, p; t_1)$ as Bob's maximal continuation value from time t to t_1, conditional on no breakdown by t and belief p, when Alice exits at t_1 and takes action 1 absent a breakdown. The terminal condition is

$$J_\varepsilon(t_1, p; t_1) = 1 \qquad \forall\, p \in [0, 1]. \tag{29}$$

A dynamic programming argument yields, for a.e. $t < t_1$,

$$0 = \max_{b\in[0,1]} \Big\{ \partial_t J_\varepsilon(t,p;t_1) - \lambda(1-b)p(1-p)\,\partial_p J_\varepsilon(t,p;t_1) - p\lambda(1-b)\,J_\varepsilon(t,p;t_1) - \gamma b - \varepsilon b^2 \Big\}. \tag{30}$$

For each (t,p) the maximand is strictly concave in b (because $\varepsilon > 0$), hence the maximizer is unique. With $[x]_0^1 \equiv \min\{1, \max\{0,x\}\}$,

$$b_\varepsilon(t,p) = \left[\frac{p\lambda\Big(J_\varepsilon(t,p;t_1) + (1-p)\partial_p J_\varepsilon(t,p;t_1)\Big) - \gamma}{2\varepsilon} \right]_0^1. \tag{31}$$

Step 4b: delay–and–ramp shape. Fix any onset time $s \in [0,t_1]$ and consider strategies with $b(t) = 0$ for $t < s$. On $[s,t_1]$, the exchange argument from Step 1 (now with the single local deadline t_1) implies that any optimal control can be chosen weakly increasing, strictly increasing where interior, and once it reaches the cap 1 it remains at 1 for the remainder of $[s,t_1]$. Denote by $\beta^{\varepsilon,s}$ the resulting (unique a.e.) optimal control on $[s,t_1]$, with $\beta^{\varepsilon,s}(t) = 0$ for $t < s$. This is the *delay–and–ramp* strategy associated with onset s and deadline t_1.

Step 5 (Candidate equilibrium profile). Guided by Steps 0–4, we look for an equilibrium in which:

(i) Along the no-breakdown history, Alice stops at a deterministic time $t_1^\varepsilon > 0$.

(ii) Bob draws once at a deterministic time $s_1^\varepsilon \in [0, t_1^\varepsilon]$ between two pure strategies: the null strategy 0 and the delay–and–ramp strategy $\beta^{\varepsilon,s_1^\varepsilon}$ (relative to deadline t_1^ε), with mixing probability $\alpha^\varepsilon \in [0,1)$.

(iii) Off path (if Alice continues sampling beyond t_1^ε), both pure strategies are extended by setting $\beta(t) \equiv 1$ for all $t > t_1^\varepsilon$, so continuation beyond t_1^ε is uninformative and strictly dominated for Alice.

Steps 6–7 choose $(s_1^\varepsilon, t_1^\varepsilon, \alpha^\varepsilon)$ so that (i)–(iii) indeed form a manipulation equilibrium.

Step 6 (Choosing s_1^ε and t_1^ε to satisfy Bob's incentive conditions). Fix any candidate exit time $t_1 > 0$ and any onset time $s \in [0,t_1)$. Because both pure strategies in the support prescribe no manipulation before s, Bob's private belief at s (conditional on no breakdown) is the common posterior generated under $\beta \equiv 0$; denote it by p_s.

Define Bob's gain from *starting* manipulation at time s (given exit at t_1) as

$$N_\varepsilon(s,t_1) \equiv J_\varepsilon(s,p_s;t_1) - J^{\text{null}}(s,p_s;t_1), \tag{32}$$

$$J^{\text{null}}(s,p_s;t_1) = (1-p_s)\cdot 1 + p_s \cdot e^{-\lambda(t_1-s)} = 1 - p_s\big(1-e^{-\lambda(t_1-s)}\big). \tag{33}$$

Here $J_\varepsilon(s,p_s;t_1)$ is the value of the control problem (30)–(29) starting from (s,p_s).

Step 6a: continuity and a uniform bound. Fix $\bar{t} > t_1^0 + 1$ and restrict attention to $0 \le s \le t_1 \le \bar{t}$. Standard continuity results for finite-horizon control problems imply that $(\varepsilon, s, t_1) \mapsto J_\varepsilon(s,p_s;t_1)$ is continuous on compact subsets, hence so is N_ε.

Moreover, for any fixed (s,t_1) and any feasible $\beta(\cdot) \in [0,1]$, the only difference between $\varepsilon = 0$ and $\varepsilon > 0$ is the additional running cost $\varepsilon\beta(t)^2$, incurred only while no breakdown has occurred. Since $\beta(t)^2 \le 1$, the additional expected cost is at most $\varepsilon(t_1 - s)$. Thus

$$N_0(s,t_1) - \varepsilon(t_1-s) \;\le\; N_\varepsilon(s,t_1) \;\le\; N_0(s,t_1). \tag{34}$$

Step 6b: existence of $(s_1^\varepsilon, t_1^\varepsilon)$ near (s_1^0, t_1^0). In the linear-cost equilibrium, Bob is indifferent at s_1^0, so $N_0(s_1^0,t_1^0) = 0$ and s_1^0 maximizes $s \mapsto N_0(s,t_1^0)$. Therefore, there exist $\eta > 0$ and $\kappa > 0$ such that

$$N_0(s,t_1^0) \le -\kappa \qquad \text{for all } s \in [0, t_1^0 - \eta] \text{ with } |s - s_1^0| \ge \eta. \tag{35}$$

Combining (35) with (34) implies that for all sufficiently small $\varepsilon > 0$ and all t_1 sufficiently close to t_1^0,

$$N_\varepsilon(s,t_1) < 0 \qquad \text{for all } s \in [0, t_1 - \eta] \text{ with } |s - s_1^0| \ge \eta. \tag{36}$$

For each (ε, t_1) close to $(0, t_1^0)$, let $s_\varepsilon(t_1)$ be a maximizer of $s \mapsto N_\varepsilon(s,t_1)$ over the compact set $[s_1^0 - \eta, s_1^0 + \eta] \cap [0, t_1 - \eta]$, and define

$$G_\varepsilon(t_1) \equiv N_\varepsilon\big(s_\varepsilon(t_1), t_1\big). \tag{37}$$

By Berge's maximum theorem, G_ε is continuous in (ε, t_1) and $G_0(t_1^0) = 0$.

In the linear-cost model,

$$N_0(s,t_1) = p_s\big(1 - e^{-\lambda(t_1-s)}\big) - \gamma(t_1 - s), \tag{38}$$

so if $\Delta^0 \equiv t_1^0 - s_1^0 > 0$ then $N_0(s_1^0,t_1^0) = 0$ implies $p_{s_1^0}(1 - e^{-\lambda\Delta^0}) = \gamma\Delta^0$ and therefore

$$G_0'(t_1^0) = \partial_{t_1} N_0(s_1^0, t_1^0) = p_{s_1^0}\lambda e^{-\lambda\Delta^0} - \gamma = \gamma\Big(\frac{\lambda\Delta^0}{e^{\lambda\Delta^0} - 1} - 1\Big) < 0. \tag{39}$$

Hence G_0 crosses 0 transversely at t_1^0: there exists $\delta > 0$ such that $G_0(t_1^0 - \delta) > 0$ and $G_0(t_1^0 + \delta) < 0$. By continuity, the same sign change holds for G_ε for all sufficiently small $\varepsilon > 0$, so by the intermediate value theorem there exists $t_1^\varepsilon \in (t_1^0 - \delta, t_1^0 + \delta)$ such that $G_\varepsilon(t_1^\varepsilon) = 0$. Define $s_1^\varepsilon \equiv s_\varepsilon(t_1^\varepsilon)$. Then

$$N_\varepsilon(s_1^\varepsilon, t_1^\varepsilon) = 0 \qquad \text{and} \qquad N_\varepsilon(s, t_1^\varepsilon) \leq 0 \;\; \text{for all } s \in [0, t_1^\varepsilon - \eta]. \tag{40}$$

Step 7 (Choosing α^ε and completing the equilibrium). Fix $(s_1^\varepsilon, t_1^\varepsilon)$ from Step 6. Let β^ε denote the delay–and–ramp best response $\beta^{\varepsilon, s_1^\varepsilon}$ relative to deadline t_1^ε (Step 4b), and consider the mixed strategy $\phi^{\varepsilon,\alpha}$ that assigns probability α to β^ε and probability $1 - \alpha$ to the null strategy 0, with the draw realized at time s_1^ε.

Step 7a: Alice's continuation constraint. For $t < t_1^\varepsilon$, define

$$q_{\varepsilon,\alpha}(t) \equiv \mathbb{P}^{\varepsilon,\alpha}(\omega = -1,\ \phi_r(t) = 0 \mid \tau > t), \qquad \tilde{p}_t^{\varepsilon,\alpha} \equiv \mathbb{P}^{\varepsilon,\alpha}(\omega = -1 \mid \tau > t).$$

A one-step deviation calculation (as in the linear-cost case) yields that Alice weakly prefers continuing at time t iff

$$c \;\leq\; \lambda\, q_{\varepsilon,\alpha}(t)\, \tilde{p}_t^{\varepsilon,\alpha} - \dot{\tilde{p}}_t^{\varepsilon,\alpha}. \tag{41}$$

Because β^ε is backloaded (weakly increasing on $[s_1^\varepsilon, t_1^\varepsilon]$ and equal to 0 before s_1^ε), the right-hand side of (41) is weakly decreasing in t along the no-breakdown history. Hence it suffices to impose (41) at $t \uparrow t_1^\varepsilon$. Define

$$\bar{\alpha}^\varepsilon \equiv \sup\Big\{\alpha \in [0,1) :\; c \leq \liminf_{t \uparrow t_1^\varepsilon} \big(\lambda q_{\varepsilon,\alpha}(t) \tilde{p}_t^{\varepsilon,\alpha} - \dot{\tilde{p}}_t^{\varepsilon,\alpha}\big)\Big\}. \tag{42}$$

Then any $\alpha \leq \bar{\alpha}^\varepsilon$ guarantees that Alice weakly prefers continuing for all $t < t_1^\varepsilon$.

Step 7b: Alice's initiation constraint. Let $U_A^\varepsilon(\alpha)$ be Alice's ex ante payoff at $t = 0$ from the rule "sample until t_1^ε or a breakdown" against $\phi^{\varepsilon,\alpha}$, and let $U_A^{\text{stop}}(p_0) \equiv \max\{p_0, 1 - p_0\}$ be the payoff from stopping immediately at $t = 0$. The map $\alpha \mapsto U_A^\varepsilon(\alpha)$ is weakly decreasing, and $(\varepsilon, \alpha) \mapsto U_A^\varepsilon(\alpha)$ is continuous on compact subsets of $[0, \infty) \times [0, 1)$. Define the initiation cutoff

$$\alpha_{\text{start}}^\varepsilon \equiv \sup\big\{\alpha \in [0,1) :\; U_A^\varepsilon(\alpha) \geq U_A^{\text{stop}}(p_0)\big\}. \tag{43}$$

Then Alice weakly prefers initiating sampling at $t = 0$ whenever $\alpha \leq \alpha_{\text{start}}^\varepsilon$. Moreover, $\alpha_{\text{start}}^\varepsilon \to \alpha^0$ as $\varepsilon \downarrow 0$ (by monotonicity and continuity, as in the binding/slack discussion in Step 5 of the revised proof).

Step 7c: choice of α^ε and equilibrium verification. Choose

$$\alpha^\varepsilon \equiv \min\{\alpha^0,\ \bar{\alpha}^\varepsilon,\ \alpha^\varepsilon_{\text{start}}\}. \tag{44}$$

By assumption, the linear-cost equilibrium satisfies $\alpha^0 > 0$ and $s_1^0 < t_1^0$. Then by continuity and by the fact that Alice's sampling payoff is strictly decreasing in α, there exists $\underline{\alpha} \in (0, \alpha^0)$ and $\bar{\varepsilon} > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon})$,

$$\bar{\alpha}^\varepsilon \geq \underline{\alpha}, \qquad \alpha^\varepsilon_{\text{start}} \geq \underline{\alpha}, \qquad \text{and hence} \qquad \alpha^\varepsilon = \min\{\alpha^0, \bar{\alpha}^\varepsilon, \alpha^\varepsilon_{\text{start}}\} \geq \underline{\alpha} > 0.$$

Therefore, Alice initiates sampling (since $\alpha^\varepsilon \leq \alpha^\varepsilon_{\text{start}}$) and continues until t_1^ε (since $\alpha^\varepsilon \leq \bar{\alpha}^\varepsilon$). Extend both pure strategies off path by setting $\beta(t) \equiv 1$ for all $t > t_1^\varepsilon$, so Alice strictly prefers stopping at t_1^ε. Bob is indifferent between the null strategy and β^ε at s_1^ε and has no profitable onset deviation by (40); conditional on manipulating, β^ε is optimal by construction (Step 4). Therefore $(\sigma^\varepsilon, \phi^\varepsilon)$ is a manipulation equilibrium with on-path play $(s_1^\varepsilon, \alpha^\varepsilon, t_1^\varepsilon)$. □

O5 (Non)Existence of Wald Equilibrium

The following proposition provides conditions for (non)existence of the Wald equilibrium.

Proposition 7 (Existence of Wald equilibrium)**.** *Suppose $c < \lambda/2$.*

1. *If $\gamma < c$, the Wald equilibrium with $t_1^w > 0$ does not exist for any $p \in (\underline{p}^w, \overline{p}^w)$.*
2. *There exists $\hat{\gamma} \in (c, \gamma^*]$ such that the Wald equilibrium with $t_1^w > 0$ exists for all $p \in (\underline{p}^w, \overline{p}^w)$ if and only if $\gamma \geq \hat{\gamma}$.*
3. *If $\gamma > c$, the Wald equilibrium with $t_1^w > 0$ exists for priors $p \in (\underline{p}^w, \gamma/\lambda]$.*

Proof. Suppose that Alice chooses the Wald stopping time

$$t_1^w = \frac{1}{\lambda} \log\left(\frac{\lambda - c}{c} \cdot \frac{p}{1-p}\right),$$

which is obtained through Bayes' updating and the stopping belief c/λ. Alice will (continue to) sample at belief p (i.e., $t_1^w(p) > 0$) only if Bob's gain from manipulation satisfies $N(\hat{p}; \gamma) \leq 0$ for all $\hat{p} \in (\underline{p}^w, p]$. Note that

$$N(p;\gamma) = p(1 - e^{-\lambda t_1^w}) - \gamma t_1^w = \frac{\lambda p - c}{\lambda - c} - \frac{\gamma}{\lambda} \log\left(\frac{\lambda - c}{c} \cdot \frac{p}{1-p}\right).$$

Its derivative is given by

$$\frac{\partial N(p;\gamma)}{\partial p} = \frac{\lambda}{\lambda - c} - \frac{\gamma}{\lambda}\left(\frac{1}{p} + \frac{1}{1-p}\right)$$

with

$$N(\underline{p}^w;\gamma) = 0 \quad \text{and} \quad \frac{\partial N(\underline{p}^w;\gamma)}{\partial p} = \frac{\lambda(c-\gamma)}{c(\lambda - c)}.$$

Consequently, if $\gamma < c$, $N(p;\gamma) > 0$ for all p sufficiently close to $\underline{p}^w$. Therefore, as the belief drops sufficiently close to $\underline{p}^w$, Bob will start manipulation at the full rate, and the standard Wald equilibrium with $t_1^w > 0$ cannot exist.

Next suppose $\gamma \geq c$. Note that

$$\frac{\partial N(p;\gamma)}{\partial p} \leq 0 \iff \gamma \geq \frac{\lambda^2 p(1-p)}{\lambda - c}.$$

Since $N(\underline{p}^w;\gamma) = 0$, if $\gamma \geq \gamma^*$ with

$$\gamma^* \equiv \frac{\lambda^2}{4(\lambda - c)},$$

then $N(p;\gamma) \leq 0$ for all $p \in (\underline{p}^w, \overline{p}^w)$. Now define

$$S(\gamma) \equiv \sup_{p \in (\underline{p}^w, \overline{p}^w)} N(p;\gamma).$$

Because $N(p;\cdot)$ is continuous and strictly decreasing in γ pointwise, S is continuous and strictly decreasing in γ on $[c,\infty)$. Moreover, we have previously shown that

$$S(c) > 0 \quad \text{and} \quad S(\gamma^*) \leq 0.$$

Hence there exists a unique cutoff $\hat{\gamma} \in (c, \gamma^*]$ such that

$$S(\gamma) \leq 0 \iff \gamma \geq \hat{\gamma}.$$

Therefore, the standard Wald equilibrium exists for all $p \in (\underline{p}^w, \overline{p}^w)$ if and only if $\gamma \geq \hat{\gamma}$.

Finally, suppose that $\gamma \geq c$ and $p \in (\underline{p}^w, \gamma/\lambda]$. For any positive stopping time t_1, as illustrated in Figure 2, we have $p(1 - e^{-\lambda t_1}) < \gamma t_1$. That is, Bob has a strict incentive not to manipulate at time 0. In the absence of manipulation, Bob's belief p_t is strictly decreasing over time. Therefore, if Bob does not want to manipulate at time 0, he does not want to manipulate at any later time $t \in (0, t_1)$ either. That is, Bob's strict best response to any $t_1 > 0$ is not to manipulate at any time $t \in [0, t_1)$. Consequently, it is an equilibrium for Alice to adopt her optimal sampling strategy as in the standard Wald model as if Bob were absent. □

Remark 1 (On the assumption $c < \gamma < \lambda/2$). Throughout the paper, we impose $c < \gamma < \lambda/2$. This restriction isolates the parameter region in which both Wald and manipulation equilibria may arise.

If $\gamma < c$, then the Wald equilibrium does not exist by Proposition 7. Since Theorem 4 shows that a Wald equilibrium exists whenever a manipulation equilibrium exists, it follows that no manipulation equilibrium exists either.

If $\gamma > \lambda/2$, then $\gamma/\lambda > 1/2$. In any manipulation equilibrium with $s_1 = 0$, Bob's indifference condition implies that Alice's stopping belief must equal γ/λ, which therefore exceeds 1/2. This is impossible, because by Lemma 5 Alice never stops at a belief above 1/2. Hence no manipulation equilibrium exists in this case either. □

O6 Numerical Examples

We conclude with three numerical examples. Example 1 compares static and sequential sampling under two priors and shows that static sampling can dominate when the prior is low, even under simultaneous timing. Example 2 illustrates coexistence of Wald and manipulation equilibria and shows that the Wald equilibrium can fail on part of the prior range where the classical Wald problem still prescribes positive sampling. Example 3 shows that, even at the critical prior p^*, the worst manipulation equilibrium may exist while static sampling remains strictly better for Alice.

Example 1 (Static versus sequential sampling)**.** Fix $(\lambda, \gamma, c) = (1, 1/5, 1/20)$, so $p^* \approx 0.754$. If $p_0 = 3/4$, then $p_0 < p^*$. The worst sequential manipulation equilibrium has on-path play $(0, \overline{\alpha}, t_1)$ with $t_1 \approx 3.653$, $\overline{\alpha} \approx 0.107$, and Alice's payoff $V(p_0) \approx 0.809$. In static sampling, the deterrence threshold is $\widehat{\gamma} \approx 0.258 > \gamma$, so Alice chooses the deterrent sample size $t_a = t_1$ and obtains payoff ≈ 0.798. Thus sequential sampling is still better at this higher prior.

If instead $p_0 = 1/2$, then the worst sequential manipulation equilibrium has $t_1 \approx 2.232$, $\overline{\alpha} \approx 0.342$, and $V(p_0) \approx 0.704$. In static sampling, $\widehat{\gamma} \approx 0.195 < \gamma$, so Alice can choose the benchmark static size $t_a = t^s \approx 2.303$ and obtains payoff ≈ 0.835. Thus static sampling strictly dominates sequential sampling, reversing the classical ranking.

The same conclusion survives under simultaneous timing. If Alice chooses t_a while Bob simultaneously randomizes between $t_b = 0$ and $t_b = t_a$, then her static payoff at $p_0 = 1/2$ is 0.804, still above the sequential payoff 0.704. Hence lower priors can make sequential sampling more vulnerable to manipulation, whereas higher priors preserve its usual advantage. □

Example 2 (Coexistence and failure of the Wald equilibrium)**.** Fix $(\lambda, c, \gamma) = (1, 0.25, 0.28)$, so $\underline{p} = \gamma/\lambda = 0.28$ and $\underline{p}^w = c/\lambda = 0.25$. In the standard Wald problem (without Bob), Alice's stopping time is

$$t_1^w(p) = \ln\left(3\frac{p}{1-p}\right),$$

and the upper cutoff $\overline{p}^w$ solving $V^w(\overline{p}^w) = \overline{p}^w$ is

$$\overline{p}^w \approx 0.5996.$$

A Wald equilibrium with Bob requires Bob not to gain from full manipulation over $[0, t_1^w(p)]$, namely

$$N^w(p;\gamma) = \frac{p - 0.25}{0.75} - 0.28\ln\left(3\frac{p}{1-p}\right) \leq 0.$$

This function is decreasing near 0.25 and increasing on $(0.3, 0.7)$, so the largest prior satisfying the inequality is the nontrivial root

$$p^\dagger \approx 0.3587.$$

Hence the Wald equilibrium exists iff $p \in (\underline{p}^w, p^\dagger]$, even though the classical Wald problem prescribes positive sampling up to $\overline{p}^w \approx 0.5996$.

Manipulation equilibria are necessarily non-passive in this example. Indeed, if

$$\Gamma(t) = 1 - e^{-t} - t(1 - 0.28\,t),$$

then the unique positive solution to $\Gamma(t) = 0$ yields $p^* \approx 0.660 > \overline{p}^w$. Thus any candidate manipulation equilibrium satisfies

$$p = \frac{0.28\,t_1}{1 - e^{-t_1}}, \qquad \overline{\alpha}(p) = 1 - \frac{1}{p}\frac{0.25}{0.75e^{-t_1} + 0.25}.$$

Along this branch, Alice's gain from sampling relative to immediate stopping is

$$F(t_1) = \frac{0.25 \cdot 0.75\big(1 - e^{-t_1} - t_1 e^{-t_1}\big)}{0.75e^{-t_1} + 0.25} > 0 \qquad \text{for all } t_1 > 0.$$

So value matching never pins down the upper existence cutoff. Instead, the cutoff is determined by coexistence with the Wald equilibrium: for each $p \in (\underline{p}, p^\dagger)$, we have

$t_1 > 0$, $F(t_1) > 0$, and $\Gamma(t_1) < 0$, so a manipulation equilibrium exists; for $p > p^\dagger$ the Wald equilibrium fails, and since a Wald equilibrium must exist whenever a manipulation equilibrium does, no manipulation equilibrium can exist there. Consequently,

$$\overline{p} = p^\dagger \approx 0.3587 < \frac{1}{2},$$

the manipulation-equilibrium existence set is $(\underline{p}, \overline{p})$, and the endpoint is not attained because $\overline{\alpha}(\overline{p}) = 0$. Thus Wald and manipulation equilibria coexist on $(0.28, 0.3587)$, while for $p \in (0.3587, 0.5996]$ the classical Wald problem still calls for sampling but no Wald equilibrium exists once Bob is present. □

Example 3 (Static dominates at the critical prior p^*)**.** Fix $(\lambda, c, \gamma) = (1, 0.085, 0.24)$. The equation $\Gamma(t) = 1 - e^{-t} - t(1 - 0.24\, t) = 0$ has a unique positive solution $t_1^* \approx 2.7472$, which implies

$$p^* = \frac{0.24\, t_1^*}{1 - e^{-t_1^*}} \approx 0.7045.$$

Set $p_0 = p^*$. The worst sequential manipulation equilibrium has on-path play $(0, \overline{\alpha}(p^*), t_1^*)$, where

$$\overline{\alpha}(p^*) = 1 - \frac{1}{p^*} \frac{0.085}{0.915 e^{-t_1^*} + 0.085} \approx 0.1601.$$

Alice's equilibrium payoff is

$$\hat{V}_0\big(0, \overline{\alpha}(p^*), t_1^*\big) \approx 0.7068.$$

Since $p^* > 1/2$ and $0.7068 > p^* \approx 0.7045$, Alice strictly prefers to start sampling, so the worst manipulation equilibrium indeed exists at $p_0 = p^*$.

In static sampling, the benchmark size and deterrence threshold are

$$t^s = \log\left(\frac{p^*}{0.085}\right) \approx 2.1148, \qquad \widehat{\gamma} = \frac{p^* - 0.085}{\log(p^*/0.085)} \approx 0.2929.$$

Because $\gamma = 0.24 < \widehat{\gamma}$, the static equilibrium uses the deterrent sample size $t_a = t_1^*$ and Bob does not manipulate. Alice's static payoff is

$$U^{\text{stat}}(p^*) = (1 - p^*) + p^*(1 - e^{-t_1^*}) - 0.085\, t_1^* = (1 - p^*) + (0.24 - 0.085)t_1^* \approx 0.7213.$$

Therefore

$$U^{\text{stat}}(p^*) \approx 0.7213 > \hat{V}_0\big(0, \overline{\alpha}(p^*), t_1^*\big) \approx 0.7068.$$

Hence, even at the critical prior p^*, the worst sequential manipulation equilibrium exists but static sampling is strictly better for Alice. □