Dynamic Disclosure with(out) Timestamps*

Aaron Kolb† Beixi Zhou‡

August 18, 2026

Click here for latest version

Abstract

We study the role of timestamps in a dynamic disclosure game. At a random date, a sender privately obtains one piece of hard evidence about a hidden, evolving binary state and can disclose it at any later date. When evidence carries a timestamp, the unique equilibrium features immediate disclosure of good evidence and disclosure of bad evidence after a deterministic, timestamp-dependent delay. We also show that the principle of worst case beliefs conditional on nondisclosure holds in this timestamps setting. Without timestamps, if the prior is high or the sender is patient, equilibrium must feature delayed disclosure of good evidence, and we construct an equilibrium in which good evidence is initially delayed, then stochastically released, and eventually disclosed immediately. With an impatient sender, we construct an equilibrium in which bad evidence is disclosed with delay in isolated bursts. Timestamps prevent pretending old good evidence is fresh and allow proving bad evidence is old, thereby accelerating disclosure of good evidence and facilitating disclosure of bad evidence.

*We thank Brendan Daley, Bart Lipman, Elliot Lipnowski, George Mailath, Curtis Taylor, and seminar audiences at Boston University, Duke University, Indiana University, 2026 AEA CSWEP Session on Economic Theory, Bonn Christmas Microeconomic Theory Conference, University of Georgia Microeconomic Theory Mini-Conference, IIOC 2026, Pennsylvania Economics Theory Conference, and University of Miami Theory Day for helpful comments and conversations.

†Kelley School of Business, Indiana University, kolba@iu.edu.

‡Department of Economics, University of Pittsburgh, beixi.zhou@pitt.edu.

1 Introduction

A key feature of many evidentiary documents is whether they carry a timestamp. Evidence may therefore convey not only information about the underlying state, but also information about when that information was generated. Medical test results, financial filings such as bankruptcies and audit reports, and legal records like signed contracts and digital forensic logs, all bear verifiable dates. By contrast, many commonly used documents are not timestamped. For example, undated recommendation letters provide an evaluation without indicating when it was written; sensationalist news media sometimes reuse old photos to support current narratives when original dates are difficult or costly for the public to verify.

There has been growing attention in the literature to how the timing of disclosure affects beliefs and incentives in dynamic environments. However, existing work typically abstracts from whether the act of disclosure itself reveals the time at which the evidence was generated. It remains an open question whether the presence or absence of timestamps plays a role in strategic disclosure over time.

We address this question by introducing and studying a simple continuous-time sender-receiver model of dynamic disclosure of hard evidence about an evolving state. Specifically, there is a hidden state θ_t that evolves according to Markov switching dynamics between two states, 1 ("good") and 0 ("bad"). A sender privately obtains one piece of hard evidence at some random, exponentially distributed date τ and can choose to disclose it (or not) at any later time. We consider two distinct environments: *timestamps* and *no timestamps*. In the timestamps environment, upon disclosure, the receiver sees both θ_τ and τ; in the no-timestamps environment, the receiver only sees θ_τ and must infer τ itself. The sender's flow payoff is the receiver's posterior belief that the state is 1 discounted at rate r.

This simple environment generates rich equilibrium dynamics that differ qualitatively between the timestamps and no-timestamps environments.

In the timestamps environment, we show that there is a unique equilibrium, and good evidence is disclosed immediately while bad evidence is disclosed after a deterministic, timestamp-dependent delay. To understand the intuition, consider first the incentive for disclosing good evidence. If the sender possesses good evidence $\theta_\tau = 1$ at τ and discloses it immediately, the receiver's belief jumps to 1 and decays according to the hidden Markov switching dynamics thereafter. If instead the sender were to delay until some time $t > \tau$, the receiver would retroactively update about θ_τ, so her belief path after t would be exactly the same as under immediate disclosure. The only effect of delay is to lower the receiver's belief during the interval $[\tau, t)$, while evidence is being withheld. Timestamps prevent the sender from pretending good evidence is fresh and thus the sender cannot benefit from delaying disclosure of good evidence.

In addition to proving that good evidence is fresh, timestamps allow the sender to prove that old bad evidence is in fact old. Therefore, when the sender discloses old bad evidence, beliefs update retroactively downward to 0, but drift up according to Markov switching after that. Furthermore, since evidence only arrives once, once bad evidence has been disclosed, the receiver no longer expects good evidence to be disclosed. This shuts down a "no news is bad news" effect which raises the trajectory of future beliefs. For each bad evidence arrival date, there is a unique date at which the receiver's belief after disclosure of that evidence crosses the no-disclosure belief path from below. This is the optimal time to disclose bad evidence. Since the sender discloses at the first instant it is myopically optimal to do so, the equilibrium is independent of the sender's discount rate. Additionally, we show that the Minimum Principle (Acharya, DeMarzo and Kremer, 2011; Guttman, Kremer and Skrzypacz, 2014) applies in our setting: at all times, the receiver's equilibrium belief conditional on no disclosure is the most pessimistic belief consistent with Bayes' rule among all possible disclosure strategies.

By contrast, in the no-timestamps environment, we show that unless the prior belief is sufficiently low, equilibrium must feature delayed disclosure of some good evidence. Toward a contradiction, suppose the receiver conjectured immediate disclosure of good evidence at all times, and the sender happens to obtain good evidence at time 0. By disclosing it immediately, the receiver's belief jumps up to 1 and begins decaying due to the Markov switching. By instead disclosing after a small delay, the receiver's belief is lower in the interim, but, importantly, it still jumps to 1 after disclosure since the receiver does not observe the timestamp and assumes the evidence is fresh, and then it begins decaying. Thus, delay involves a trade-off between foregoing a higher belief in the short run and inducing a slightly higher belief at all times in the long run. When the prior belief is above a cutoff, this trade-off favors delay.[1]

Under some conditions, there exists an equilibrium with no disclosure of bad evidence. A sufficient condition is that the arrival rate of evidence is sufficiently small. This condition holds for a larger range of parameter values when the sender is very impatient. Intuitively, when the sender is very impatient, she places a large weight on the immediate negative effect of bad evidence disclosure and has a stronger incentive to delay.[2]

[1] In standard disclosure models, equilibria often feature cutoff strategies, with highest types disclosing immediately. The delayed disclosure in our model does not contradict this theme, because the notion of "type" applies not only to whether evidence is good or bad but also its timestamp, and from this perspective, there is no highest type — since for any arrival time of good evidence, there are always future arrival times for the sender to mimic.

[2] In contrast, impatience increases the sender's incentive to disclose good evidence early in the no-timestamps case. With timestamps, recall that the unique equilibrium is independent of the discount rate.

To cleanly characterize delayed disclosure of good evidence, we specialize to the case where the bad state is absorbing, which is sufficient for ensuring bad evidence is never disclosed. When the prior is high, or the sender is sufficiently patient, we construct an equilibrium consisting of three distinct phases, which we label *stockpiling*, *purging*, and *transparency*. In the stockpiling phase, the sender does not disclose good evidence, creating a stockpile of (possible arrival dates of) undisclosed good evidence from the receiver's perspective. At some date, the sender would prefer to disclose good evidence. However, equilibrium cannot have a bang-bang disclosure strategy with respect to good evidence, with an abrupt transition to a phase of immediate disclosure, because at the first instant of disclosure, the receiver would be expecting all old good evidence to be disclosed, and therefore would not revise her belief to 1. The sender could simply wait an instant longer and separate from these older evidence types. This tension — where a conjecture of too much disclosure incentivizes delay — is naturally resolved with mixing in the purging phase. During this phase, the sender stochastically discloses old evidence at an intensity that depends on calendar time but not on the evidence's timestamp. Over time, the disclosure intensity increases, which causes the stockpile to deplete, shifting the distribution of evidence to more recent dates. This means that disclosure is eventually interpreted more favorably, which sustains the sender's indifference to delay in the first place. By the end of this phase, the disclosure intensity explodes and the stockpile fully (but continuously) empties. At that point, the receiver's belief is sufficiently low that the sender is willing to disclose any good evidence immediately.

The sender's incentive to delay in the purging phase depends not only on the receiver's belief about the current state, but also on the receiver's beliefs about the sender's evidence, since the latter belief determines the receiver's belief after disclosure and the evolution of the receiver's belief conditional on no disclosure. Technically, this presents a challenge in that the receiver's current belief about the state is not the only state variable. Moreover, the receiver must form a belief distribution at each instant about the continuum of past evidence arrival dates. However, with timestamp-independent disclosure rates, we show that these beliefs can be summarized by a few stock variables that track the joint probability that the sender has good or bad evidence and the state is good, conditional on no disclosure. We describe the rest of the construction in detail in Section 4.2.

Finally, we analyze disclosure of bad evidence, focusing on an impatient sender, which sustains immediate disclosure of good evidence. We construct an equilibrium in which good evidence is disclosed immediately, while bad evidence is disclosed with delay in isolated bursts. Intuitively, a transparency phase cannot exist for bad evidence, since all disclosures would be interpreted as fresh. The delays between bursts ensure that disclosed evidence has sufficiently aged on average that disclosure is optimal.

Our analysis points to a natural policy implication. Without timestamps, there can be delayed disclosure of good evidence, so institutions that require timestamps can help the receiver get information sooner. Moreover, timestamps can facilitate the disclosure of bad evidence by eliminating equilibria in which no bad evidence is ever disclosed and ensuring that all bad evidence is eventually disclosed.

Related literature

Our paper contributes to a large literature on strategic disclosure of hard information by a sender without commitment, pioneered by Grossman and Hart (1980), Grossman (1981), and Milgrom (1981). While the canonical static models predict full unraveling, with essentially all sender types disclosing to separate themselves from lower types, Dye (1985) showed that when there is uncertainty about whether the sender has evidence to disclose, unraveling need not be complete, as low types with evidence can refrain from disclosing it to mimic high types that do not have evidence. A long literature on mostly static models has explored various alternative explanations for nondisclosure.

A small and recent literature has studied disclosure in dynamic settings. In Acharya et al. (2011), a sender chooses when to disclose information about a fixed state of the world when there is exogenous public news about that state. The authors show that exogenous news creates a real options problem for the sender: by waiting to disclose, the sender retains the option to not disclose in case very favorable news arrives.

In Guttman et al. (2014), the sender can obtain multiple pieces of evidence sequentially across two stages, but the state (firm value) is constant and independent of the evidence arrival times. The authors show that late disclosures are interpreted more favorably. Antić and Pei (2026) study selective disclosure in an overlapping generations model with multiple, noisy pieces of evidence. In their model, evidence does not carry a timestamp. Although the state is fixed, the absence of timestamps is relevant in that disclosed signals convey information about the history of evidence arrival and what signals could have been concealed. Our focus is different; evidence arrival time and the state are intimately linked due to Markov switching. Specifically, more recent good evidence is more indicative of the state being good currently, and the opposite is true of bad evidence. This allows us to study optimal disclosure timing even when there is only one piece of evidence.

Kremer, Schreiber and Skrzypacz (2024) study disclosure of evidence about an evolving state. In their model, evidence is effectively timestamped, but the sender must either disclose it immediately or never; delayed disclosure is not permitted, consistent with Securities and Exchange Commission rules. In our model, delayed disclosure is permitted and is central to the analysis: with timestamps, bad evidence is optimally disclosed only after a delay, while without timestamps, good evidence may be delayed in order to appear fresh. Thus, relative to Kremer et al. (2024), our paper isolates how the ability to delay disclosure interacts with the presence or absence of timestamps. A further difference relates to the threshold for disclosing evidence. In Kremer et al. (2024), this threshold lies below the current mean belief of the receiver. In our model, with timestamps, bad evidence is disclosed only once it is old enough that disclosure doesn't lower the receiver's belief about the current state. Furthermore, without timestamps, good evidence is sometimes withheld even though disclosure would immediately raise the receiver's belief.

Other papers study dynamic persuasion with commitment; e.g., Ely (2017) and Ely, Georgiadis and Rayo (2025). In the latter paper, feedback is given to the sender (receiver) according to a bang-bang rule: first silence, then transparency. Although the results are not directly comparable due to fundamental model differences, it is notable that in our no-timestamps environment, the bang-bang policy cannot occur in equilibrium, and instead there is a purging phase between our stockpiling (silence) and transparency phases.

2 Model

The game is played between a sender and a receiver in continuous time over an infinite horizon. There is a hidden, binary underlying state $\theta_t \in \{0, 1\}$ that evolves according to Markov switching; when the state is i, it switches to $j \neq i$ at constant exponential rate λ_i. We assume $\lambda_1 > 0$ and $\lambda_0 \geq 0$. Neither the sender nor the receiver observes the realization of the state process, and they share a common prior over the initial state $p_0 = \Pr(\theta_0 = 1) \in (0, 1)$.

At some random date τ, the sender privately obtains hard evidence of the current state θ_τ. The receiver does not directly observe τ or θ_τ. We assume that τ is distributed exponentially with parameter μ and is independent of the state process $(\theta_t)_{t\geq 0}$. The sender can receive evidence at most once. The sender then chooses to disclose this evidence at any date $s \geq \tau$ or never disclose it, which we denote by $s = \infty$. We consider two environments: *timestamps* and *no timestamps*. In the timestamps environment, when evidence is disclosed, the receiver observes both the realization θ_τ and the timestamp τ. In the no-timestamps environment, the receiver only observes the realization θ_τ (and must infer τ).

A strategy for the sender consists of a collection of cumulative distribution functions $H(\cdot; \theta_\tau, \tau, t)$ on $[t, \infty]$, for $\theta_\tau \in \{0, 1\}$ and $0 \leq \tau \leq t$, where $H(s; \theta_\tau, \tau, t)$ specifies the probability that a sender who obtained evidence θ_τ at time τ and has not disclosed it before time t plans to disclose by time s. Strategies must satisfy the usual internal consistency requirement. We say that a strategy is *age-independent* if for all τ, τ', t with $0 \leq \tau, \tau' \leq t$, and $x \in \{0,1\}$ $H(\cdot; x, \tau, t) = H(\cdot; x, \tau', t)$. In other words, the continuation strategy is independent of the evidence's age conditional on the current time. Much of our analysis in the no-timestamps environment focuses on age-independent equilibria.

A belief system for the receiver specifies at each time t, given the public history, a joint distribution over the current state θ_t, and the sender's evidence history E_t, where $E_t = \emptyset$ denotes no evidence has arrived and $E_t = (\tau, x) \in [0,t] \times \{0,1\}$ if evidence arrived at time τ and has type $x = \theta_\tau$. We denote by p_t the receiver's posterior belief at time t that $\theta_t = 1$; the process $p = (p_t)_{t \geq 0}$ is $[0,1]$-valued and càdlàg. The sender obtains a flow payoff equal to p_t and discounts the future at constant rate $r > 0$.

A perfect Bayesian equilibrium (henceforth, equilibrium) is a strategy and belief system such that the sender's strategy is sequentially rational and the belief process is consistent with Bayes' rule whenever possible. Since evidence arrives stochastically, nondisclosure is always on path, and Bayes' rule applies. Disclosure can be off-path. In the timestamps case, the receiver's belief after off-path disclosure is already determined by the timestamp. In the no-timestamps case, there is flexibility in off-path beliefs. We say that off-path beliefs satisfy the *support restriction* if at each time t, and for each type of evidence, they place zero probability on timestamps for which disclosure before time t is expected with probability 1.

Belief evolution. Let $p^* := \lambda_0/(\lambda_0 + \lambda_1)$ denote the long-run steady-state belief. Given the continuous-time Markov chain, two belief processes will be useful throughout. First,

$$\phi_t := p^* + (p_0 - p^*)e^{-(\lambda_0+\lambda_1)t}$$

is the ex ante probability that $\theta_t = 1$, which converges monotonically to p^* from above (below) when $p_0 > p^*$ ($p_0 < p^*$). Second, adopting the convention that G ("good") refers to state $\theta = 1$ and B ("bad") to $\theta = 0$, define $q_t^{G,s} := \Pr(\theta_t = 1 \mid \theta_s = 1)$ and $q_t^{B,s} := \Pr(\theta_t = 1 \mid \theta_s = 0)$ as the probabilities that $\theta_t = 1$ conditional on the state at $s \leq t$ being good or bad, respectively. Then for all $t \geq s \geq 0$, both beliefs evolve according to

$$\dot{q}_t^{i,s} = \lambda_0(1 - q_t^{i,s}) - \lambda_1 q_t^{i,s}$$

with initial conditions $q_s^{G,s} = 1$ and $q_s^{B,s} = 0$. As $t \to \infty$, $q_t^{G,s}$ converges to p^* from above and $q_t^{B,s}$ converges to p^* from below.

3 Timestamps

3.1 Equilibrium characterization

In the timestamps environment, there exists a unique equilibrium. Conditional on evidence arriving, every realization, good or bad, is eventually disclosed. In particular, good evidence immediately and bad evidence after a finite, timestamp-dependent delay. If the sender receives good evidence $\theta_s = 1$ at time s, delaying disclosure cannot make the evidence appear more recent and therefore cannot improve the receiver's post-disclosure belief. If the sender receives bad evidence $\theta_s = 0$ at time s, he initially withholds it and discloses it at time $s + a^*(s)$. Bad evidence becomes less damaging for beliefs about the current state as it ages, the sender can therefore potentially benefit from disclosing bad evidence once it is sufficiently old.

Theorem 1. *There exists a unique equilibrium. If the sender ever possesses good evidence (on or off path), he discloses it immediately. If he possesses bad evidence at time s, he discloses it at time $s + a^*(s)$, where $a^*(s) \in [0, \infty)$ for all $s \geq 0$ is the unique positive solution to*

$$q^{B,s}_{s+a^*(s)} = p^{\dagger}_{s+a^*(s),s},$$

where

$$q^{B,s}_t = \frac{\lambda_0}{\lambda_0 + \lambda_1}\left(1 - e^{-(\lambda_0+\lambda_1)(t-s)}\right) \quad \textit{and} \quad p^{\dagger}_{t,s} = \frac{e^{-\mu t}\phi_t + \int_s^t \mu e^{-\mu\tau}(1-\phi_\tau)q^{B,\tau}_t d\tau}{e^{-\mu t} + \int_s^t \mu e^{-\mu\tau}(1-\phi_\tau) d\tau}.$$

If the sender still possesses this evidence after $s + a^*(s)$ (off path), he discloses it immediately.

The equation in Theorem 1 identifies the date at which bad evidence becomes worth disclosing. The left-hand side, $q^{B,s}_{s+a^*(s)}$, is the receiver's belief at the proposed disclosure date, $s + a^*(s)$, after observing bad evidence generated at time s. The right-hand side, $p^{\dagger}_{s+a^*(s),s}$, is an auxiliary no-disclosure belief: it is the belief the receiver would hold at that date under the conjecture that all good evidence has already been disclosed, all bad evidence with timestamps before s has already been disclosed, and no bad evidence with timestamps after s has yet been disclosed. Thus $s + a^*(s)$ is the date at which disclosing bad evidence timestamped at s first becomes as good as continued nondisclosure.

The equilibrium no-disclosure belief, denoted by p_t, is therefore piecewise. The first time the sender discloses bad evidence is $a^*(0)$, so for $t < a^*(0)$, no bad evidence is disclosed, hence $p_t = p^{\dagger}_{t,0}$. For $t \geq a^*(0)$, bad evidence with earlier timestamps has already been disclosed, whereas bad evidence with later timestamps is still being withheld. Hence, $p_t = p^{\dagger}_{s+a^*(s),s}$ where $t = s + a^*(s)$.

Figure 1 illustrates the equilibrium belief dynamics under a high prior (left panel) and a low prior (right panel). The black curve is the receiver's belief conditional on no disclosure, p_t. If good evidence arrives at time s_1, the sender discloses it immediately and the receiver's belief jumps to 1 and then follows the blue path. If bad evidence arrives at time s, where $s = s_1$ or s_2, the sender waits until the red post-disclosure belief path first reaches the black no-disclosure path. This crossing date is $s + a^*(s)$.

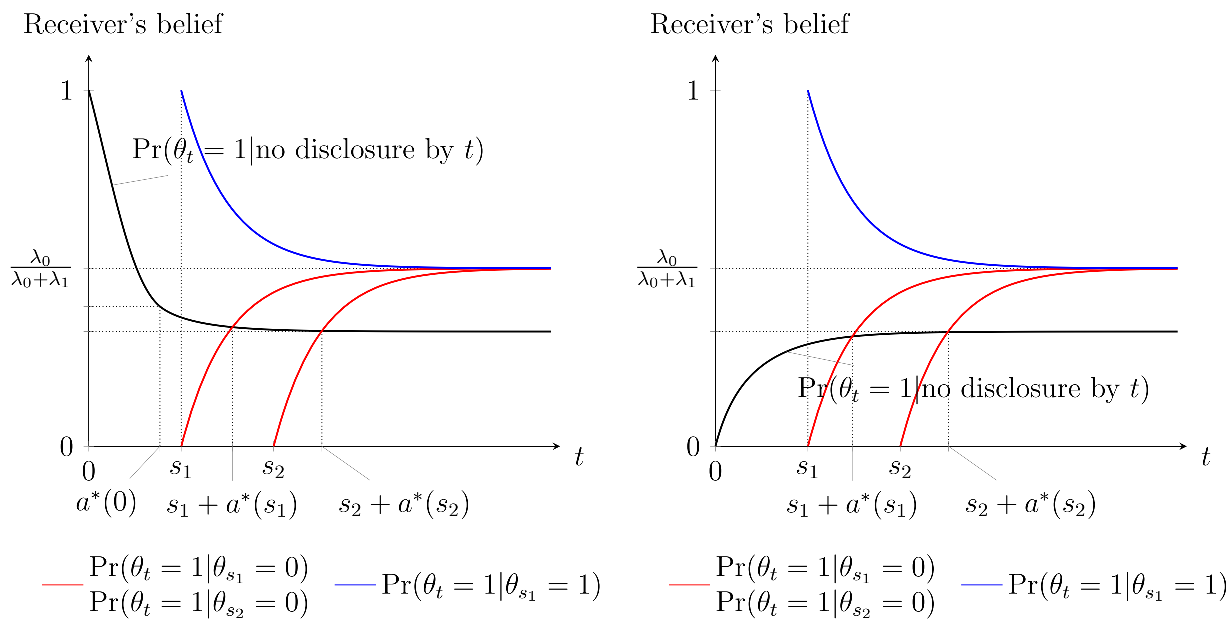

Figure 1: Equilibrium belief evolution for priors $p_0 = 1$ (left panel) and $p_0 = 0$ (right panel), and parameters $\lambda_0 = 1, \lambda_1 = 1, \mu = 0.5$, and $r = 0.1$.

We now describe the incentives underlying the equilibrium strategy and the belief paths. The argument uses two structural features of the timestamp environment. First, the receiver's post-disclosure belief is pinned down by the timestamp, regardless of whether the disclosure occurs on or off the equilibrium path. Second, because evidence arrives at most once, after any disclosure, the receiver does not expect further evidence and will update her belief post-disclosure according to the Markov process, $\Pr(\theta_t = 1 \mid \theta_s = i)$ for $i \in \{0, 1\}$.

Consider first good evidence, taking as given the equilibrium delay of bad evidence. If the sender discloses good evidence immediately, the receiver's belief jumps to 1 and then decreases according to Markov switching. Delaying disclosure does not improve the post-disclosure belief: once the evidence is disclosed, the timestamp leads the receiver to update as if evidence timestamped at s had been disclosed. Delay only lowers the sender's payoff before disclosure: the receiver's belief follows the no-disclosure path, where "no news is bad news" as bad evidence is being withheld. Hence good evidence is disclosed immediately.

The key role of timestamps for good evidence is that they prevent the sender from pretending old evidence is fresh: if good evidence is disclosed after a delay, the receiver's belief follows the same path as if it had been disclosed immediately.

It is worth noting that although the above intuition is built on a best response argument, immediate disclosure of good evidence holds in any equilibrium. That said, it is not a dominant strategy: it is not a best response to every possible conjecture of the receiver. For example, if the receiver expects bad evidence to be disclosed immediately, no disclosure would be interpreted favorably and the sender may prefer to delay good evidence. This, however, does not arise in equilibrium.

The logic for delaying disclosure of bad evidence is more subtle. Fix immediate disclosure of good evidence. First, fresh bad evidence is too damaging to disclose immediately. If bad evidence arrives at time s, disclosure at that same date sends the receiver's belief to 0, while the no-disclosure belief is strictly positive as it preserves the possibility that no evidence has arrived. The sender therefore initially delays disclosure. Second, disclosure of bad evidence cannot be delayed forever. To see why, compare the belief induced by nondisclosure with the belief induced by disclosing bad evidence. For nondisclosure, Since good evidence would have been disclosed immediately, no disclosure rules out the sender having already received good evidence, but it may mean that the sender is withholding bad evidence. Thus, "no news is bad news." By contrast, as bad evidence ages, its implication for the current state becomes less pessimistic because the state may have switched from bad to good. In addition, because evidence arrives at most once, disclosure of bad evidence implies that no further evidence will arrive. This shuts down the "no news is bad news" effect associated with nondisclosure. Hence, once bad evidence becomes sufficiently old, disclosing it can generate a higher continuation belief than continued nondisclosure.

Therefore, delaying disclosure of bad evidence trades off two forces. As bad evidence ages, it becomes less damaging, so waiting improves the payoff from eventual disclosure. But while delaying, the sender remains in the nondisclosure pool, where the receiver cannot tell whether no disclosure means no evidence has arrived or the sender is withholding bad evidence. The optimal disclosure time is when old bad evidence has become just good enough that leaving this pool is no worse than staying in it. Formally, $s + a^*(s)$ is the first crossing between the belief induced by disclosing bad evidence timestamped at s and the no-disclosure belief. Before this crossing, the bad evidence is still too fresh, so disclosure would lower the receiver's belief and delay is optimal. After this date, disclosure leads to a strictly higher belief than no disclosure, so further delay is not optimal. The key role of timestamps for bad evidence is that they allow the sender to credibly show that bad evidence is old, making disclosure optimal after some time when its implication for the current state is not as pessimistic.

3.2 Equilibrium properties

We first examine how the timing of bad-evidence disclosure changes with its timestamp and its implications on the equilibrium beliefs. The optimal disclosure date $s + a^*(s)$ is increasing in s, so later bad evidence is disclosed later in calendar time. Moreover, the delay length $a^*(s)$ is also monotone, with the direction determined by whether the prior is above or below the steady-state belief p^*. These properties are summarized in Proposition 1 and illustrated in Figure 2.

Proposition 1. *For all $s \geq 0$, (i) $s + a^*(s)$ is increasing in s; (ii) $a^*(s)$ is monotone in s. In particular, $a^*(s)$ is strictly increasing in s if $p_0 < p^*$, strictly decreasing if $p_0 > p^*$, and constant if $p_0 = p^*$; (iii)* $\lim_{s\to\infty} a^*(s) = \alpha^*$.

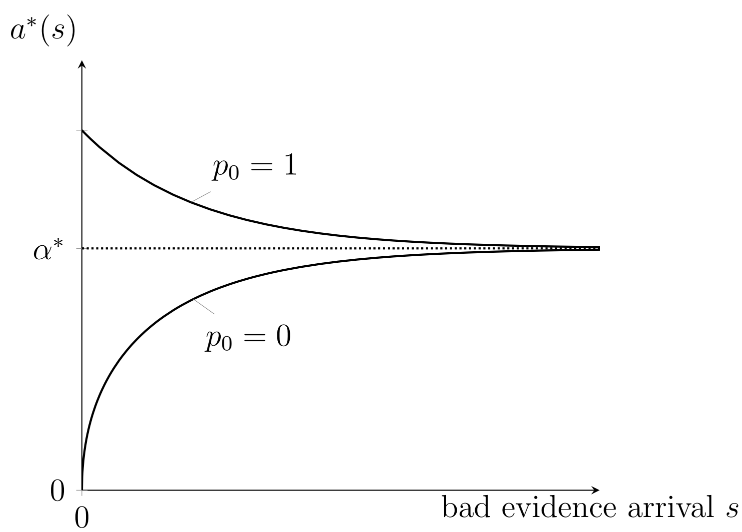

Figure 2: Optimal delay $a^*(s)$ as a function of s for priors $p_0 = 1$ and $p_0 = 0$, and parameters $\lambda_0 = 1, \lambda_1 = 1, \mu = 0.5$, and $r = 0.1$.

To understand why $s + a^*(s)$ is increasing, fix a calendar time and compare bad evidence with different timestamps. Older evidence is less damaging, while the no-disclosure belief is the same for all bad-evidence types. Thus, if more recent evidence is worth disclosing, older evidence must already be worth disclosing, so later evidence cannot be disclosed earlier.

To understand why $a^*(s)$ is monotone, fix a delay a and compare the disclosure condition at different calendar times. If bad evidence timestamped at s is disclosed at $s + a$, its disclosure value depends only on the delay a and not the calendar time because the Markov process is time-homogeneous. Calendar time therefore affects the disclosure condition only through the value of nondisclosure. For the marginal bad-evidence type, no disclosure pools the event that no evidence has arrived with the event that bad evidence arrived within the last a units of time. When $p_0 < p^*$, the unconditional belief is increasing. Consequently, no evidence is better news at later dates, and evidence arriving during the later interval is less likely to be bad. No disclosure is therefore more valuable at later dates, while the disclosure value remains unchanged, so bad evidence must age longer before the belief upon disclosure catches up, so $a^*(s)$ is increasing. When $p_0 > p^*$, the reverse holds: no disclosure becomes less valuable, disclosure catches up sooner, and $a^*(s)$ is decreasing.

The behavior of $a^*(s)$ in turn determines the behavior of the equilibrium no-disclosure belief once bad-evidence disclosure begins. As Figure 1 illustrates, for $t \geq a^*(0)$, the no-disclosure belief moves in the same direction as $a^*(s)$. When $p_0 < p^*$, $a^*(s)$ is increasing, so the marginal bad evidence disclosed at later dates is older and therefore induces a higher post-disclosure belief. Because the post-disclosure is equal to the no-disclosure belief at the optimal disclosure date, the no-disclosure belief is also higher at later dates. When $p_0 > p^*$, $a^*(s)$ is decreasing, the reverse holds.[3]

The convergence of $a^*(s)$ also determines the long-run no-disclosure belief. As $a^*(s)$ converges to α^*, the age of the marginal bad evidence converges to α^*, so the no-disclosure belief converges to the belief induced by bad evidence of that age. This limit is strictly below the unconditional steady-state belief p^*, which means no disclosure remains informative even in the long run. At large t, there is still a positive probability that the sender received bad evidence within the last α^* units of time and has not yet disclosed it. This prevents the no-disclosure belief from converging to p^*. The following corollary formalizes these results.

Corollary 1. *At each $t \geq a^*(0)$, the equilibrium no-disclosure belief p_t is strictly increasing in t if $p_0 < p^*$, strictly decreasing in t if $p_0 > p^*$, and constant in t if $p_0 = p^*$. Moreover,* $\lim_{t\to\infty} p_t = p^* \left(1 - e^{-(\lambda_0+\lambda_1)\alpha^*}\right) < p^*$.

We next show that the receiver's belief conditional on nondisclosure is minimized in the equilibrium among all possible disclosure strategies, with or without timestamps. This result extends the *Minimum Principle* from earlier work Acharya et al. (2011); Guttman et al. (2014).[4] Intuitively, for any fixed time t, the delayed disclosure of bad evidence induces a "cutoff type" in the form of the timestamp of bad evidence such that the sender is indifferent between disclosing at time t or not. All bad evidence with timestamp before this cutoff, along with all good evidence of any timestamp, is disclosed in equilibrium and would lead to a

[3] The claim that the equilibrium no-disclosure belief p_t is monotone applies only after bad-evidence disclosure begins, at $t \geq a^*(0)$; before $a^*(0)$, it may not be monotone.

[4] A related "suspicious beliefs" principle holds in Kremer et al. (2024).

higher belief at time t than p_t; conversely, all bad evidence with timestamp after the cutoff would lead to a lower belief at t than p_t. Hence, any alternative strategy would either involve nondisclosure of more favorable evidence or disclosure of less favorable evidence (or both). Both of these kinds of deviations would improve the nondisclosure pool in the sense of raising the average time-t belief conditional on nondisclosure.

Proposition 2 (The Minimum Principle)**.** *At each $t \geq 0$, the equilibrium no-disclosure belief p_t is the minimum among all disclosure strategies.*

As we shall see, the Minimum Principle does not generally hold in the no timestamps environment.

4 No timestamps

Without timestamps, the sender can attempt to manipulate the receiver's belief about the age of evidence by strategically timing its disclosure. Since this manifests differently for senders with good versus bad evidence, we divide our analysis into two parts.

4.1 Good evidence: delay and gradual purging

Without timestamps, good evidence need not be disclosed immediately in equilibrium. In a conjectured equilibrium with immediate disclosure of good evidence at all times, the sender can have an incentive to delay disclosure to pretend the evidence is fresher. We begin with a simple, but general, result that says that immediate disclosure occurs only if the prior belief is sufficiently low; for higher priors, there *must* be delayed disclosure of good evidence in equilibrium, in contrast to the timestamps environment.

Proposition 3. *If $p_0 > \frac{\lambda_0 + r}{\lambda_0 + \lambda_1 + r}$, then there is no equilibrium with immediate disclosure of good evidence at all times.*

To understand Proposition 3 and the cutoff belief, suppose the sender obtains good evidence at time 0. By disclosing this evidence immediately, the posterior belief jumps to 1 and then decays toward the long run steady state based on the Markov switching dynamics. If the sender deviates and discloses at some small time $\varepsilon > 0$, then on $[0, \varepsilon)$, the posterior belief is lower than had the sender disclosed at time 0. Since beliefs are right-continuous, to a first-order approximation, the loss from this delay is $(1 - p_0)\varepsilon$. However, when the receiver sees the disclosure at time ε, she updates to 1 assuming that the evidence is fresh; the belief then proceed to decay in the same way as before. In other words, delay results in a lower belief for the receiver in the short run, but a higher belief at all times after the disclosure. Compared to having disclosed at time 0, the belief at time ε is higher by approximately $\lambda_1\varepsilon = 1-(1-\lambda_1\varepsilon)$, where $-\lambda_1$ is the slope of the post-disclosure belief at time 0. Moreover, this higher belief results in a higher belief at all future times, albeit discounted, and the net present value of this difference is $\lambda_1\varepsilon\frac{1}{\lambda_0+\lambda_1+r}$. Put together, delay is optimal if $(1-p_0)\varepsilon < \lambda_1\varepsilon\frac{1}{\lambda_0+\lambda_1+r}$, which rearranges to the condition in Proposition 3.

Proposition 3 also implies that, for high beliefs, the Minimum Principle fails. In the timestamps environment, that principle holds because the sender discloses any evidence that would raise the receiver's belief and delays disclosure of any evidence that would lower the belief. Without timestamps, the lack of immediate disclosure of good evidence means that there is at least partial pooling between sender types with good evidence and the type with no evidence.

Although Proposition 3 implies that equilibrium must feature delayed disclosure of good evidence, it does not say *how* such delay occurs. To that end, and to isolate the role of good evidence, suppose $\lambda_0 = 0$; that is, the bad state is absorbing. (We return to $\lambda_0 > 0$ later in the section.) Then disclosure of bad evidence is strictly dominated, since it guarantees a flow payoff of 0 forever after.

The following class of equilibria is central to our analysis.

Definition 1. *A three-phase equilibrium is characterized by two deterministic times $\underline{T}, \overline{T}$ with $0 \le \underline{T} \le \overline{T} < \infty$ such that*

- Stockpiling phase: *On $[0, \underline{T})$, there is no disclosure of good evidence. Belief updating immediately after off-path disclosure of good evidence is passive.*
- Purging phase: *On $(\underline{T}, \overline{T})$, the sender discloses good evidence at rate $\kappa_t > 0$ independent of its timestamp, where κ is continuous.*
- Transparency phase: *On $[\overline{T}, \infty)$, the sender discloses good evidence immediately.*

We now present the main result of this section, which establishes a unique equilibrium within the three-phase class. Define $\overline{p} := \frac{r+\frac{\lambda_1}{2}}{r+\lambda_1}$ and $\underline{p} := \frac{r}{r+\lambda_1}$.

Theorem 2. *Assume $\lambda_0 = 0$. For all p_0, λ_1, μ and r, there exists a unique three-phase equilibrium. Moreover,*

- *If $p_0 > \overline{p}$, all three phases are nonempty.*
- *If $p_0 \in (\underline{p}, \overline{p}]$, then only the purging and transparency phases are nonempty.*
- *If $p_0 \le \underline{p}$, then only the transparency phase is nonempty.*

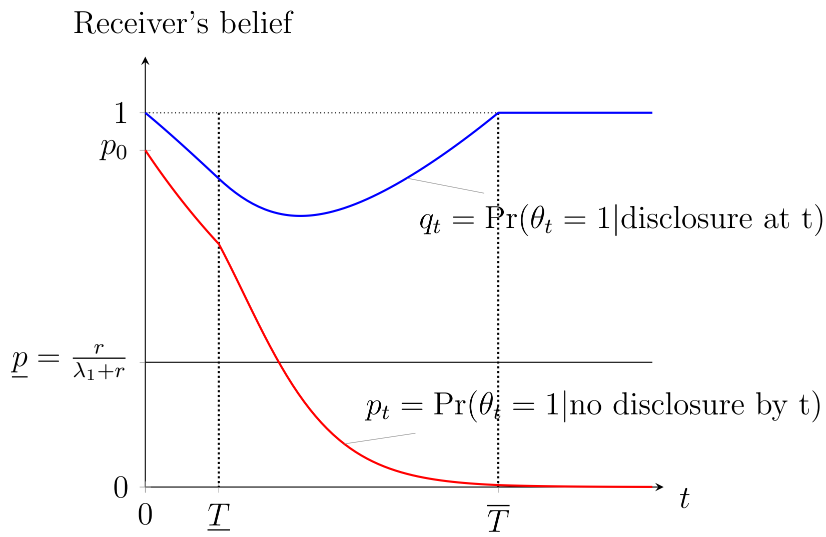

Figure 3: Belief paths in the three-phase for $p_0 > \overline{p}$, with vertical dashed lines at $\underline{T}$ and $\overline{T}$ for parameters $p_0 = 0.9, \lambda_0 = 0, \lambda_1 = 1, \mu = 3$, and $r = 0.5$.

Figure 3 illustrates the equilibrium beliefs when $p_0 > \overline{p}$; p_t denotes the receiver's belief about θ_t conditional on no disclosure, and q_t denotes the receiver's posterior immediately after disclosure of good evidence.[5] In the stockpiling phase, delay leads to an accumulation or stock of undisclosed evidence from the receiver's perspective. The receiver's belief p conditional on no disclosure falls due to Markov switching, but beyond that, "no news is no news" as the receiver anticipates delay. Meanwhile, q falls due to the aging of evidence, but this fall is dampened by the arrival of new evidence.

The purging phase begins at the time when, if the receiver had counterfactually continued to conjecture delay, the sender would strictly prefer to disclose. However, delay followed by a strict preference to disclose at time $\underline{T}$ cannot be sustained in equilibrium. If the receiver conjectured such a bang-bang disclosure strategy, then after disclosure at time $\underline{T}$, the receiver's belief would jump to a value less than 1, based on its belief distribution of past arrival dates of good evidence. By waiting until time $\underline{T} + \varepsilon$ to disclose, for $\varepsilon > 0$ arbitrarily small, the sender can induce a posterior belief arbitrarily close to 1. This marginal delay yields a discrete gain at arbitrarily small cost, so the sender would prefer to deviate. Mixing in the purging phase resolves this dilemma.

At the beginning of the purging phase, p has a downward kink as the rate of disclosure is bounded away from 0. During this phase, q initially continues to fall as the rate of disclosure is small and does not offset the aging of the stock of undisclosed evidence. However, later in the purging phase, p falls sufficiently low that indifference can only be sustained with an

[5]Note that while the belief jumps to q_t at time t, after that, it evolves according to Markov switching; i.e., the receiver's belief does not continue to be given by q at all future times after the disclosure.

increasing q; this requires κ to increase. Moreover, for q to approach its maximum value 1, it must be that the composition of the pool of undisclosed evidence shifts toward fresh evidence; this requires κ to tend to infinity as $t \to \overline{T}$. When the transparency phase begins, it is no longer possible to maintain indifference, and the sender discloses good evidence immediately as it arrives.

In Figure 3, it is noteworthy that the belief crosses $\underline{p}$ before the end of the purging phase. In other words, the sender continues to (stochastically) delay even though the receiver's belief p_t conditional on non-disclosure is sufficiently low that, had the game started at the same belief, there would be immediate disclosure at all times. The reason is that the incentive to delay depends not only on the current belief p_t but also the post-disclosure value q_t and its evolution; the presence of a stockpile of undisclosed evidence means that immediate disclosure of good evidence does not convince the receiver that the current state is 1, and delay can be weakly optimal. Thus, the variable q is a critical additional factor in the sender's best response problem.

4.2 Three-phase equilibrium construction

Recall the definition of q_t as the receiver's belief immediately after disclosure of good evidence. Clearly, $q_0 = 1$. If $t \in (0, \underline{T})$, disclosure at time t is off-path, and passive belief updating means that the receiver's posterior beliefs over the timestamp is equal to the prior belief about the timestamp given that the evidence was good and there has been no disclosure before time t. Rather than maintain a belief over the whole interval of past times, it is convenient to introduce "stock" variables

$$G_{i,t} = \Pr(\tau \leq t \text{ and } \theta_\tau = 1 \text{ and } \theta_t = i | \text{no disclosure up to } t) \qquad \text{for } i \in \{0,1\},$$

representing the probability that the sender possesses good evidence *and* the state is currently i, conditional on no disclosure yet.

It follows that for $t \in (0, \underline{T})$,

$$q_t = \frac{G_{1,t}}{G_{1,t} + G_{0,t}}.$$

The same characterization also applies during the purging phase, since the disclosure rate of good evidence is independent of its timestamp.

With abuse of notation, let p denote the (deterministic) belief process conditional on no disclosure.

In either the stockpiling or purging phase, (G_1, G_0, p) evolve according to

$$\begin{aligned}\dot{G}_{1,t} &= -G_{1,t}\lambda_1 + (p_t - G_{1,t})\mu - G_{1,t}\kappa_t(1 - G_{1,t} - G_{0,t}) \\ \dot{G}_{0,t} &= G_{1,t}\lambda_1 - G_{0,t}\kappa_t(1 - G_{1,t} - G_{0,t}) \\ \dot{p}_t &= -\lambda_1 p_t - G_{1,t}\kappa_t + p_t(G_{0,t} + G_{1,t})\kappa_t,\end{aligned}$$

with initial values $(0, 0, p_0)$. In the G_1 and G_0 ODEs, the λ_1 terms reflect that Markov switching from state 1 to state 0 increases G_0 while decreasing G_1. The $(p_t - G_{1,t})\mu$ term represents the arrival of good evidence, which occurs at rate μ when the state is good and the sender does not already possess good evidence.[6] The $-\kappa_t G_{i,t}$ terms reflect outflows due to disclosure, and the adjustment factor $(1 - G_{1,t} - G_{0,t})$ accounts for the conditioning on no disclosure yet. The p-ODE terms have an analogous interpretation.[7]

In the stockpiling phase, $\kappa \equiv 0$, and the system has a closed form solution. The local incentive compatibility condition for delay is

$$\dot{q}_t \geq r q_t - (\lambda_1 + r)p_t. \tag{1}$$

Intuitively, a marginal delay is beneficial when the current flow payoff p_t is large, the posterior after disclosing good evidence q_t is low, or when $\dot{q}_t$ is large (or not too negative). Provided $p_0 > \overline{p}$, the local delay IC condition holds until some time $\underline{T}$, when the sender becomes locally indifferent. This marks the end of the stockpiling phase.

Throughout the purging phase, the IC condition must hold with equality. The system in (G_1, G_0, p, q, κ) is a mixed differential-algebraic system rather than a standard initial value problem (IVP) since κ is not given exogenously, and instead we have an additional equation linking q to G_0 and G_1. To solve the system, we eliminate G_0 and G_1 using $S := G_0 + G_1$ and $q = G_1/S$, and then we reduce the system to an IVP in two variables (p, q) by solving for S and κ as functions of (p, q). The dynamics of this system are illustrated in Figure 4. The lightly shaded, outlined region, R, is where S and κ are well defined and κ is nonnegative. In the proof, we show that (p, q) must exit this region at some finite time $\overline{T}$ — marking the end of the purging phase — and it does so through $q_{\overline{T}} = 1$. Hence, as $t \to \overline{T}$, the purging rate κ_t explodes and the stockpile empties, $S_t \to 0$.

After the purging phase, the stockpile is empty and the receiver conjectures immediate

[6] When we extend to the case $\lambda_0 > 0$, we must also consider the possibility that the sender already possesses bad evidence.

[7] When the bad state is not absorbing, we also must introduce analogous variables $B_{i,t}$ corresponding to bad evidence. The variable B_1 enters the law of motion for p when $\lambda_0 > 0$ since it is possible that the state is good but was bad in the past and the sender already obtained bad evidence, preventing the arrival of good evidence. When $\lambda_0 = 0$, B_1 is identically 0.

disclosure at all times. Hence, $q_t = 1$ for all $t \geq \overline{T}$. Since the receiver's belief p_t is below the threshold $\underline{p}$, immediate disclosure is optimal.

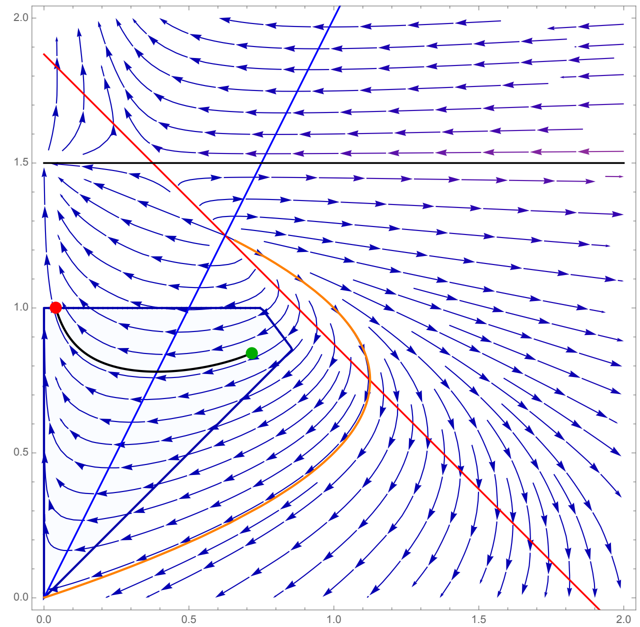

Figure 4: Phase diagram for reduced system in (p, q). The red and blue lines are isoclines for p and q, respectively, where their derivatives change sign. The ODEs are well defined below the black horizontal line. The orange curve is the set of initial conditions for which the solution converges to the absorbing steady state $(0, 0)$. The solution first exits R through $q_{\overline{T}} = 1$.

4.3 Bad evidence: delay and disclosure bursts

Throughout this section, we assume $\lambda_0 > 0$; if $\lambda_0 = 0$, then disclosing bad evidence cannot be a part of any equilibrium. We focus on equilibria with two properties: (i) age-independent disclosure of bad evidence, and (ii) immediate disclosure of good evidence. Property (i) is natural since the hidden timestamp of bad evidence is not directly payoff relevant to the sender; note that good evidence disclosure in the three-phase equilibrium of Theorem 2 was also age-independent. Property (ii) serves two purposes: first, to isolate the effect of timestamps on bad evidence disclosure, as good evidence is also disclosed immediately in the timestamps environment, and second, to allow us to hold fixed the treatment of good evidence while varying parameter values to better understand how bad evidence incentives change. To accommodate property (ii), we will study large discount rates, under which, as we will verify, immediate disclosure of good evidence is indeed optimal whenever it is conjectured.[8]

Since we are now assuming that the bad state is not absorbing, the sender can potentially benefit from disclosing bad evidence because it eliminates the penalty for not disclosing good evidence in the future. However, in contrast to the timestamps environment, the sender cannot prove that old bad evidence is indeed old and wait until disclosure is myopically beneficial; after the disclosure of bad evidence, the receiver updates based on the distribution of arrival dates. With age-independent disclosure, disclosing bad evidence unambiguously and immediately decreases the receiver's belief about the state. The reason is that bad evidence disclosure proves that the sender indeed had bad evidence, and the receiver's distribution of arrival dates simply scales up proportionally to account for the mass that was previously assigned to the probability that no evidence has arrived yet.

Thus, the sender now faces a trade-off between long-run benefit and short-run cost, and the sender's optimal disclosure timing depends critically on his discount rate. This is in stark contrast to the timestamps case, where the sender's strategy is independent of the discount rate.

The following result says that by taking r sufficiently large, bad evidence can only be disclosed after an arbitrarily long delay period. It also establishes that the only fixed good evidence strategy that would survive for arbitrarily large r is precisely immediate disclosure of good evidence, justifying our focus on equilibria with this strategy.

Proposition 4. *For any fixed good evidence strategy with positive probability of delay at some time, if r is sufficiently large, it cannot be part of an equilibrium. Moreover, for any $T > 0$, if r is sufficiently large, then in all age-independent equilibria, there is zero probability of bad evidence disclosure in $[0, T]$.*

Despite the fact that bad evidence is being withheld on $[0, T]$ (for r sufficiently large), effectively pooling bad evidence types with the no-evidence type, the Minimum Principle (Proposition 2) guarantees that at all times, the nondisclosure belief is weakly higher than in the timestamps environment. The reason is that on $[0, T]$, the bad evidence that is disclosed in the timestamps environment is precisely evidence which raises the receiver's belief (at all future times).

Proposition 4 on its own does not specify how or whether bad evidence is disclosed after time T. It turns out that different off-path beliefs can support different equilibria. Using belief threats that specify any off-path disclosure of bad evidence must be fresh, there exists an equilibrium with no disclosure of bad evidence when r is sufficiently large.

[8] Immediate disclosure of good evidence would be trivially and uniquely optimal for all discount rates in case $\lambda_1 = 0$, which we have ruled out by assumption.

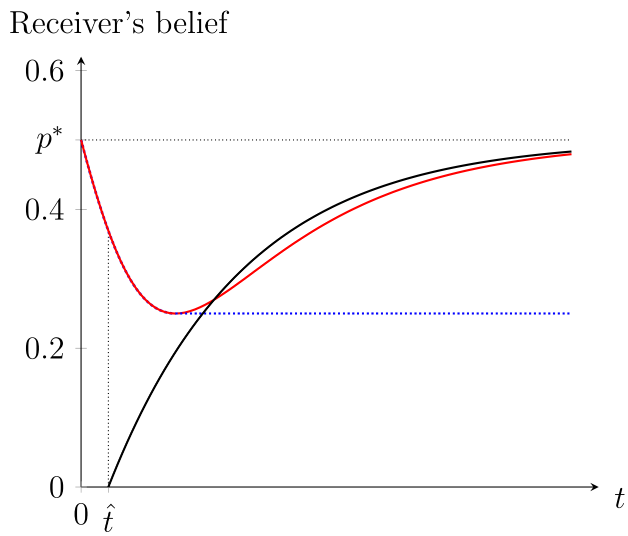

Figure 5: The receiver's beliefs in an equilibrium where the sender discloses good evidence immediately and never discloses bad evidence. The red curve shows the belief following no disclosure and the solid black curve shows the belief following disclosure off-path at $\hat{t}$. As a reference, the dotted blue curve is the receiver's belief given no disclosure in the unique timestamps equilibrium. The parameters are $\lambda_0 = 1, \lambda_1 = 1, \mu = 6$, $p_0 = p^* = 1/2$, and $r = 4$.

However, the following lemma shows that under passive belief updating, there is no equilibrium in which bad evidence is concealed forever; there is some time T at which the agent would strictly prefer to disclose bad evidence.

Lemma 1. *Under passive beliefs after off-path disclosure of bad evidence, there is no equilibrium in which good evidence is immediately disclosed and bad evidence is never disclosed. Under threat beliefs, such an equilibrium exists if r is sufficiently large.*

Under the support restriction, it follows that there cannot be an equilibrium with immediate disclosure of all bad evidence from that point forward; Lemma 2 establishes this formally. In such an equilibrium, all bad evidence disclosed after the first instant would be interpreted as fresh, so after disclosure, the belief would drop to 0 and evolve according to the "no news is no news" Markov switching dynamics. By not disclosing, the belief stays positive and still evolves according to those same dynamics: the receiver expects evidence to be disclosed at rate μ in either state, so again no news is no news. Hence, disclosure would result not only in a lower expected payoff but a lower flow payoff at all future times.

Lemma 2. *There is no equilibrium in which good evidence is always disclosed immediately and bad evidence is disclosed immediately at all $t \geq t_0$ for some $t_0 \geq 0$.*

A similar failure of "bang-bang" good evidence disclosure occurred in Section 4.1, but the underlying reason is fundamentally different. There, an atom of disclosure gives the sender an incentive to marginally delay, and this dilemma is resolved through continuous mixing prior to the transparency phase. Here, the problem is that an empty stockpile after t_0 reduces incentives for types just after t_0 to disclose, reinforcing the incentive for earlier types to disclose at t_0 itself. A smooth emptying of that stockpile by t_0 through mixing would not resolve this problem.

In sum, the preceding observations point to looking for equilibria for large r in which good evidence is immediately disclosed, and bad evidence is initially held back, but later released in a pure but *intermittent* fashion. Our main result for bad evidence, Theorem 3, shows there exists an equilibrium with exactly these properties: bad evidence is disclosed along an infinite sequence of isolated *bursts*. We utilize a stationary prior $p_0 = p^*$, which admits a stationary equilibrium, to simplify the characterization and facilitate comparison to the timestamps environment,

Theorem 3 (Bad evidence in isolated bursts)**.** *Assume $p_0 = p^*$. Then if r is sufficiently large, there exists an equilibrium in which good evidence is disclosed immediately and bad evidence is disclosed at the next time in the set $\{0, \Delta, 2\Delta, \dots\}$, where $\Delta > 0$. Passive off-path beliefs with the support restriction are sufficient to support this equilibrium.*

The existence of such an equilibrium rests on two key incentive constraints. The first constraint is that in between bursts, if the sender has bad evidence, he must be willing to wait until the next burst; this constraint places an upper bound on Δ.[9] The second constraint is that the sender must be willing to disclose bad evidence at any given burst rather than wait until the next one; this constraint places a lower bound on Δ. These upper and lower bounds both tend to infinity as $r \to \infty$, and the proof shows that they have the correct ranking when r is large.

The equilibrium in Theorem 3 exhibits stationary behavior illustrated in Figure 6. The unconditional belief ϕ_t remains constant at p^* by definition. In between bursts, only good evidence is disclosed, so no news is bad news, causing the no-disclosure belief to decay, but this decay can be offset eventually by upward mean-reversion due to the hidden Markov switching, resulting in the hook shape. Since each bursts results in the disclosure of all possessed evidence, the belief resets to p^*. Meanwhile, the passive off-path resets to 0 after each burst due to the empty stockpile, and it rises as the stock of bad evidence ages.

As a consequence of Lemma 3, which we discuss later, the sender's expected payoff is the same in both the timestamps equilibrium and the burst equilibrium we have constructed.

[9] Of course, satisfying this constraint for passive off-path beliefs means it would also be satisfied for .

By Proposition 2, the sender's flow payoff prior to any disclosure is higher in the burst equilibrium. Furthermore, his payoff after disclosure of good evidence is identical in both equilibria. It follows that, for each sample path in which good evidence eventually arrives, the sender prefers the bursts equilibrium. Indifference then implies that averaging over the paths in which bad evidence eventually arrives, the sender prefers the timestamps equilibrium.

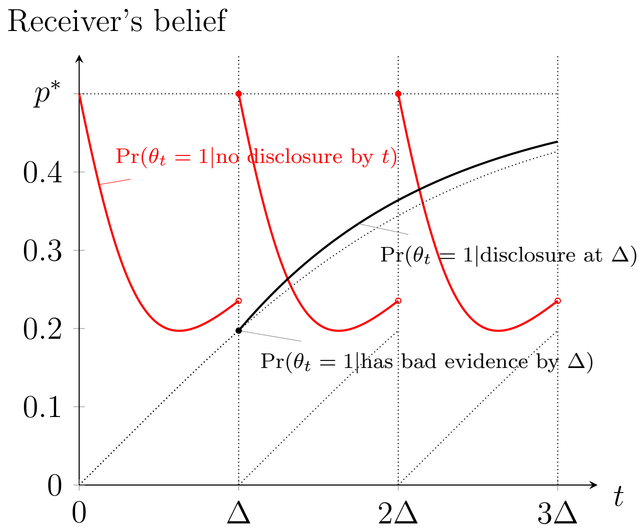

Figure 6: The receiver's beliefs in a stationary burst equilibrium. The red curve shows the belief following no disclosure, the dotted black curve shows the off-path belief following disclosure between bursts, and the solid black curve shows the belief following disclosure at the burst time Δ. The parameters are $\lambda_0 = 1, \lambda_1 = 1, \mu = 10, p_0 = p^* = 1/2$, $r = 5$, and $\Delta = 0.4$.

5 Discussion

5.1 Modeling assumptions

Our model assumes the state evolves over time. Holding fixed the rest of the model, this assumption is in fact necessary for timestamps to play a nontrivial role; with a constant state, the sender would immediately disclose good evidence and never disclose bad evidence, for all values of the other parameters.

We assume that evidence only arrives once. This is appropriate in settings where hard evidence is not continuously available, either because producing multiply pieces of evidence is infeasible or too costly. We conjecture that (i) delayed disclosure of good evidence and (ii) some disclosure of bad evidence would continue to be present in a richer model that allowed multiple arrivals of evidence. However, such a model would be more technically challenging

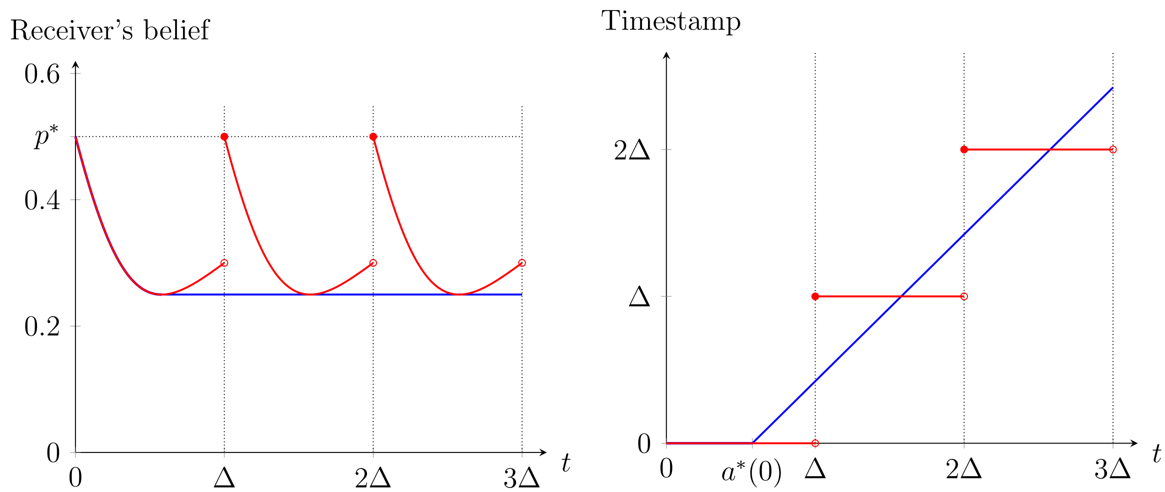

Figure 7: The left panel compares the receiver's no-disclosure beliefs, and the right panel compares the most recently disclosed bad evidence, under the unique timestamp equilibrium (blue) and a stationary burst equilibrium without timestamps (red). The parameters are $\lambda_0 = 1, \lambda_1 = 1, \mu = 6$, $p_0 = p^* = 1/2$, $r = 4$, and $\Delta = 0.6$.

as the receiver would have to maintain beliefs about a higher-dimensional evidence status of the sender.

We assume that the sender has no private information about the state other than the evidence she obtains. Since the sender's preferences are state-independent, equilibria in our model would also be equilibria in a model where the sender observes the state at all times. The latter model would allow for stronger off-path belief threats; for example, the receiver could assume that if good evidence is unexpectedly disclosed, the sender must know that the state is *currently* bad.

We assume that the only information available to the receiver (other than time itself and knowledge of the hidden Markov switching) is disclosure by the sender. If the model included exogenous public news, the evidence's timestamp would help the sender forecast continuation payoffs, changing the disclosure incentives.

5.2 Robustness and other equilibria

Theorem 2 uses passive belief updating off-path. A natural question is whether more pessimistic off-path beliefs can sustain a three-phase equilibrium with a longer stockpiling phase. The following result shows that they cannot; while the incentive to delay can indeed be sustained longer with lower off-path beliefs, the construction for the purging phase uniquely pins down the start time $\underline{T}$. The intuition is that the indifference condition at the beginning of the purging phase constrains the stock of undisclosed good evidence and its composition (i.e., the proportion corresponding to the state being good at time $\underline{T}$). These values depend on the length of the stockpiling phase, but not the off-path beliefs during it.

Corollary 2. *Fixing parameter values with $\lambda_0 = 0$, the lengths of the stockpiling and purging phases are the same across all three-phase equilibria.*

Although we do not establish uniqueness within the general class of equilibria, the sender's expected payoff is the same in all equilibria, given the parameter values. More strongly, following result establishes that as long as the receiver's conjecture is correct (independent of sender optimality), the sender's expected payoff is the same, and thus the sender can also not benefit from commitment.

Lemma 3. *Fix parameter values. Fix any strategy of the sender, and let $(p_t)_{t\geq 0}$ denote the receiver's posterior belief process when that strategy is followed and the receiver's conjecture is correct. Then for all $t \geq 0$, $\mathbb{E}[p_t] = \phi_t = p^* + (p_0 - p^*)e^{-(\lambda_0+\lambda_1)t}$, which is independent of the sender's strategy.*

The proof is simply the following: by definition, p_t is the receiver's time-t expectation of the state θ_t, and by the law of iterated expectations, $\mathbb{E}[p_t] = \mathbb{E}[\theta_t] = \phi_t$, the receiver's belief based only on knowledge of the Markov switching dynamics.

Our restriction to $\lambda_0 = 0$ in the no-timestamps environment is useful as it isolates the analysis of good evidence disclosure by shutting down disclosure of bad evidence. That said, our results can be extended to $\lambda_0 > 0$. The following result gives a sufficient condition for there to continue to be an equilibrium with no disclosure of bad evidence.[10]

Proposition 5. *Assume $\mu < r$. Then for all sufficiently small $\lambda_0 > 0$, there exists an equilibrium with the same three-phase structure for good evidence and no disclosure of bad evidence.*

Within the no-timestamps environment, good and bad evidence have an interesting qualitative difference. With good evidence, if the receiver is conjecturing an atom of disclosure, the sender has an incentive to wait until just after the atom so that the evidence appears fresh. In this sense, the conjecture and the best reply are strategic substitutes. In contrast, with bad evidence, if the receiver conjectures an atom of disclosure, delay until just after the atom becomes *less* attractive, since fresh bad evidence is very harmful for beliefs. In this sense, a conjecture of fast disclosure of bad evidence can reinforce and incentive to disclose bad evidence early.

[10]Numerically, the equilibria in Theorem 2 are robust to small $\lambda_0 > 0$, and we conjecture that this can be proven by continuity with respect to λ_0.

A Proofs

A.1 Proof of Theorem 1

The proof proceeds in the following main steps. First, we show that in any equilibrium, good evidence is disclosed immediately (Lemma 4). Second, we show that bad evidence with any timestamp τ is disclosed after some (possibly infinite) time $D(\tau)$ increasing in τ (Corollary 3). Third, we characterize $D(\tau)$ uniquely.

Toward what follows, recall that $q_t^{G,\tau}$ and $q_t^{B,\tau}$ denote the belief after disclosure at time t of good and bad evidence, respectively, with timestamp τ. They admit closed-form solutions

$$q_t^{G,\tau} = \frac{\lambda_0}{\lambda_0 + \lambda_1} + \frac{\lambda_1}{\lambda_0 + \lambda_1} e^{-(\lambda_0+\lambda_1)(t-\tau)} \quad \text{and} \quad q_t^{B,\tau} = \frac{\lambda_0}{\lambda_0 + \lambda_1} - \frac{\lambda_0}{\lambda_0 + \lambda_1} e^{-(\lambda_0+\lambda_1)(t-\tau)},$$

and are continuous in t and τ. Moreover, for a fixed t, $q_t^{G,\tau}$ is strictly increasing in τ, while $q_t^{B,\tau}$ is strictly decreasing in τ. Fix an arbitrary conjecture about the sender's strategy, and let p be the belief path conditional on no disclosure.

A.1.1 Immediate disclosure of good evidence with timestamps

Lemma 4. *In any equilibrium, if the sender possesses good evidence at any time, he discloses it immediately.*

In the rest of this section, we prove Lemma 4.

Lemma 5. *If $p_t < q_t^{G,0}$ for all times, then immediate disclosure of good evidence is optimal regardless of its timestamp.*

Proof. Suppose the sender has good evidence with timestamp τ and hasn't disclosed it before time t. If the sender discloses at time $t' \geq t$, her flow payoff is p_s for $s \in [t, t')$ and $q_t^{G,\tau} \geq q_t^{G,0} > p_t$ for all $s \in [t', \infty)$. Immediate disclosure maximizes the flow payoff at all times. □

Lemma 6. *If the sender possesses bad evidence at any time t and $p_t > q_t^{B,0}$, then there exists $\varepsilon > 0$ such that the sender does not disclose it before time $t + \varepsilon$.*

Proof. Suppose the sender has bad evidence with timestamp τ and hasn't disclosed it before time t. Since p and q^B are right-continuous, there exists $\varepsilon > 0$ such that $p_s > q_s^{B,0} \geq q_s^{B,\tau}$ for all $s \in [t, t+\varepsilon)$. Hence, by disclosing at time $t+\varepsilon$ instead of t, the flow payoff is strictly higher on $[t, t+\varepsilon)$ and is unchanged on $[t+\varepsilon, \infty)$. □

Define $t^G = \inf\{t \geq 0 : p_t \geq q_t^{G,0}\}$ and $t^B = \inf\{t \geq 0 : p_t \leq q_t^{B,0}\}$. By right-continuity of p_t, $q_t^{G,0}$, and $q_t^{B,0}$ in t, if $p_0 \in (0,1)$, then $t^G, t^B > 0$. Right-continuity also implies $p_{t^G} \geq q_{t^G}^{G,0}$ if t^G is finite, and likewise $p_{t^B} \leq q_{t^B}^{B,0}$ if t^B is finite.

We now show that $t^G = \infty$. Clearly, since $q_t^{G,0} \neq q_t^{B,0}$ for all t, we cannot have $t^G = t^B < \infty$. Hence, if $t^G < \infty$, we have either $t^G < t^B$ or $t^B < t^G$.

The following lemma establishes that if t^G is finite, p must jump above $q^{G,0}$ from below $q^{B,0}$ at time t^G.

Lemma 7. *If $t^G < \infty$, then $t^B < t^G$ and* $\sup\{t < t^G : p_t \leq q_t^{B,0}\} = t^G$.

Proof. Toward a contradiction, suppose there exists $\varepsilon > 0$ such that for all $t \in [t^G - \varepsilon, t^G)$, $p_t > q_t^{B,0}$. By right-continuity, we can choose this ε such that $p_t > q_t^{B,0}$ for all $t \in [t^G - \varepsilon, t^G + \varepsilon)$. To simplify notation, set $t_0 := t^G - \varepsilon$ and $t_1 := t^G + \varepsilon$. Then by Lemma 6, there is no disclosure of bad evidence on $[t_0, t_1)$. Hence, on $[t_0, t_1)$, p is bounded above by the solution to $\dot{y}_t = \lambda_0(1 - y_t) - \lambda_1 y_t$ with initial condition $y_{t_0} = p_{t_0}$. That is, $y_t = \mathbb{P}_{t_0}(\theta_t = 1)$ is the unconditional belief based on Markov switching, starting from p_{t_0}. To prove this, let $\mathbb{E}_t$ denote the receiver's expectation operator at time t, and $\mathbb{P}_t$ the probability measure at time t. By the law of iterated expectations, for all $t \in [t_0, t_1)$,

$$y_t = \mathbb{P}_{t_0}(\theta_t = 1) = \mathbb{P}_{t_0}(\text{disclosure by } t)\mathbb{P}_{t_0}(\theta_t = 1|\text{disclosure by } t) + \mathbb{P}_{t_0}(\text{no disclosure by } t)p_t.$$

Because $p_t > q_t^{B,0}$, no bad evidence is disclosed in $[t_0, t_1)$, so any disclosure is of good evidence. The lowest post-disclosure belief from good evidence is $q_t^{G,0}$; good evidence with any later timestamp yields a strictly higher belief. So $\mathbb{P}_{t_0}(\theta_t = 1|\text{disclosure by } t) \geq q_t^{G,0}$.

Moreover, $y_t < q_t^{G,0}$. This is because y and q^G satisfy the same linear ODE on $[t_0, t_1)$ and differ only in their at the initial time; we have $q_{t_0}^{G,0} > y_{t_0}$, so by the comparison theorem, $y_t < q_t^{G,0}$ for all $t \geq t_0$.

Putting these facts together, we have

$$\begin{aligned} y_t &= \mathbb{P}_{t_0}(\text{disclosure by } t)\mathbb{P}_{t_0}(\theta_t = 1|\text{disclosure by } t) + \mathbb{P}_{t_0}(\text{no disclosure by } t)p_t \\ &\geq \mathbb{P}_{t_0}(\text{disclosure by } t)q_t^{G,0} + \mathbb{P}_{t_0}(\text{no disclosure by } t)p_t \\ &> \mathbb{P}_{t_0}(\text{disclosure by } t)y_t + \mathbb{P}_{t_0}(\text{no disclosure by } t)p_t. \end{aligned}$$

Since $\mathbb{P}_{t_0}(\text{no disclosure by } t) > 0$, the last inequality implies that $p_t < y_t < q_t^{G,0}$ for all $t \in [t_0, t_1)$, contradicting the definition of t^G. □

Lemma 8. *It must be that $t^G = \infty$.*

Proof. By Lemma 7, if $t^G < \infty$, then at time t, p must jump from weakly below $q^{B,0}$ to weakly above $q^{G,0}$. Such a jump requires that good or bad evidence is conjectured to be disclosed with an atom of probability at time t^G. Now if only an atom of good evidence disclosure is conjectured, p would exhibit a downward discontinuity, since any good evidence disclosure would raise the belief given that $p_{t^G-} \le q^{B,0}_{t^G} < q^{G,0}_{t^G}$. Thus, at time t^G, the receiver must conjecture that bad evidence is disclosed with an atom of probability. But then disclosing bad evidence at time t^G would not be optimal, as there exists $\varepsilon > 0$ such that $p_{t^G} \ge q^{G,0}_{t^G} > q^{B,0}_{t^G} \ge q^{B,\tau}_{t^G}$ for all $t \in [t^G, t^G+\varepsilon)$ and all $\tau \in [0, t^G]$. In other words, the sender would strictly prefer to wait and disclose at time $t^G + \varepsilon$. ☐

Since $t^G = \infty$, we have $p_t < q^{G,0}_t$ for all $t \ge 0$, and thus it is optimal to disclose good evidence immediately.

A.1.2 Delayed disclosure of bad evidence with timestamps

Lemma 9. *In equilibrium, p is continuous.*

Proof. If p has a discontinuity at some time t_0, then it must be that at time t_0, evidence is disclosed with an atom of probability. Since good evidence is disclosed immediately by Lemma 4 and arrives according to an atomless distribution, it must be that bad evidence is disclosed with an atom of probability at time t_0. Note that for any type τ_0 with bad evidence that is willing to disclose at t_0, it must be that $q^{B,\tau_0}_{t_0} \ge p_{t_0}$, otherwise by right-continuity of p and continuity of q^{B,τ_0} we would have $p_t > q^{B,\tau_0}_t$ for all $t \in [t_0, t_0+\varepsilon)$ for some $\varepsilon > 0$, and the sender would strictly prefer to disclose at $t_0 + \varepsilon$.

We now argue that if there is a discontinuity at t_0, it must be that $p_{t_0} < p_{t_0-}$, where the left limit exists because p is càdlàg. For any $t < t_0$, let N and G denote the events that there is no disclosure or there is good evidence disclosure in $[t, t_0]$ (suppressing dependence on t and t_0). Let B^{t_0} denote the event that there is bad evidence disclosure at time t_0, and $B^{[t,t_0)}$ the event that there is bad evidence disclosure in $[t, t_0)$. Let $\mathbb{E}_t$ and $\mathbb{P}_t$ denote expectations and probabilities conditional on no disclosure prior to t. By the law of iterated expectations,

$$\mathbb{E}_t[\theta_{t_0}] = \mathbb{P}_t(N)p_{t_0} + \mathbb{P}_t(G)\mathbb{E}_t[\theta_{t_0}|G] + \mathbb{P}_t(B^{[t,t_0)})\mathbb{E}_t[\theta_{t_0}|B^{[t,t_0)}] + \mathbb{P}_t(B^{t_0})\mathbb{E}_t[\theta_{t_0}|B^{t_0}], \qquad (2)$$

where $\mathbb{E}_t[\theta_{t_0}|B^{t_0}] \ge p_{t_0}$ by the argument above. As $t \uparrow t_0$, $\mathbb{P}_t(G) \to 0$ and $\mathbb{P}_t(B^{[t,t_0)}) \to 0$, so $\mathbb{P}_t(N) + \mathbb{P}_t(B^{t_0}) \to 1$. To simplify notation, define $\alpha := \lim_{t\uparrow t_0} \mathbb{P}_t(N)$, $\beta := \lim_{t\uparrow t_0} \mathbb{P}_t(B^{t_0}) = 1-\alpha$, and $\gamma := \lim_{t\uparrow t_0} \mathbb{E}_t[\theta_{t_0}|B^{t_0}] \ge p_{t_0}$. Then $\alpha > 0$ because the probability of no disclosure in $[t, t_0)$ is at least the probability that evidence arrives after time t_0, which is strictly positive; moreover, $\beta > 0$ because there is an atom of bad evidence disclosure at time t_0.

Thus, the limit of the right-hand side of (2) as $t \uparrow t_0$ is $\alpha p_{t_0} + \beta\gamma \geq \alpha p_{t_0} + \beta p_{t_0} = p_{t_0}$. The limit of the left-hand side of (2) is p_{t_0-}. It follows that $p_{t_0-} \geq p_{t_0}$. A discontinuity at t_0 must therefore imply that $p_{t_0-} > p_{t_0}$. This further implies $\gamma = \frac{p_{t_0-} - \alpha p_{t_0}}{\beta} = \frac{p_{t_0-} - \alpha p_{t_0}}{1-\alpha} > p_{t_0-}$.

Since $\lim_{t\uparrow t_0} \mathbb{E}_t[\theta_{t_0}|B^{t_0}] = \gamma > p_{t_0-}$, we have $\mathbb{E}_t[\theta_{t_0}|B^{t_0}] > p_{t_0-}$ for t sufficiently close to t_0. As $\mathbb{E}_t[\theta_{t_0}|B^{t_0}]$ is a convex combination over τ of $q_{t_0}^{B,\tau}$ values, there exists some timestamp $\tau < t_0$ for bad evidence for which the sender is willing to wait until t_0 to disclose and $q_{t_0}^{B,\tau} > p_{t_0-}$. By definition of p_{t_0-}, for any $\delta > 0$, there exists $\varepsilon > 0$ such that if $t > t_0 - \varepsilon$, $|p_t - p_{t_0-}| < \delta$. Choosing δ sufficiently small and ε accordingly, we have $p_t < q_t^{B,\tau}$ for all $t \in (t_0 - \varepsilon, t_0)$. But then the sender would strictly prefer to disclose at any such t rather than wait until time t_0, contradicting optimality. $\square$

Lemma 10 (Single-crossing)**.** *If $p_{t_0} \leq q_{t_0}^{B,\tau_0}$ for some $0 \leq \tau_0 \leq t_0$, then $p_t < q_t^{B,\tau_0}$ for all $t > t_0$.*

Proof. Assume $p_{t_0} \leq q_{t_0}^{B,\tau_0}$ for some $0 \leq \tau_0 \leq t_0$. Define $t^* = \inf\{t > t_0 : p_t \geq q_t^{B,\tau_0}\}$. We aim to show that $t^* = \infty$. First, we rule out $t^* \in (t_0, \infty)$, and then we rule out $t^* = t_0$. Toward a contradiction, suppose $t^* \in (t_0, \infty)$.

For each $\tau \leq t^*$, define $F(\tau) = \max_{t\in[\max\{t_0,\tau\},t^*]} \left(q_t^{B,\tau} - p_t\right)$. Since q^B is continuous in both arguments and p is continuous by Lemma 9, F is continuous by the maximum theorem. Moreover, F is strictly decreasing since q^B is strictly decreasing in τ and the set $[\max\{t_0, \tau\}, t^*]$ is weakly decreasing in τ in the set inclusion sense; it satisfies $F(\tau_0) > 0$ since $t^* > t_0$ by assumption, and by the definition of t^* we cannot have $p_t \geq q_t^{B,\tau_0}$ for any $t \in (t_0, t^*)$; and it satisfies $F(t^*) < 0$ since $q_{t^*}^B(t^*) = 0 < p_{t^*}$. Thus, there exists $\overline{\tau}$ such that $F(\tau) \geq 0$ if and only if $\tau \leq \overline{\tau}$. Since $F(\overline{\tau}) = 0$, there exists some $t_1 \in [\max\{t_0, \overline{\tau}\}, t^*)$ where $q_{t_1}^{B,\overline{\tau}} = p_{t_1}$, with $q_t^{B,\overline{\tau}} \leq p_t$ for all $t \in [\max\{t_0, \overline{\tau}\}, t^*]$.

Define $\phi_t(p_{t_1}, t_1)$ as the solution to $\dot{y}_t = \lambda_0(1 - y_t) - \lambda_1 y_t$ subject to the initial condition $y_{t_1} = p_{t_1}$. Note that q_t^{B,τ_0} solves the same ODE. Since $p_{t_1} < q_{t_1}^{B,\tau_0}$, it follows that $\phi_{t^*}(p_{t_1}, t_1) < q_{t^*}^{B,\tau_0} = p_{t^*}$. Now let N, G, and B denote the events that no evidence, good evidence, and bad evidence, respectively, is disclosed in $[t_1, t^*]$. By the law of iterated expectations,

$$\mathbb{E}_{t_1}[\theta_{t^*}] = \phi_{t^*}(p_{t_1}, t_1) = \mathbb{P}_{t_1}(N)p_{t^*} + \mathbb{P}_{t_1}(G)\mathbb{E}_{t_1}[\theta_{t^*}|G] + \mathbb{P}_{t_1}(B)\mathbb{E}_{t_1}[\theta_{t^*}|B].$$

Now $\mathbb{E}_{t_1}[\theta_{t^*}|G] > q_{t^*}^{B,\tau_0} = p_{t^*}$. Hence,

$$\phi_{t^*}(p_{t_1}, t_1) \geq (1 - \mathbb{P}_{t_1}(B))p_{t^*} + \mathbb{P}_{t_1}(B)\mathbb{E}_{t_1}[\theta_{t^*}|B].$$

Recall that $p_{t^*} > \phi_{t^*}(p_{t_1}, t_1)$, and $\mathbb{P}_{t_1}(B) < 1$, so we must have

$$\begin{aligned} &\phi_{t^*}(p_{t_1}, t_1) > (1 - \mathbb{P}_{t_1}(B))\phi_{t^*}(p_{t_1}, t_1) + \mathbb{P}_{t_1}(B)\mathbb{E}_{t_1}[\theta_{t^*}|B] \\ \implies &\phi_{t^*}(p_{t_1}, t_1) > \mathbb{E}_{t_1}[\theta_{t^*}|B], \end{aligned}$$

where $\mathbb{P}_{t_1}(B) > 0$ as the first inequality cannot hold if $\mathbb{P}_{t_1}(B) = 0$. Since $\mathbb{E}_{t_1}[\theta_{t^*}|B] = \mathbb{E}_{t_1}[q_{t^*}^{B,\tau}|B]$, this means there must be some timestamp τ and some time $t \in (t_1, t^*)$ for which bad evidence is disclosed at time t and $q_{t^*}^{B,\tau} < \phi_{t^*}(p_{t_1}, t_1)$. Since $\phi_{\cdot}(p_{t_1}, t_1) = q_{\cdot}^{B,\overline{\tau}}$ on $[t_1, t^*]$ (as $q_{\cdot}^{B,\overline{\tau}}$ solves the same ODE with same initial value $q_{t_1}^{B,\overline{\tau}} = p_{t_1} = \phi_{t_1}(p_{t_1}, t_1)$), we have $q_{t^*}^{B,\tau} < q_{t^*}^{B,\overline{\tau}}$, and because $q_s^{B,\cdot}$ is strictly decreasing in the timestamp on $[0, s]$ for each fixed s, this implies $\tau > \overline{\tau}$. But then we have $F(\tau) < 0$, so $p_t > q_t^{B,\tau}$ for all $t \in (\max\{t_0, \tau\}, t^*)$, and this implies that the sender would strictly prefer to delay the disclosure until time t^*, a contradiction.

Next, we rule out the case $t^* = t_0$ by adapting arguments from the previous case. Toward a contradiction, suppose that $t^* = t_0$. Then for all $\varepsilon > 0$ there exists $\hat{t} \in (t_0, t_0 + \varepsilon)$ such that $p_{\hat{t}} \geq q_{\hat{t}}^{B,\tau_0}$. By arguments similar to those used in the previous case, there exists some timestamp $\overline{\tau} \geq \tau_0$ and time $t_1 \in [\max\{\overline{\tau}, t_0\}, \hat{t})$ such that $p_{t_1} = q_{t_1}^{B,\overline{\tau}}$ and $p_t \geq q_t^{B,\overline{\tau}}$ for all $t \in [\max\{\overline{\tau}, t_0\}, \hat{t}]$. Since $\phi_{\cdot}(p_{t_1}, t_1)$ and $q_{\cdot}^{B,\overline{\tau}}$ solve the same ODE and coincide at t_1, we have $\phi_{\cdot}(p_{t_1}, t_1) = q_{\cdot}^{B,\overline{\tau}}$ for all $s \geq t_1$, and in particular for $s = \hat{t}$.

Let N, G, and B be the events of no disclosure, good evidence disclosure, and bad evidence disclosure on $[t_1, \hat{t}]$. By similar logic to that used in the proof for the $t^* \in (t_0, \infty)$ case, we have $\phi_{\hat{t}}(p_{t_1}, t_1) = q_{\hat{t}}^{B,\overline{\tau}} \leq q_{\hat{t}}^{B,\tau_0} \leq p_{\hat{t}}$ and also $\mathbb{E}_{t_1}[\theta_{\hat{t}}|G] > q_{\hat{t}}^{B,\tau_0} \geq q_{\hat{t}}^{B,\overline{\tau}}$. Thus, by the law of iterated expectations,

$$\begin{aligned} \mathbb{E}_{t_1}[\theta_{\hat{t}}] = q_{\hat{t}}^{B,\overline{\tau}} &= \mathbb{P}_{t_1}(N)p_{\hat{t}} + \mathbb{P}_{t_1}(G)\mathbb{E}_{t_1}[\theta_{\hat{t}}|G] + \mathbb{P}_{t_1}(B)\mathbb{E}_{t_1}[\theta_{\hat{t}}|B] \\ &> \mathbb{P}_{t_1}(N)q_{\hat{t}}^{B,\overline{\tau}} + \mathbb{P}_{t_1}(G)q_{\hat{t}}^{B,\overline{\tau}} + \mathbb{P}_{t_1}(B)\mathbb{E}_{t_1}[\theta_{\hat{t}}|B]. \end{aligned}$$

As in the previous case, the strict inequality implies $\mathbb{P}_{t_1}(B) > 0$. Thus,

$$q_{\hat{t}}^{B,\overline{\tau}} > \mathbb{E}_{t_1}[\theta_{\hat{t}}|B].$$

Since the expectation is a convex combination of $q_{\hat{t}}^{B,\tau}$ values over $\tau \leq \hat{t}$, there must be disclosure at some time $t \in [t_1, \hat{t}]$ of bad evidence with some timestamp $\tau > \overline{\tau}$. But at t, we have $p_t \geq q_t^{B,\overline{\tau}} > q_t^{B,\tau}$. By continuity in t, this inequality holds for all times in an interval $[t, t + \varepsilon)$ for some $\varepsilon > 0$, and the sender would strictly prefer to delay disclosure until $t + \varepsilon$, contradicting optimality. This rules out $t^* = t_0$. $\square$

Corollary 3 (Deterministic delays)**.** *In equilibrium, bad evidence with timestamp τ is disclosed at some time $D(\tau) \in (\tau, \infty]$. Moreover, there exists $\overline{\tau} \in (0, \infty]$ such that for $\tau \in [0, \overline{\tau})$, $D(\tau)$ is strictly increasing and finite, and for $\tau \geq \overline{\tau}$, $D(\tau) = \infty$.*

Proof. Suppose that at some time t_0, the sender has bad evidence with timestamp $\tau_0 \leq t_0$. Let $D(\tau_0, t_0) := \inf\{t > t_0 : p_t \leq q_t^{B,\tau_0}\}$. (Note that $D(\tau_0, t_0) > \tau_0$, since $q_{\tau_0}^{B,\tau_0} = 0 < p_{\tau_0}$.) We claim that disclosing at time $D(\tau_0, t_0)$ is optimal. If $D(\tau_0, t_0) = \infty$, then we have $p_t > q_t^{B,\tau_0}$ for all $t \geq t_0$, and clearly it is optimal for the sender to never disclose the evidence. If $D(\tau_0, t_0) < \infty$, then we have (i) $p_t > q_t^{B,\tau_0}$ for all $t < D(\tau_0, t_0)$, which implies disclosing at time $D(\tau_0, t_0)$ yields a higher payoff than disclosing at any earlier time, and (ii) $p_t < q_t^{B,\tau_0}$ for all $t > D(\tau_0, t_0)$, which implies that disclosing at time $D(\tau_0, t_0)$ yields a higher payoff than disclosing at any later time. On path, it is therefore optimal to disclose at time $D(\tau) := D(\tau, \tau)$.

Define $\overline{\tau} := \inf\{\tau \geq 0 : D(\tau) = \infty\}$. Since $q_t^{B,\tau}$ is strictly increasing in t and strictly decreasing in τ, we have that $D(\tau)$ is increasing on $[0, \infty)$ and strictly increasing on $[0, \overline{\tau})$. $\square$

We now turn to characterizing $\tau \mapsto D(\tau)$. Note that, in equilibrium, when $p_t = q_t^{B,s}$ for some s, p_t admits a direct characterization via Bayes' rule:

$$p_t = p_{t,s}^{\dagger} := \frac{e^{-\mu t}\phi_t + \int_s^t \mu e^{-\mu\tau}(1 - \phi_\tau) q_t^{B,\tau} d\tau}{e^{-\mu t} + \int_s^t \mu e^{-\mu\tau}(1 - \phi_\tau) d\tau}. \tag{3}$$

As mentioned in the main text, $p_{t,s}^{\dagger}$ is an auxiliary belief process that specifies the probability that $\theta_t = 1$ conditional on no disclosure up to time t, under a conjecture that (i) good evidence is always disclosed immediately and (ii) bad evidence with timestamp before s is disclosed before t, and no evidence with timestamp after s is disclosed before t. We show that $p_{t^{\dagger}(s)} = p_{t^{\dagger}(s),s}^{\dagger}$ for all $s \geq 0$.[11]

Lemma 11. *If $\lambda_0, \lambda_1 > 0$, then for all $s \geq 0$, there exists $t^{\dagger}(s)$ that solves the following equation in t:*

$$p_{t,s}^{\dagger} = q_t^{B,s}. \tag{4}$$

Moreover, $t^{\dagger}(s)$ is unique and lies in (s, ∞).

[11] For fixed s, $p_{t,s}^{\dagger}$ does not coincide at all times with the equilibrium belief process p_t, as these belief processes are based on different conjectures about the disclosure of bad evidence.

Proof. Expanding (4) and multiplying through by the denominator yields

$$e^{-\mu t}\phi_t + \int_s^t \mu e^{-\mu\tau}(1-\phi_\tau)q_t^{B,\tau}d\tau = q_t^{B,s}\left[e^{-\mu t} + \int_s^t \mu e^{-\mu\tau}(1-\phi_\tau)d\tau\right]$$

$$\iff \phi_t - q_t^{B,s} = \int_s^t \mu e^{-\mu(\tau-t)}(1-\phi_\tau)(q_t^{B,s} - q_t^{B,\tau})d\tau. \tag{5}$$

We establish existence by the intermediate value theorem. At $t = s$, the left-hand side is $\phi_s - 0 > 0$, while the right-hand side vanishes. As $t \to \infty$, the left-hand side tends to $p^* - p^* = 0$. Meanwhile, for fixed $\varepsilon > 0$, the right-hand side is bounded below by

$$\int_{t-\varepsilon}^t \mu e^{-\mu(\tau-t)}(1-\phi_\tau)(q_t^{B,s} - q_t^{B,\tau})d\tau \geq \int_{t-\varepsilon}^t \mu e^{-\mu\varepsilon}(1-\phi_\tau)(q_t^{B,s} - q_t^B(t-\varepsilon))d\tau$$

As $t \to \infty$, we have $\phi_t \to p^* < 1$ under the assumption that $\lambda_1 > 0$, and we also have $q_t^{B,s} \to p^* > 0$ under the assumption that $\lambda_0 > 0$. In addition, $q_t^{B,t-\varepsilon} < p^*(\lambda_0 + \lambda_1)\varepsilon$. Hence, for small ε and large t the integrand is strictly positive, and we conclude that the right-hand side of (5) is positive for large t. Since both the left- and right-hand side of (5) are continuous in t, there exists a solution $t^\dagger(s) \in (s, \infty)$ (and no solution at the ends of this interval).

Uniqueness follows from showing that $p_{t,s}^\dagger$ crosses $q_t^{B,s}$ from above at $t^\dagger(s)$. Given Markov switching, $\dot{q}_t^{B,s} = \lambda_0(1 - q_t^{B,s}) - \lambda_1 q_t^{B,s}$. For fixed s, we can write the evolution of $p_{t,s}^\dagger$ as

$$\dot{p}_{t,s}^\dagger = \lambda_0(1-p_{t,s}^\dagger) - \lambda_1 p_{t,s}^\dagger - \frac{\mu e^{-\mu t}\phi_t}{e^{-\mu t} + \int_s^t \mu e^{-\mu\tau}(1-\phi_\tau)\,d\tau}(1-p_{t,s}^\dagger). \tag{6}$$

At $t^\dagger(s)$, $p_{t^\dagger(s),s}^\dagger = q_{t^\dagger(s)}^{B,s}$, so

$$\dot{q}_{t^\dagger(s)}^{B,s} - \dot{p}_{t^\dagger(s),s}^\dagger = \frac{\mu e^{-\mu t^\dagger(s)}\phi_{t^\dagger(s)}}{e^{-\mu t^\dagger(s)} + \int_s^{t^\dagger(s)} \mu e^{-\mu\tau}(1-\phi_\tau)\,d\tau}(1-p_{t,s}^\dagger) > 0.$$

The result follows. □

A.2 Proof of Proposition 1

By definition, $t^\dagger(s) = s + a^*(s)$. Recall the equilibrium condition $q_{s+a^*(s)}^{B,s} = p_{s+a^*(s),s}^\dagger$. The equilibrium belief following non-disclosure by type t, denoted p_t, is given by

$$p_t = \begin{cases} p_{t,0}^\dagger & \text{if } t \leq a^*(0), \\ p_{s+a^*(s),\,s}^\dagger & \text{if } t \geq a^*(0), \end{cases}$$

where in the second case, $s + a^*(s) = t$.

We first show $a^*(s)$ is differentiable in s. Define

$$H(s,a) := q^{B,s}_{s+a} - p^{\dagger}_{s+a,s}.$$

By Lemma 11, $a^*(s)$ is the unique solution to $H(s,a) = 0$ for each s. Note that $H(s,a)$ is continuously differentiable in s and a because all functions $\phi, q^{B,\cdot}_{\cdot}$, and the exponential functions are smooth. so

$$\frac{\partial H(s,a)}{\partial a} = \left[\frac{\partial q^{B,s}_t}{\partial t} - \frac{\partial p^{\dagger}_{t,s}}{\partial t}\right]\Bigg|_{t=s+a}.$$

As is shown in the proof of Lemma 11, at $a = a^*(s)$,

$$\frac{\partial H(s,a)}{\partial a} = \frac{\mu e^{-\mu t}\phi_t}{e^{-\mu t} + \int_s^t \mu e^{-\mu\tau}(1-\phi_\tau)d\tau}(1 - q^{B,s}_t)\Bigg|_{t=s+a^*(s)} > 0.$$

By the implicit function theorem, $a^*(s)$ is continuously differentiable at each $s \geq 0$.

Differentiate the equilibrium condition $q^{B,s}_{s+a^*(s)} = p^{\dagger}_{s+a^*(s),s}$ implicitly in s, we get

$$\frac{dq^{B,s}_t}{ds}\Bigg|_{t=s+a^*(s)} = \left[\frac{\partial q^{B,s}_t}{\partial t}\left(1 + \frac{da^*(s)}{ds}\right) + \frac{\partial q^{B,s}_t}{\partial s}\right]\Bigg|_{t=s+a^*(s)}$$

and

$$\frac{dp^{\dagger}_{t,s}}{ds}\Bigg|_{t=s+a^*(s)} = \left[\frac{\partial p^{\dagger}_{t,s}}{\partial t}\left(1 + \frac{da^*(s)}{ds}\right) + \frac{\partial p^{\dagger}_{t,s}}{\partial s}\right]\Bigg|_{t=s+a^*(s)}.$$

Because $\partial q^{B,s}_t/\partial s = -\partial q^{B,s}_t/\partial t$, so

$$\frac{dq^{B,s}_t}{ds}\Bigg|_{t=s+a^*(s)} = \frac{\partial q^{B,s}_t}{\partial t}\frac{da^*(s)}{ds}\Bigg|_{t=s+a^*(s)}. \tag{7}$$

Moreover, at $t = s + a^*(s)$, $p^{\dagger}_{t,s} = q^{B,s}_t$, so

$$\frac{\partial p^{\dagger}_{t,s}}{\partial s}\Bigg|_{t=s+a^*(s)} = \frac{e^{-\mu s}\mu(1-\phi_s)(p^{\dagger}_{t,s} - q^{B,s}_t)}{e^{-\mu t} + \int_s^t \mu e^{-\mu\tau}(1-\phi_\tau)d\tau}\Bigg|_{t=s+a^*(s)} = 0.$$

Hence,

$$\frac{dp^{\dagger}_{t,s}}{ds}\Bigg|_{t=s+a^*(s)} = \frac{\partial p^{\dagger}_{t,s}}{\partial t}\left(1 + \frac{da^*(s)}{ds}\right)\Bigg|_{t=s+a^*(s)}. \tag{8}$$

In equilibrium, (7)=(8), which implies

$$\frac{da^*(s)}{ds} = \left.\frac{\partial p^\dagger_{t,s}/\partial t}{\partial q^{B,s}_t/\partial t - \partial p^\dagger_{t,s}/\partial t}\right|_{t=s+a^*(s)} \text{ and } 1 + \frac{da^*(s)}{ds} = \left.\frac{\partial q^{B,s}_t/\partial t}{\partial q^{B,s}_t/\partial t - \partial p^\dagger_{t,s}/\partial t}\right|_{t=s+a^*(s)}. \quad (9)$$

We have shown that $\partial q^{B,s}_t/\partial t - \partial p^\dagger_{t,s}/\partial t > 0$, so $1 + da^*(s)/ds > 0$. That is, $s + a^*(s)$ is increasing in s for $s \geq 0$ and strictly increasing if $s > 0$. This proves (i).

From (9), $\frac{da^*(s)}{ds}$ has the same sign as $\partial p^\dagger_{t,s}/\partial t|_{t=s+a^*(s)}$. In what follows, we derive necessary and sufficient conditions under which $\partial p^\dagger_{t,s}/\partial t|_{t=s+a^*(s)}$ is (strictly) positive.

Let $t = s + a$ with $a \geq 0$ fixed. Then

$$\frac{dp^\dagger_{s+a,s}}{ds} = \left.\frac{\partial p^\dagger_{t,s}}{\partial t}\right|_{t=s+a} + \left.\frac{\partial p^\dagger_{t,s}}{\partial s}\right|_{t=s+a}.$$

Evaluating at the equilibrium $a = a^*(s)$. As shown before, $\frac{\partial p^\dagger_{t,s}}{\partial s}|_{t=s+a^*(s)} = 0$. So at $a = a^*(s)$,

$$\left.\frac{dp^\dagger_{s+a,s}}{ds}\right|_{a=a^*(s)} = \left.\frac{\partial p^\dagger_{t,s}}{\partial t}\right|_{t=s+a^*(s)}.$$

Then, it suffices to derive necessary and sufficient conditions under which $dp^\dagger_{s+a,s}/ds|_{t=s+a^*(s)}$ is (strictly) positive for any $s \geq 0$. By definition,

$$p^\dagger_{s+a,s} = \frac{e^{-\mu(s+a)}\phi_{s+a} + \int_s^{s+a} \mu e^{-\mu\tau}(1-\phi_\tau) q^{B,\tau}_{s+a} d\tau}{e^{-\mu(s+a)} + \int_s^{s+a} \mu e^{-\mu\tau}(1-\phi_\tau) d\tau}.$$

Via a change of variable, let $\tau = s + v$ where $v \in [0, a]$. Then $d\tau = dv$ (since s is fixed). After some simplifying, we get

$$p^\dagger_{s+a,s} = \frac{e^{-\mu a}\phi_{s+a} + \int_0^a \mu e^{-\mu v}(1-\phi_{s+v}) q^{B,s+v}_{s+a} dv}{e^{-\mu a} + \int_0^a \mu e^{-\mu v}(1-\phi_{s+v}) dv}. \quad (10)$$

Note that $q^{B,s+v}_{s+a} = p^*\left(1 - e^{-(\lambda_0+\lambda_1)(a-v)}\right)$ is independent of s. This means $p^\dagger_{s+a,s}$ changes with s only through its changes in $\phi_\cdot$. Differentiate both sides with respect to s and note that $d\phi_{s+x}/ds = d\phi_t/dt|_{t=s+x} =: \dot{\phi}_{s+x}$. After simplifying,

$$\frac{dp^\dagger_{s+a,s}}{ds} = \frac{e^{-\mu a}\dot{\phi}_{s+a} + \int_0^a \mu e^{-\mu v}\dot{\phi}_{s+v}\left(p^\dagger_{s+a,s} - q^{B,s+v}_{s+a}\right) dv}{e^{-\mu a} + \int_0^a \mu e^{-\mu v}(1-\phi_{s+v}) dv}.$$

Evaluate at $a = a^*(s)$. Because the denominator is positive, $\frac{dp^\dagger_{s+a,s}}{ds}\Big|_{a=a^*(s)}$ has the same sign as the numerator:

$$e^{-\mu a^*(s)}\dot{\phi}_{s+a^*(s)} + \int_0^{a^*(s)} \mu e^{-\mu v}\dot{\phi}_{s+v}\left(p^\dagger_{s+a^*(s),s} - q^{B,s+v}_{s+a^*(s)}\right) dv.$$

Because $p^\dagger_{s+a^*(s),s} = q^{B,s}_{s+a^*(s)}$ so $p^\dagger_{s+a^*(s),s} - q^{B,s+v}_{s+a^*(s)} = q^{B,s}_{s+a^*(s)} - q^{B,s+v}_{s+a^*(s)} \geq 0$ for all $v \in [0, a^*(s)]$ with equality only at $v = 0$. Moreover, $\dot{\phi}_t = -(\lambda_0+\lambda_1)(p_0 - p^*)e^{-(\lambda_0+\lambda_1)t}$ for all t, so $\dot{\phi}_t > 0$ if and only if $p_0 < p^*$, and $\dot{\phi}_{s+a^*(s)}$ and $\dot{\phi}_{s+v}$ have the same sign. Thus, the numerator is positive if and only if $p_0 < p^*$.

Putting all arguments together, $\frac{da^*(s)}{ds}$ has the same sign as $\frac{\partial p^\dagger_{t,s}}{\partial t}|_{t=s+a^*(s)}$ which is equal to $\frac{dp^\dagger_{s+a,s}}{ds}|_{a=a^*(s)}$. The above argument shows $\frac{dp^\dagger_{s+a,s}}{ds}|_{a=a^*(s)}$ is positive if and only if $p_0 < p^*$. This proves (ii).

To prove (iii), first note that by Lemma 11 and (ii), $a^*(s)$ is monotone and bounded. By the monotone convergence theorem, $a^*(s)$ converges. In what follows, we first derive α^* and then show $a^*(s)$ converges to α^*.

Let $p_0 = p^*$. Then $\phi_{s+a} = p^*$ and

$$p^\dagger_{s+a,s} = \frac{e^{-\mu a}p^* + (1-p^*)p^* \int_0^a \mu e^{-\mu v} q^{B,s+v}_{s+a} dv}{e^{-\mu a} + (1-p^*)\int_0^a \mu e^{-\mu v} dv},$$

which is independent of s. Moreover, $q^{B,s}_{s+a} = p^*\left(1 - e^{-(\lambda_0+\lambda_1)a}\right)$, so $H(s,a) = q^{B,s}_{s+a} - p^\dagger_{s+a,s}$ is also independent of s. To emphasize this independence, let $H^*(a) := H(s,a)|_{p_0=p^*}$. By Lemma 11, there exists a unique finite solution to $H^*(a) = 0$, denote it by α^*. In words, α^* is the equilibrium delay for all s when $p_0 = p^*$. Note that $H(s,0) = 0 - p^* < 0$ and because $H(s,a)$ is continuous in a, $\alpha^* > 0$.

In what follows, we show that for any p_0 and fixed $0 < a < \infty$, $\lim_{s\to\infty} H(s,a) \to H^*(a)$.

If $p_0 = p^*$, $H(s,a) = H^*(a)$ for all s by definition. Take any $p_0 \neq p^*$ and fix any $0 < a < \infty$. Take the limit as $s \to \infty$, recall that $q^{B,s}_{s+a} = p^*\left(1 - e^{-(\lambda_0+\lambda_1)a}\right)$ is independent of s. In $p^\dagger_{s+a,s}$, which is defined by equation (10), $\lim_{s\to\infty}\phi_{s+a} = p^*$ by definition. For the integral term, denote the integrand by $I(s,v) = \mu e^{-\mu v}(1-\phi_{s+v})q^{B,s+v}_{s+a}$, which converges pointwise to $I^*(v) := \lim_{s\to\infty} I(s,v) = \mu e^{-\mu v}(1-p^*)q^{B,s+v}_{s+a}$. Note that

$$|I(s,v) - I^*(v)| = \mu e^{-\mu v}\left|p^* - \phi_{s+v}\right| q^{B,s+v}_{s+a} \leq \mu\left|p_0 - p^*\right| e^{-(\lambda_0+\lambda_1)s}p^*$$

where the inequality follows from $|p^* - \phi_{s+v}| = |p_0 - p^*|\, e^{-(\lambda_0+\lambda_1)(s+v)} \leq |p_0 - p^*|\, e^{-(\lambda_0+\lambda_1)s}$, $e^{-\mu v} \leq 1$, and $q^{B,s+v}_{s+a} \leq p^*$. Note that the right-hand side is independent of v and converges to 0 as $s \to \infty$. Formally, $\sup_{v\in[0,a]} |I(s,v) - I^*(v)| \leq \mu\left|p^* - \phi_{s+v}\right| e^{-(\lambda_0+\lambda_1)s}p^* \to 0$, which means $I(s,v)$ converges uniformly to $I^*(v)$, and

$$\lim_{s\to\infty}\int_0^a \mu e^{-\mu v}(1-\phi_{s+v})q_{s+a}^{B,s+v}dv = \int_0^a \lim_{s\to\infty}\mu e^{-\mu v}(1-\phi_{s+v})q_{s+a}^{B,s+v}dv$$
$$= \int_0^a \mu e^{-\mu v}(1-p^*)q_{s+a}^{B,s+v}dv.$$

The same argument applies to the denominator: $\lim_{s\to\infty}\int_0^a \mu e^{-\mu v}(1-\phi_{s+v})dv = \int_0^a \mu e^{-\mu v}(1-p^*)dv > 0$. Therefore, $p^\dagger_{s+a,s} \to p^\dagger_{s+a,s}|_{p_0=p^*}$ and so $\lim_{s\to\infty} H(s,a) \to H^*(a)$.

Next, fix any $\varepsilon \in (0,\alpha^*)$. By Lemma 11, $H^*(a)$ has a unique zero at α^* and crosses from negative to positive. Hence

$$H^*(\alpha^*-\varepsilon) < 0 < H^*(\alpha^*+\varepsilon).$$

Since $H(s,a) \to H^*(a)$ for each fixed a, there exists S such that for all $s \geq S$,

$$H(s,\alpha^*-\varepsilon) < 0 < H(s,\alpha^*+\varepsilon).$$

Again by Lemma 11, $a^*(s)$ is the unique zero of $H(s,\cdot)$ for each s. Moreover, the crossing is from negative to positive, so $H(s,a) < 0$ for $a < a^*(s)$ and $H(s,a) > 0$ for $a > a^*(s)$. Therefore, for all $s \geq S$,

$$\alpha^* - \varepsilon < a^*(s) < \alpha^* + \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, it follows that $\lim_{s\to\infty} a^*(s) = \alpha^*$.

A.3 Proof of Corollary 1

By Proposition 1(i), $s + a^*(s)$ is increasing in s, which implies that for each $t \geq a^*(0)$, there is a unique $\tau(t) \geq 0$ satisfying $\tau(t) + a^*(\tau(t)) = t$ for each $t \geq a^*(0)$. Moreover $\tau(t)$ is increasing in t. The equilibrium condition at each t says

$$p^\dagger_{t,s} = q_t^{B,\tau(t)} = p^*\left(1 - e^{-(\lambda_0+\lambda_1)(t-\tau(t))}\right) = p^*\left(1 - e^{-(\lambda_0+\lambda_1)a^*(\tau(t))}\right).$$

The right-hand side is strictly increasing in $a^*(\tau(t))$. Because $\tau(t)$ is increasing in t, $p^\dagger_{t,s}$ therefore inherits the monotonicity of $a^*(s)$. Proposition 1 (ii) then implies that $p^\dagger_{t,s}$ is strictly increasing when $p_0 < p^*$, strictly decreasing when $p_0 > p^*$, and constant when $p_0 = p^*$.

Finally, as $t \to \infty$, we have $\tau(t) \to \infty$. By Proposition 1 (iii), $a^*(s) \to \alpha^*$, $\lim_{t\to\infty} p^\dagger_{t,s} = p^*\left(1 - e^{-(\lambda_0+\lambda_1)\alpha^*}\right) < p^*$, where the inequality follows because α^* is finite.

A.4 Proof of Proposition 2

For each $t \geq a^*(0)$, define $\tau(t)$ implicitly by $\tau(t) + a^*(\tau(t)) = t$; that is, $\tau(t)$ is the most recent timestamp for bad evidence that is disclosed by date t. For $t < a^*(0)$, set $\tau(t) = -\infty$.

Now fix any $t \geq 0$. Let ν_s^G and ν_s^B denote the probability in equilibrium that good evidence and bad evidence, respectively, obtained at time s is disclosed by time t. Fix an arbitrary alternative disclosure strategy and let $\tilde{\nu}_s^G$ and $\tilde{\nu}_s^B$ denote the analogous probabilities. Note that $\nu_s^G = 1$ for all $s \in [0, t]$, and $\nu_s^B = 1$ for all $s \leq \tau(t)$ and $\nu_s^B = 0$ for all $s > \tau(t)$. Let p_t and $\tilde{p}_t$ denote the beliefs conditional on no disclosure in the equilibrium and under the alternative disclosure strategy, respectively. Finally, let N_t and $\tilde{N}_t$ denote the probability of no disclosure under the equilibrium and alternative strategies, respectively.

By construction, we have (i) $p_t \leq q_t^{G,s}$ for all $s \in [0, t]$, (ii) $p_t \leq q_t^{B,s}$ for $s \in [0, \tau(t)]$, and (iii) $q_t^{B,s} \leq p_t$ for $s \in (\tau(t), t]$.

By the law of iterated expectations,

$$\phi_t = N_t p_t + \int_0^t \phi_s \mu e^{-\mu s} \nu_s^G q_t^{G,s}\, ds + \int_0^t (1 - \phi_s) \mu e^{-\mu s} \nu_s^B q_t^{B,s}\, ds \tag{11}$$

and also

$$\phi_t = \tilde{N}_t \tilde{p}_t + \int_0^t \phi_s \mu e^{-\mu s} \tilde{\nu}_s^G q_t^{G,s}\, ds + \int_0^t (1 - \phi_s) \mu e^{-\mu s} \tilde{\nu}_s^B q_t^{B,s}\, ds, \tag{12}$$

where the belief at time t after disclosure (at any time) of good or bad evidence with timestamp s, namely $q_t^{G,s}$ or $q_t^{B,s}$, is independent of the conjectured strategy. Equating the right hand sides of (11) and (12) and subtracting p_t from both yields

$$\begin{aligned}
&0 + \int_0^t \phi_s \mu e^{-\mu s} \nu_s^G (q_t^{G,s} - p_t)\, ds + \int_0^t (1 - \phi_s) \mu e^{-\mu s} \nu_s^B (q_t^{B,s} - p_t)\, ds \\
&= \tilde{N}_t (\tilde{p}_t - p_t) + \int_0^t \phi_s \mu e^{-\mu s} \tilde{\nu}_s^G (q_t^{G,s} - p_t)\, ds + \int_0^t (1 - \phi_s) \mu e^{-\mu s} \tilde{\nu}_s^B (q_t^{B,s} - p_t)\, ds \\
\iff & \tilde{N}_t (\tilde{p}_t - p_t) = \int_0^t \phi_s \mu e^{-\mu s} (\nu_s^G - \tilde{\nu}_s^G)(q_t^{G,s} - p_t)\, ds \\
&\qquad + \int_0^t (1 - \phi_s) \mu e^{-\mu s} (\nu_s^B - \tilde{\nu}_s^B)(q_t^{B,s} - p_t)\, ds.
\end{aligned} \tag{13}$$

Now for all $s \in [0, t]$, we have $1 = \nu_s^G \geq \tilde{\nu}_s^G$ and $q_t^{G,s} - p_t \geq 0$, so the first integral on the right hand side of (13) is nonnegative. We then turn to the integrand of the second integral. For $s \in [0, \tau(t)]$, we have $1 = \nu_s^B \geq \tilde{\nu}_s^B$ and $q_t^{B,s} \geq p_t$, so the integrand is nonnegative. And for $s \in (\tau(t), t]$, we have $0 = \nu_s^B \leq \tilde{\nu}_s^B$ and $q_t^{B,s} \leq p_t$, so again the integrand is nonnegative.

Hence, the right hand side of (13) is nonnegative. As $\tilde{N}_t \geq e^{-\mu t} > 0$, we must have $\tilde{p}_t \geq p_t$.

A.5 Proof of Proposition 3

Suppose there exists such an equilibrium, and suppose the sender obtains good evidence at time 0. Let $F(T)$ denote the sender's discounted payoff from disclosing at time T. By disclosing immediately, the sender's payoff is $F(0) = \int_0^\infty e^{-rt}\left[p^* + (1-p^*)e^{-(\lambda_0+\lambda_1)t}\right] dt = \frac{\lambda_0+r}{r(\lambda_0+\lambda_1+r)}$. By instead disclosing at some time $\varepsilon > 0$, the sender's payoff is $F(\varepsilon) = \int_0^\varepsilon e^{-rt}p_t\, dt + \int_\varepsilon^\infty e^{-rt}\left[p^* + (1-p^*)e^{-(\lambda_0+\lambda_1)(t-\varepsilon)}\right] dt = \int_0^\varepsilon e^{-rt}p_t\, dt + e^{-r\varepsilon}F(0)$. If $p_0 > \frac{\lambda_0+r}{\lambda_0+\lambda_1+r}$, then since p_t is right-continuous, we have for sufficiently small $\varepsilon > 0$ that

$$\int_0^\varepsilon e^{-rt}p_t\, dt > \int_0^\varepsilon e^{-rt}\left(\frac{\lambda_0+r}{\lambda_0+\lambda_1+r}\right) dt = \frac{1-e^{-r\varepsilon}}{r}\frac{\lambda_0+r}{\lambda_0+\lambda_1+r} = (1-e^{-r\varepsilon})F(0).$$

Hence, $F(\varepsilon) > (1-e^{-r\varepsilon})F(0) + e^{-r\varepsilon}F(0) = F(0)$. Thus, immediate disclosure is not optimal at time 0.

A.6 Proof of Theorem 2

A.6.1 Setup and Laws of Motion

We use the notation and laws of motion introduced in Section 4.2. In particular, G_i denotes the joint probability that there is undisclosed good evidence and the time-t state is i; $S_t = G_{0,t} + G_{1,t}$ denotes the total stock of undisclosed good evidence; $q_t = G_{1,t}/S_t$ is the receiver's belief after disclosure of good evidence; and p_t is the no-disclosure belief. During a stockpiling phase, $\kappa_t = 0$.

During either the stockpiling or purging phase, (p, G_1, G_0) have the following laws of motion:

$$\dot{G}_{1,t} = -G_{1,t}\lambda_1 + (p_t - G_{1,t})\mu - G_{1,t}\kappa_t(1 - G_{1,t} - G_{0,t}) \tag{14}$$
$$\dot{G}_{0,t} = G_{1,t}\lambda_1 - G_{0,t}\kappa_t(1 - G_{1,t} - G_{0,t}) \tag{15}$$
$$\dot{p}_t = -\lambda_1 p_t - G_{1,t}\kappa_t + p_t(G_{0,t} + G_{1,t})\kappa_t, \tag{16}$$

We use the change of variables $S = G_0 + G_1$ and $q = G_1/S$ and differentiate to obtain

$$\dot{S}_t = (p_t - S_t q_t)\mu - S_t\kappa_t(1 - S_t) \tag{17}$$
$$\dot{q}_t = -\lambda_1 q_t + (1 - q_t)\mu\frac{p_t - S_t q_t}{S_t} \tag{18}$$

During the transparency phase, we have $q_t \equiv 1$ and p evolves according to

$$\dot{p}_t = -\lambda_1 p_t - \mu p_t(1 - p_t). \tag{19}$$

Note that p is decreasing during the transparency phase.

A.6.2 Local incentive compatibility

We will make heavy use of local incentive compatibility. Suppose the sender possesses good evidence at time t_0 and plans to disclose it at time $t \geq t_0$. Then for $s \geq t$, we have $p_s = q_t e^{-\lambda_1(s-t)}$. The sender's payoff is

$$F(t; t_0) := \int_{t_0}^{t} e^{-r(s-t_0)} p_s \, ds + \int_{t}^{\infty} e^{-r(s-t_0)} q_t e^{-\lambda_1(s-t)} \, ds.$$

The agent is indifferent disclosing at time t_0 and time $t_0 + \epsilon$ for all sufficiently small ϵ if and only if $F_t(t; t_0)|_{t=t_0} = 0$; equivalently,

$$0 = p_{t_0} - q_{t_0} + \int_{t_0}^{\infty} (\lambda_1 q_{t_0} + \dot{q}_{t_0}) e^{-r(s-t_0)} e^{-\lambda_1(s-t_0)} \, ds = p_{t_0} - q_{t_0} + (\lambda_1 q_{t_0} + \dot{q}_{t_0}) \frac{1}{\lambda_1 + r}.$$

Rearranging and replacing t_0 with t yields the local indifference condition

$$\dot{q}_t = r q_t - (\lambda_1 + r) p_t. \tag{20}$$

If $\dot{q}_t > r q_t - (\lambda_1 + r) p_t$, then for small enough ϵ, the sender strictly prefers to disclose at $t + \epsilon$ rather than at time t. If $\dot{q}_t < r q_t - (\lambda_1 + r) p_t$, then the agent strictly prefers to disclose at time t than at time $t + \epsilon$ for all sufficiently small ϵ.

A.6.3 Identifying $\underline{p}$ and $\overline{p}$

In this section we prove the following lemma.

Lemma 12. *Any three-phase equilibrium must start in the stockpiling phase if $p_0 > \overline{p}$, the transparency phase if $p_0 < \underline{p}$, and the purging phase if $p_0 \in (\underline{p}, \overline{p})$, where $\overline{p} := \frac{r + \lambda_1/2}{\lambda_1 + r}$ and $\underline{p} := \frac{r}{\lambda_1 + r}$.*

Proof. By definition, any three-phase equilibrium must start in one of the three phases. We consider each possibility separately and derive necessary conditions on p_0. We then conclude the proof by observing that the conditions are mutually exclusive, so that p_0 determines the starting phase as in the proposition.

Starting in stockpiling phase: We argue that if the game starts in the stockpiling phase, then a sender who obtains good evidence at time 0 must be willing to wait an instant, and this implies $p_0 > \overline{p} = \frac{r+\frac{\lambda_1}{2}}{\lambda_1+r}$.

A necessary and sufficient condition to start in the stockpiling phase is that for some time $\underline{T} > 0$, waiting until $\underline{T}$ to disclose must be weakly preferred to disclosing at time t for a sender who obtains good evidence at time $t \in [0, \underline{T})$. Hence, recalling (20), a necessary condition for starting in stockpiling is

$$\dot{q}_{\underline{T}} \geq r q_{\underline{T}} - (\lambda_1 + r) p_{\underline{T}} \tag{21}$$

for some $\underline{T} > 0$.

Using $\kappa_t \equiv 0$, the laws of motion for beliefs admit a (unique) closed-form solution:

$$G_{1,t} = e^{-\lambda_1 t} p_0 (1 - e^{-\mu t}) \tag{22}$$

$$G_{0,t} = e^{-\lambda_1 t} p_0 \left(\frac{\lambda_1 e^{-t\mu} + \mu e^{\lambda_1 t}}{\mu + \lambda_1} - 1 \right) \tag{23}$$

$$p_t = e^{-\lambda_1 t} p_0 \tag{24}$$

$$q_t = \frac{G_{1,t}}{G_{0,t} + G_{1,t}} = \frac{(e^{\mu t} - 1)(\lambda_1 + \mu)}{(e^{(\lambda_1+\mu)t} - 1)\mu}. \tag{25}$$

Using (24), we get the necessary condition that for some $t > 0$,

$$p_0 \geq b(t) := \frac{e^{\lambda_1 t}(r q_t - \dot{q}_t)}{\lambda_1 + r},$$

where $b(t)$ is defined using (25). We claim that $b'(t) > 0$ for all $t > 0$. Direct calculation shows that

$$b'(t) = \frac{e^{\lambda_1 t}(\lambda_1 + \mu)}{(e^{(\lambda_1+\mu)t} - 1)^3 (\lambda_1 + r)\mu} P(t), \qquad \text{where}$$
$$P(t) := \mu(\lambda_1 + \mu + r)\, e^{t(2\lambda_1+2\mu)} - (\lambda_1 + \mu)(2\lambda_1 + \mu + r)\, e^{t(\lambda_1+2\mu)}$$
$$+ (2\lambda_1^2 + 3\lambda_1\mu + \lambda_1 r + \mu^2 - \mu r)\, e^{t(\lambda_1+\mu)} - (\lambda_1 + \mu)(\mu - r)\, e^{t\mu} - \lambda_1 r].$$

Note that for both $\mu \geq r$ and $\mu < r$, $P(t)$ has three sign changes and thus at most three real roots. Moreover, it is easy to check that there is a triple root at $t = 0$. Hence, there are no other real roots. Also, since the leading term is positive, $P(t)$ is positive for sufficiently large t. We conclude that for all $t > 0$, we have $P(t) > 0$ and therefore $b'(t) > 0$. It follows that $p_0 \geq b(t)$ for some $t > 0$ only if

$$p_0 > b(0) = \frac{r \cdot q_0 - \dot{q}_0}{\lambda_1 + r} = \frac{r \cdot 1 + \frac{\lambda_1}{2}}{\lambda_1 + r} = \overline{p}.$$

This is exactly the inequality asserted in Lemma 12.

For later use, let us define $c(t) = \dot{q}_t - rq_t + (\lambda_1 + r)p_t$. Assume $p_0 > \overline{p}$. We show that c crosses 0 from above at the unique point t where $p_0 = b(t)$. We have

$$\begin{aligned} c(t) &= (\lambda_1 + r)e^{-\lambda_1 t}(p_0 - b(t)) \\ \implies c'(t) &= -\lambda_1 c(t) - b'(t)(\lambda_1 + r)e^{-\lambda_1 t}. \end{aligned}$$

When $c(t) = 0$, the right hand side has the same sign as $-b'(t) < 0$, as desired.

Starting in transparency phase:

If it starts in the transparency phase, then we have $q_0 \equiv 1$ and $\dot{q}_0 \equiv 0$, so at time 0 the left hand side of (20) is 0 and the right hand side is $r - (\lambda_1 + r)p_0$. Now for immediate disclosure to be optimal, we must have $0 \leq r - (\lambda_1 + r)p_0$, or equivalently, $p_0 \leq \frac{r}{\lambda_1 + r} = \underline{p}$.

Starting in purging phase: The sender must be indifferent between disclosing good evidence at time 0 and disclosing at time t for all sufficiently small $t > 0$, the local IC constraint (20) must hold for all $t \in [0, \epsilon)$ for sufficiently small $\epsilon > 0$.

At $t = 0$, we have $q_0 = 1$ and $\dot{q}_0 \leq 0$ as q is bounded above by 1. This implies $p_0 \geq \frac{r}{\lambda_1 + r}$. Moreover, we can rule out $p_0 = \frac{r}{\lambda_1 + r}$ as follows. Note that this would require $\dot{q}_0 = 0$. Since q is bounded above by 1, we have $\ddot{q}_0 \leq 0$. If indifference holds for all $t \in [0, \varepsilon)$, then $\ddot{q}_0 = r\dot{q}_0 - (\lambda_1 + r)\dot{p}_0 = 0 - (\lambda_1 + r)\dot{p}_0 \geq (\lambda_1 + r)\lambda_1 p_0 > 0$, where the bound $\dot{p}_0 \leq -\lambda_1 p_0$ follows from $p_t \leq \phi_t = \phi_0 e^{-\lambda_1 t}$. This contradicts $\ddot{q}_0 \leq 0$, so we conclude that the game can begin with a purging phase only if $p_0 > \frac{r}{\lambda_1 + r}$.

We argue that $\dot{q}_0 \geq -\frac{\lambda_1}{2}$. Let $X_t = t - \tau$ be the age of evidence that is disclosed at time t. Then

$$\dot{q}_0 = -\lambda_1 \lim_{t \to 0} \frac{\mathbb{E}[X_t]}{t}.$$

Expanding and writing $\gamma_s := \phi_s \mu e^{-\mu s}$ yields

$$\lim_{t \to 0} \frac{\mathbb{E}[X_t]}{t} = \lim_{t \to 0} \frac{\int_0^t (t - s)\gamma_s e^{-\int_s^t \kappa_u \, du} \, ds}{t \int_0^t \gamma_s e^{-\int_s^t \kappa_u \, du} \, ds} = \lim_{t \to 0} \frac{\int_0^t (t - s) e^{-\int_s^t \kappa_u \, du} \, ds}{t \int_0^t e^{-\int_s^t \kappa_u \, du} \, ds},$$

where we have used that γ is continuous. Since the density $\frac{e^{-\int_s^t \kappa_u\,du}}{\int_0^t e^{-\int_x^t \kappa_u\,du}\,dx}$ is nondecreasing in s, while the function $t-s$ is decreasing, the weighted average is bounded above by replacing the density with a uniform density. Thus,

$$\lim_{t\to 0}\frac{\mathbb{E}[X_t]}{t} \le \lim_{t\to 0}\frac{\int_0^t (t-s)\,ds}{t\int_0^t\,ds} = \lim_{t\to 0}\frac{\frac{t^2}{2}}{t^2} = \frac{1}{2}.$$

We conclude that $\dot{q}_0 \ge -\frac{\lambda_1}{2}$, and thus

$$p_0 = \frac{r-\dot{q}_0}{\lambda_1+r} \le \frac{r+\frac{\lambda_1}{2}}{\lambda_1+r}.$$

Summary We have shown that starting in the stockpiling phase requires $p_0 > \overline{p}$, the purging phase $p_0 \in (\underline{p}, \overline{p}]$, and the transparency phase $p_0 \le \underline{p}$. Since a three-phase equilibrium must start in one of these three phases, the only possibility is stockpiling for $p_0 > \overline{p}$, transparency for $p_0 < \underline{p}$, and purging for $p_0 \in (\underline{p}, \overline{p})$. □

We now establish existence and uniqueness within the class of three-phase equilibria for p_0 in each of the three intervals.

A.6.4 Low p_0.

Suppose $p_0 \le \frac{r}{\lambda_1+r}$ and that the receiver conjectures that bad evidence is never disclosed, while good evidence is disclosed immediately.

Then $q_t \equiv 1$ and p evolves according to (19). Note that because p is strictly decreasing, if $p_0 \le \underline{p}$ then this inequality remains true at all future times. Hence, immediate disclosure of evidence is optimal at all histories, including after off-path delay of disclosure.

A.6.5 High p_0.

Suppose that $p_0 > \frac{r+\lambda_1/2}{r+\lambda_1}$. Fix any candidate three-phase equilibrium. Since it must start in the stockpiling phase, the phases are characterized by cutoff times $0 < \underline{T}_0 < \overline{T}_0 < \infty$. Let $(\kappa_t)_{t\in[\underline{T}_0,\overline{T}_0)}$ be the disclosure intensity during the purging phase. We derive necessary conditions that uniquely determine these objects and then verify that the sender's strategy is optimal.

Let (p, q, S) denote the belief/stock variables in the candidate equilibrium. It is useful to define auxiliary variables (p^0, S^0, q^0) as the belief/stock variables under an alternative conjecture that good evidence and bad evidence are never disclosed (with q^0 determined by passive belief updating).

Since both conjectures prescribe no disclosure on $[0, \underline{T}_0)$, and the variables (p_t, q_t, S_t) and (p_t^0, q_t^0, S_t^0) are continuous at $\underline{T}_0$ (since there are no atoms of disclosure at $\underline{T}_0$ in either case), we have

$$(p_t, q_t, S_t) = (p_t^0, q_t^0, S_t^0) \qquad \text{for all } t \leq \underline{T}_0.$$

These have closed form expressions given by (22)-(25) where $q_0 = 1$ and $q_t = G_{1,t}/(G_{0,t}+G_{1,t})$ for $t > 0$.

Lemma 13. *We have $\underline{T}_0 = \underline{T}$, where $\underline{T}$ is the unique time at which local IC (20) holds under the no-disclosure beliefs:*

$$\dot{q}_t^0 = rq_t^0 - (\lambda_1 + r)p_t^0.$$

Proof. In the purging phase, q must satisfy local indifference,

$$\dot{q}_t = rq_t - (\lambda_1 + r)p_t \qquad \text{for } t \in [\underline{T}_0, \overline{T}_0), \tag{26}$$

where we use the right-derivative.

Also, by (18) at $t = \underline{T}_0$,

$$\dot{q}_{\underline{T}_0} = -\lambda_1 q_{\underline{T}_0} + (1 - q_{\underline{T}_0})\mu \frac{p_{\underline{T}_0} - S_{\underline{T}_0} q_{\underline{T}_0}}{S_{\underline{T}_0}}, \tag{27}$$

$$\dot{q}_{\underline{T}_0}^0 = -\lambda_1 q_{\underline{T}_0}^0 + (1 - q_{\underline{T}_0}^0)\mu \frac{p_{\underline{T}_0}^0 - S_{\underline{T}_0}^0 q_{\underline{T}_0}^0}{S_{\underline{T}_0}^0}. \tag{28}$$

Since $(p_{\underline{T}_0}, S_{\underline{T}_0}, q_{\underline{T}_0}) = (p_{\underline{T}_0}^0, S_{\underline{T}_0}^0, q_{\underline{T}_0}^0)$, (27) and (28) imply

$$\dot{q}_{\underline{T}_0} = \dot{q}_{\underline{T}_0}^0.$$

Combining with (26) and again using $(p_{\underline{T}_0}, S_{\underline{T}_0}, q_{\underline{T}_0}) = (p_{\underline{T}_0}^0, S_{\underline{T}_0}^0, q_{\underline{T}_0}^0)$ gives

$$\dot{q}_{\underline{T}_0}^0 = rq_{\underline{T}_0}^0 - (\lambda_1 + r)p_{\underline{T}_0}^0. \tag{29}$$

Recall that from the construction of $\overline{p}$ earlier, for $p > \overline{p} = b(0)$, there is a unique value of $\underline{T}_0 > 0$ at which $p_0 = b(\underline{T}_0)$, which is equivalent to (29). We denote this unique value by $\underline{T}$. Hence, in any three-phase equilibrium, $\underline{T}_0 = \underline{T}$; in other words, the stockpiling phase must have length $\underline{T}$. □

We now characterize the purging phase solution: κ_t and $\overline{T}$. The initial values for $(p_{\underline{T}}, G_{1,\underline{T}}, G_{0,\underline{T}}, q)$ are already given from the end of the stockpiling phase. In the purging phase, (p, G_1, G_0, q, κ) must satisfy the local indifference condition (20) and the laws of motion (14)-(16). In addition, q must also satisfy, by definition, $q_t = G_{1,t}/(G_{1,t} + G_{0,t})$ at each t in $[\underline{T}_0, \overline{T}_0)$.[12].

This system first reduces to the system in (p, q, S, κ) given by (16), (17), (18), and (20). In the next two steps, we eliminate S_t and κ_t.

Next, we solve for S_t by using the right side of (20) to replace $\dot{q}_t$ in (18):

$$\begin{aligned}
\dot{q}_t = -\lambda_1 q_t + (1 - q_t)\mu \frac{p_t - S_t q_t}{S_t} &= r q_t - (\lambda_1 + r)p_t \\
\implies S_t = S(p_t, q_t) &:= \frac{\mu p_t(1 - q_t)}{q_t[\lambda_1 + r + \mu(1 - q_t)] - (\lambda_1 + r)p_t}.
\end{aligned} \tag{30}$$

A sufficient condition for $S(p_t, q_t)$ to be well defined is that $p_t < q_t \leq 1$, which we will verify.

To solve for κ_t, differentiate (30), substitute in the expressions for $\dot{p}$ and $\dot{q}$ from (16) and (20), and equate the resulting expression to the right hand side of (17). Solving for κ_t then yields

$$\kappa_t = \kappa(p_t, q_t) := \frac{A_t(q_t[\lambda_1 + r + \mu(1 - q_t)] - (\lambda_1 + r)p_t)}{(q_t - p_t)(1 - q_t)[\lambda_1 + r + \mu(1 - q_t)]}, \qquad \text{where} \tag{31}$$

$$A_t := A(p_t, q_t) := r + \lambda_1 + \mu + (r - \mu)q_t - 2p_t(\lambda_1 + r). \tag{32}$$

As long as $0 < p_t < q_t < 1$, which we will verify in our solution, denominator will be positive, so $\kappa(p_t, q_t)$ will be well defined. We will also need to verify that $\kappa_t \geq 0$ at all times for it to be feasible as an arrival rate. Since $0 < p_t < q_t < 1$ implies the second factor in the numerator is positive, it will be sufficient to show that $A_t \geq 0$ for all $t \in [\underline{T}, \overline{T})$.

After eliminating S_t and κ_t in this way, we obtain the reduced IVP consisting of the pair of ODEs

$$\dot{p}_t = p_t \frac{2(\lambda_1 + r)\mu p_t + \mu q_t(\lambda_1 + \mu - r) - (\lambda_1 + \mu)(\lambda_1 + r + \mu)}{\lambda_1 + r + \mu(1 - q_t)} \tag{33}$$

$$\dot{q}_t = r q_t - (\lambda_1 + r)p_t, \tag{34}$$

with initial conditions $(p_{\underline{T}}, q_{\underline{T}})$ from the end of the stockpiling phase.

Note that the right hand sides of (33) and (34) are time-invariant, C^1 functions of (p_t, q_t) on the domain $U := \{(p, q) \in \mathbb{R}^2 : q < 1 + \frac{\lambda + r}{\mu}\}$. In particular, they are locally Lipschitz continuous, uniformly in time. By the Picard-Lindelöf theorem, there exists a unique solution to the IVP with initial conditions on some interval $[\underline{T}, \underline{T}+\varepsilon]$ with $\varepsilon > 0$. Let $[\underline{T}, T^{\max})$ denote the maximal interval of existence (and uniqueness), where $T^{\max}$ is possibly ∞.

[12] Since $\underline{T}_0 > 0$, $G_{1,\underline{T}_0}/(G_{1,\underline{T}_0} + G_{0,\underline{T}_0})$ is well defined.

Define $R \subset U$ by

$$R := \{(p, q) : 0 < p < q < 1 \text{ and } A(p, q) > 0\}.$$

Then on R, (i) κ_t is well defined via (31) and strictly positive, and (ii) S_t is well defined via (30). Define the (possibly infinite) time $T := \inf\{t \in [\underline{T}, T^{\max}) : (p_t, q_t) \notin R\}$.

Lemma 14. *We have that $(p_{\underline{T}}, q_{\underline{T}})$ lies in the interior of R. Hence, $T > \underline{T}$.*

Proof. In the stockpiling phase, we have

$$\dot{p}_t = -\lambda_1 p_t \tag{35}$$

$$\dot{q}_t = -\lambda_1 q_t + (1 - q_t)\mu \frac{p_t - S_t q_t}{S_t} \tag{36}$$

$$\dot{S}_t = (p_t - S_t q_t)\mu, \tag{37}$$

with initial conditions $(p_0, q_0, S_0) = (p_0, 1, 0)$. That p_t is strictly positive for all $t \in [0, \underline{T}]$ is standard; this follows from the comparison theorem and is also evident from the closed form solution.

It is useful to show that $S_t > 0$ until $\underline{T}$. We have $\dot{S}_0 > 0$, so $S_t > 0$ for all sufficiently small t. Now write (37) as $\dot{S}_t = f(t, S_t)$, and define $y_t \equiv 0$. Then $S_t > y_t$ and $\dot{S}_t - f(t, S_t) = 0 > -p_t\mu = y_t - f(t, y_t)$. By the comparison theorem implies $S_t > y_t = 0$ for all $t > 0$.

The comparison theorem also implies $q_t > 0$ for all $t \in [0, \underline{T}]$; moreover, $q_t < 1$ for all $t \in (0, \underline{T}]$, since $\dot{q}_0 = -\lambda_1 < 0$ and the RHS of the q-ODE is negative whenever $q_t = 1$.

Also note that $N_t := p_t - S_t q_t = p_t - G_{1,t}$, which has ODE $\dot{N}_t = -(\lambda_1 + \mu)N_t$. Since $N_0 = p_0 - G_{1,0} = p_0 \in (0, 1)$, we have $N_t = p_t - G_{1,t} > 0$ for all $t \in [0, \underline{T})$.

Define $\delta_t := q_t - p_t$, where $\delta_0 = 1 - p_0 > 0$. We have

$$\dot{\delta}_t = -\lambda_1 \delta_t + (1 - q_t)\mu \frac{N_t}{S_t}. \tag{38}$$

The inequalities just established imply the second term is positive for all $t \in (0, \underline{T}]$. Hence, by the comparison theorem, we have $\delta_t > 0$ for all $t \in [0, \underline{T}]$.

To prove that $A_{\underline{T}_0} > 0$, recall from earlier in the proof that $c'(\underline{T}) < 0$. Hence, $c(\underline{T})$ crosses 0 from above at $\underline{T}$. Now $c(\underline{T}) = 0$ and (36) imply

$$0 = \dot{q}_t - r q_t + (\lambda_1 + r) p_t = \dot{q}_t + \lambda_1 q_t - (1 - q_t)\mu \frac{p_t - S_t q_t}{S_t} \tag{39}$$

$$\implies S_t = S(p_t, q_t) \tag{40}$$

Now when $c(t) = 0$, some algebra shows that $c'(t) = -\frac{(\lambda_1+r)(q_t-p_t)}{1-q_t}A_t$. Since $c(\underline{T}) = 0$ by construction and $c'(\underline{T}) < 0$, it follows that $A_{\underline{T}} > 0$. □

The following lemma establishes that if the solution exits R at finite time T, it does so through $q_T = 1$.

Lemma 15. *If $T < T^{\max}$, then $q_T > p_T > 0$ and $A_T > 0$. Hence, $q_T = 1$.*

Proof. Throughout the proof, suppose $T < T^{\max}$. Since the right hand side of (33) contains a factor of p_t, the comparison theorem implies $p_t > 0$ for all $t \in [\underline{T}, T^{\max})$, and in particular, $p_T > 0$.

Note that for all $t \in [\underline{T}, T)$, we have $(p_t, q_t) \in R$ and therefore κ_t is well-defined and strictly positive. Moreover, S_t is well-defined and strictly positive. Hence, by construction, the ODEs for (G_1, G_0, p) are all satisfied.

The original ODE for p can be written

$$\dot{p}_t = -\lambda_1 p_t - S_t\kappa_t(q_t - p_t). \tag{41}$$

Defining $\delta_t = q_t - p_t$ and subtracting this ODE for p from the indifference condition for q, we have

$$\dot{\delta}_t = \delta_t(r + S_t\kappa_t). \tag{42}$$

It follows that δ is strictly increasing on $[\underline{T}, T)$, so $\delta_T > 0$; in other words, $q_T > p_T$.

We now establish that $A_T > 0$. Since $A_t > 0$ for all $t < T$, we must have $A_T \geq 0$. Toward a contradiction, suppose $A_T = 0$. We claim that $\dot{A}_T > 0$. A_t is differentiated as

$$\begin{aligned}\dot{A}_t &= (r-\mu)\dot{q}_t - 2\dot{p}_t(\lambda_1 + r)\\ &= (r-\mu)[rq_t - (\lambda_1+r)p_t]\\ &\quad - 2(\lambda_1+r)\left[p_t\frac{2(\lambda_1+r)\mu p_t + \mu q_t(\lambda_1+\mu-r) - (\lambda_1+\mu)(\lambda_1+r+\mu)}{\lambda_1+r+\mu(1-q_t)}\right].\end{aligned}$$

Now if $A_t = 0$, then $p_t = \frac{r+\lambda_1+\mu+(r-\mu)q_t}{2(\lambda_1+r)}$. Substituting this into the expression for $\dot{A}_t$ and simplifying yields

$$\begin{aligned}\dot{A}_t|_{A_t=0} &= C_0 + C_1 q_t, \qquad \text{where}\\ C_0 &:= \frac{1}{2}(2\lambda_1+\mu-r)(\lambda_1+r+\mu)\\ C_1 &:= \frac{1}{2}(r-\mu)(2\lambda_1+\mu+r).\end{aligned}$$

Since this is linear in q_t, it suffices to sign at extreme values. Using $q_t = 1$ yields $C_0 + C_1 = \frac{1}{2}\lambda_1(3r + 2\lambda_1 + \mu) > 0$. To evaluate at a lower bound on q_t, recall that $A_T = 0$ implies

$$p_T = \frac{r + \lambda_1 + \mu + (r - \mu)q_T}{2(\lambda_1 + r)} < q_T,$$

as we have already shown $p_T < q_T$. This in turn yields $q_T > \underline{q} := \frac{\lambda_1+\mu+r}{2\lambda_1+\mu+r}$. Hence,

$$\begin{aligned} C_0 + C_1\underline{q} &= \frac{1}{2}(2\lambda_1 + \mu - r)(\lambda_1 + r + \mu) + \frac{1}{2}(r - \mu)(\lambda_1 + \mu + r) \\ &= \lambda_1(\lambda_1 + r + \mu) > 0. \end{aligned}$$

As $q_t \in (\underline{q}, 1)$ it follows that $C_0 + C_1 q_t$ lies between two positive values obtained at the extreme values $\underline{q}$ and 1, and is therefore positive.

We conclude that $\dot{A}_T|_{A_T=0} > 0$, so A_t cannot intersect 0 from above at $t = T$. Thus, $A_T > 0$.

Now the 0 function solves the p-ODE, so by the comparison theorem, $p_t > 0$ for all $t \in [\underline{T}, T)$.

Having established that $q_T > p_T > 0$ and $A_T > 0$, the only remaining possibility is that $q_T = 1$. □

We now show that $T < T^{\max}$, so Lemma 15 indeed applies.

If it is not true that $T < T^{\max}$, there are two alternative possibilities: either $T = T^{\max} < \infty$ or $T = T^{\max} = \infty$. Lemma 16 rules out the first possibility, and (17) rules out the second.

Lemma 16. *If $T = T^{\max}$, then both are infinite.*

Proof. If $T = T^{\max}$ and T is finite, then we have $q_T = 1$, so (p_T, q_T) is in the interior of U by its definition. This implies that the solution extends uniquely beyond T, so $T^{\max} > T$, a contradiction. □

Lemma 17. *It does not hold that $T = T^{\max} = \infty$.*

Proof. Suppose $T = T^{\max} = \infty$. Then a unique solution exists for all time and remains in R. But the ODE for $\delta = q - p$ has solution

$$\delta_t = \delta_{\underline{T}} e^{\int_{\underline{T}}^{t}(r + S_s\kappa_s)\,ds} \geq \delta_{\underline{T}} e^{r(t-\underline{T})}$$

where we have used that $S_t > 0$ and $\kappa_t > 0$ for all t, since $(p_t, q_t) \in R$. In particular, δ reaches 1 in finite time. This contradicts the fact that in R, $0 < p_t < q_t < 1$. □

To summarize, we have $T < T^{\max}$, and the unique maximal solution exits R at time T with $q_T = 1$. We define $\overline{T} = T$. It follows that as $t \to \overline{T}$, $\kappa_t \to \infty$ and $S_t \to 0$.

Because q_t hits 1 from below at $\overline{T}$, we have $\dot{q}_{\overline{T}-} \geq 0$. As (20) holds throughout the purging phase and (p, q) are continuous at $\overline{T}$, it follows that $p_{\overline{T}} \leq \frac{rq_{\overline{T}}}{\lambda_1 + r} = \underline{p}$. From time $\overline{T}$ onward, we have $q_t = 1$, and optimality of the sender's strategy follows exactly

Standard verification shows that the sender's strategy is sequentially rational.

A.6.6 Moderate p_0

Suppose $p_0 \in (\underline{p}, \overline{p}]$. In this case, the equilibrium must begin in the purging phase. The system is the same as in the purging phase for the high-p_0 case, except for minor technical modifications due to the initial conditions $(G_0, G_1, S_0, q_0) = (0, 0, 0, 1)$. This means that κ_t given by (31) is not finite at time 0, and the right side of (18) is not well defined at time 0. Nonetheless, there is a unique solution of the reduced system in (p, q), S is well defined by (30), and κ is well defined by (31) except at time 0. There is a unique pair (G_1, G_0) that is continuous at time 0, satisfies the initial conditions, and satisfies (14)-(15) at all $t > 0$. Moreover, (16), (17), and (18) are satisfied at all $t > 0$ by construction.

A.7 Proof of Proposition 5

The proof has two main steps. First, we begin by proving the following corollary to Lemma 20, which establishes optimality of non-disclosure of bad evidence; it is stronger than needed for our purposes because it allows for arbitrary good evidence disclosure strategies, whereas our construction will involve a particular (three-phase) good evidence disclosure strategy. Second, we use a continuity argument to extend the good evidence disclosure strategy of Theorem 2.

For the first step, we use the following intermediate result.

Lemma 18. *Fix an arbitrary strategy H such that bad evidence is never disclosed. Let H^0 denote the strategy under which good evidence is disclosed immediately and bad evidence is never disclosed. Let p and p^0 denote the no-disclosure belief paths under H and H^0, respectively. We claim that for all $t \geq 0$, $p_t \geq p_t^0$.*

Proof of Lemma 18. The proof uses similar logic to that of Proposition 2. Fix $t \geq 0$. For strategy H, define N as the ex ante probability that no evidence is disclosed by time t. For each $s \in [0, t]$, let ν_s^G denote the probability that good evidence arriving at time s is disclosed by time t. Let $\nu_s^{G,0} \equiv 1$ and N^0 have analogous definitions to ν_s^G and N, but under strategy H^0 instead of H.

We now apply the law of iterated expectations to H and H^0. Even though timestamps are not observed by the receiver, we can condition on the true evidence-arrival time in the law-of-iterated-expectations decomposition. Thus,

$$\phi_t = Np_t + \int_0^t \mu\phi_s e^{-\mu s}\nu_s^G q_t^{G,s}\,ds \qquad \text{and} \qquad \phi_t = N^0 p_t^0 + \int_0^t \mu\phi_s e^{-\mu s}\nu_s^{G,0} q_t^{G,s}\,ds,$$

where we have used that bad evidence is not disclosed under either H or H^0.

Equating the right hand sides and subtracting p_t^0 throughout yields

$$\begin{aligned} N(p_t - p_t^0) + \int_0^t \mu\phi_s e^{-\mu s}\nu_s^G (q_t^{G,s} - p_t^0)\,ds &= 0 + \int_0^t \mu\phi_s e^{-\mu s}\nu_s^{G,0}(q_t^{G,s} - p_t^0)\,ds \\ \implies N(p_t - p_t^0) &= \int_0^t \mu\phi_s e^{-\mu s}(\nu_s^{G,0} - \nu_s^G)(q_t^{G,s} - p_t^0)\,ds. \end{aligned}$$

Since $N \geq e^{-\mu t} > 0$ and $1 = \nu_s^{G,0} \geq \nu_s^G$ for all s and $q_t^{G,s} \geq \phi_t \geq p_t^0$, we conclude that $p_t \geq p_t^0$. □

Corollary 4. *Fix a disclosure strategy in which bad evidence is never disclosed, and suppose that after off-path disclosure of bad evidence, the receiver updates his belief to* 0*. If* $p_0 > \frac{\lambda_0}{\lambda_0+\lambda_1+r}$ *and* $\mu < \left(\frac{\lambda_0+\lambda_1+r}{\lambda_1+r}\right) r$*, then it is optimal for the sender to never disclose bad evidence, independent of the good-evidence disclosure strategy.*

Proof of Corollary 4. From the proof of Lemma 20, under the given off-path beliefs, delaying disclosure of bad evidence is locally incentive compatible if $p_t \geq \frac{\lambda_0}{\lambda_0+\lambda_1+r}$ for all $t \geq 0$. Hence, it suffices to show that this inequality holds at all times.

Given Lemma 18, we only need to show that $p_t^0 \geq \frac{\lambda_0}{\lambda_0+\lambda_1+r}$ for all $t \geq 0$. The proof of Lemma 20 establishes that this condition is implied by the inequalities $p_0 > \frac{\lambda_0}{\lambda_0+\lambda_1+r}$ and $\mu < \left(\frac{\lambda_0+\lambda_1+r}{\lambda_1+r}\right) r$ taken together. □

For the second step, we now characterize the disclosure of good evidence and prove its optimality. During the stockpiling and purging phases, the laws of motion are

$$\dot{G}_{1,t} = G_{0,t}\lambda_0 - G_{1,t}\lambda_1 + (p_t - G_{1,t} - B_{1,t})\mu - G_{1,t}\kappa_t(1 - G_{1,t} - G_{0,t}) \tag{43}$$

$$\dot{G}_{0,t} = -G_{0,t}\lambda_0 + G_{1,t}\lambda_1 - G_{0,t}\kappa_t(1 - G_{1,t} - G_{0,t}) \tag{44}$$

$$\dot{B}_{1,t} = B_{0,t}\lambda_0 - B_{1,t}\lambda_1 + B_{1,t}\kappa_t(G_{0,t} + G_{1,t}) \tag{45}$$

$$\dot{B}_{0,t} = -B_{0,t}\lambda_0 + B_{1,t}\lambda_1 + (1 - p_t - G_{0,t} - B_{0,t})\mu + B_{0,t}\kappa_t(G_{0,t} + G_{1,t}). \tag{46}$$

$$\dot{p}_t = \lambda_0(1 - p_t) - \lambda_1 p_t - G_{1,t}\kappa_t + p_t(G_{0,t} + G_{1,t})\kappa_t. \tag{47}$$

During the transparency phase, the local indifference condition becomes

$$\dot{q}_{G,t} \geq (\lambda_0 + \lambda_1 + r)(q_{G,t} - p_t) + (\lambda_0 + \lambda_1)(p^* - q_{G,t}) \tag{48}$$
$$= rq_{G,t} - (\lambda_0 + \lambda_1 + r)p_t + \lambda_0. \tag{49}$$

Define $c(t; \lambda_0) = \dot{q}_{G,t} - [rq_{G,t} - (\lambda_0 + \lambda_1 + r)p_t + \lambda_0]$, generalizing $c(t)$ from the $\lambda_0 = 0$ case.

By repeating the arguments from the proof of Lemma 12, allowing for $\lambda_0 > 0$, we obtain cutoffs

$$\overline{p} := \frac{\lambda_0 + r + \frac{\lambda_1}{2}}{\lambda_0 + \lambda_1 + r}$$
$$\underline{p} := \frac{\lambda_0 + r}{\lambda_0 + \lambda_1 + r}.$$

Specifically, within the class of three-phase equilibrium with no disclosure of bad evidence, Lemma 12 goes through with these updated formulas for $\overline{p}$ and $\underline{p}$.

For $p_0 > \overline{p}$, we extend the definition of $\underline{T}$ as the first time t at which $c(t; \lambda_0) = 0$, and $\underline{T} = 0$ otherwise. Since the equations (43)-(47) and c are continuous in λ_0 and as we already showed $\frac{\partial}{\partial t} c(\underline{T}; 0) < 0$, $\underline{T}$ is continuous, and so are $(p_{\underline{T}}, q^G_{\underline{T}}, B_{1,\underline{T}}, B_{0,\underline{T}})$. After solving for $S(p, q^G, B_1)$ and $\kappa(p, q^G, B_1, B_0)$ in the same way as in the proof of Theorem 2, we obtain a system for (p, q^G, B_1, B_0) in the purging phase with initial conditions $(p_{\underline{T}}, q^G_{\underline{T}}, B_{1,\underline{T}}, B_{0,\underline{T}})$. This system and initial conditions are also continuous in λ_0 and are locally Lipschitz in (p, q^G, B_1, B_0), so the unique solution is continuous in λ_0 near 0. The feasible set is then defined by the inequalities $0 \leq B_{1,t} < p_t < q_t < 1$ which ensure that S and κ are well-defined, and $A(p, q^G, B_1, B_0; \lambda_0) > 0$, where A is continuous and coincides with A from the proof of Theorem 2 at $\lambda_0 = 0$, where it was bounded away from 0. A similar first crossing argument using the comparison theorem establishes that the solution can only exit through $q^G_t = 1$. Moreover, we still have that $\delta_t = q^G_t - p_t$ satisfies (42), implying that this exit occurs at some finite time $\overline{T}$. After $\overline{T}$, immediate disclosure is incentive compatible as before.

A.8 Proofs for Section 4.3

We begin with preliminary definitions. We define q^B_t as the belief immediately after disclosure of bad evidence. We define $q^{B,s}_t$ as the belief at time t conditional on bad evidence arriving at time s. Let σ^G and σ^B denote the time at which good or bad evidence, respectively, is disclosed (with the restriction that at most one of these is finite (and the other $+\infty$), since only one piece of evidence can arrive); let $\sigma = \min\{\sigma^G, \sigma^B\}$.

We define the following stock variables conditional on no disclosure up to time t:

$$
\begin{aligned}
B_{i,t} &= \Pr(\tau \le t \text{ and } \theta_\tau = 0 \text{ and } \theta_t = i | \sigma > t) \qquad \text{for } i \in \{0, 1\}, \\
B_t &= \Pr(\tau \le t \text{ and } \theta_\tau = 0 | \sigma > t) = B_{0,t} + B_{1,t} \\
\bar{q}_t^B &= \begin{cases} \frac{B_{1,t}}{B_t} & \text{if } B_t > 0 \\ 0 & \text{if } B_t = 0 \end{cases}.
\end{aligned}
$$

It is also sometimes convenient to work with the unnormalized (i.e., not dividing by the probability of nondisclosure) stock variables. We define

$$
\begin{aligned}
B_{i,t}^{un} &= \Pr(\tau \le t \text{ and } \theta_\tau = 0 \text{ and } \theta_t = i \text{ and } \sigma > t) \\
B_t^{un} &= \Pr(\tau \le t \text{ and } \theta_\tau = 0 \text{ and } \sigma > t) \\
N_t^{un} &= \Pr(\tau > t).
\end{aligned}
$$

Before proving Proposition 4, we state and prove the following supporting lemma, which provides a useful lower bound on the evolution of p_t. Intuitively, under age-independent bad evidence disclosure, such disclosure can only lower the belief, so nondisclosure of bad evidence can only raise it. The only remaining sources of downward pressure on the nondisclosure belief are therefore (i) Markov switching from state 1 to 0, and (ii) the fact that under immediate good evidence disclosure, no news is bad news. Since the proof is unrelated to the sender's optimization, we defer it to Section A.10.

Lemma 19. *Suppose that good evidence is disclosed immediately. Then under any age-independent bad evidence disclosure strategy, for all t, s we have $p_{t+s} \ge p_t - (\lambda_1 + \mu)s$. By implication, $p_t \ge p_{t-}$.*

Proof of Proposition 4. For the first part of the proposition, suppose by way of contradiction that there exists some fixed good evidence strategy, with a positive probability of delay, such that this strategy can be part of an equilibrium for arbitrarily large r. Let t_0 be the infimum of all good evidence types that delay with positive probability. Then for any $\epsilon > 0$, there exists $t \in [t_0, t_0 + \epsilon)$ and such that good evidence type t delays with positive probability. This means there exists some $\Delta_0 > 0$ such that there is positive probability of delay until at least $t + \Delta_0$. Moreover, there must exist $r_n \to \infty$ and an associated sequence of equilibria all of which contain this good evidence strategy, such that for each n, there exists $\Delta_n \ge \Delta_0$ such that delay until $t + \Delta_n$ is optimal. For each n along the sequence, the payoff of disclosing at time t is

$$\Pi(q^G_{n,t}) := \frac{p^*}{r_n} + \frac{q^G_{n,t} - p^*}{\Lambda + r_n}, \tag{50}$$

and the payoff of disclosing at $t + \Delta_n$ is

$$V(t, \Delta_n) := \int_t^{t+\Delta_n} e^{-r_n(s-t)} p_{n,s}\, ds + e^{-r_n \Delta_n} \Pi(q^G_{n,t+\Delta_n}).$$

We make use of the following bounds. First, by the support restriction, for each n, $q^G_{n,t} \geq \phi_{t-t_0}(1) \geq e^{-\lambda_1(t-t_0)}$, since good evidence disclosed at time t has age at most $t - t_0$. Second, $q^G_{n,t+\Delta_n} \leq 1$. Third, by Lemma 23, there is a uniform bound (across equilibria, n, and time) $p_{n,s} < p(\delta) < 1$. Incorporating these bounds yields

$$\begin{aligned} \Pi(q^G_{n,t}) &\geq \frac{p^*}{r_n} + \frac{e^{-\lambda_1(t-t_0)} - p^*}{\Lambda + r_n} \\ V(t, \Delta_n) &\leq \int_t^{t+\Delta_n} e^{-r_n(s-t)} p(\delta)\, ds + e^{-r_n \Delta_n} \Pi(1). \end{aligned}$$

Subtracting and scaling by r_n,

$$r_n(V(t, \Delta_n) - \Pi(q^G_{n,t})) \leq r_n \left(\int_t^{t+\Delta_n} e^{-r_n(s-t)} p(\delta)\, ds + e^{-r_n \Delta_n} \Pi(1) - \frac{p^*}{r_n} - \frac{e^{-\lambda_1(t-t_0)} - p^*}{\Lambda + r_n} \right).$$

Since $\Delta_n \geq \Delta_0 > 0$ for all n, the right hand side has limit

$$p(\delta) - e^{-\lambda_1(t-t_0)}.$$

By choosing t sufficiently close to t_0 at the outset, this bound is negative, and therefore for sufficiently large n, the strategy cannot be part of any equilibrium.

For the bad evidence disclosure part of the proposition, the payoff of disclosing at time t is

$$\int_0^\infty e^{-rs}[p^* - (p^* - q^B_t) e^{-(\lambda_0+\lambda_1)s}]\, ds = \frac{p^*}{r} - \frac{p^* - q^B_t}{r + \lambda_0 + \lambda_1}. \tag{51}$$

The payoff of disclosing at time $t + \epsilon$ for any $\epsilon > 0$ is at least

$$\int_0^\epsilon e^{-rs} p_{t+s}\, ds, \tag{52}$$

since the continuation flow utility after $t + \epsilon$ is nonnegative.

After subtracting (51) from (52) and scaling by r, the benefit of delay by ϵ is at least

$$\int_0^\epsilon re^{-rs}p_{t+s}\,ds - p^* + \frac{r}{r+\lambda_0+\lambda_1}(p^* - q_t^B). \tag{53}$$

By Lemma 19, this is at least

$$\int_0^\epsilon re^{-rs}[p_t - (\lambda_1+\mu)s]\,ds - p^* + \frac{r}{r+\lambda_0+\lambda_1}(p^* - q_t^B).$$

As $r \to \infty$, this converges uniformly in t to

$$p_t - p^* + (p^* - q_t^B) = p_t - q_t^B.$$

All that remains is to show that at all on-path bad evidence disclosure times $t \in [0, T]$, we have $p_t - q_t^B > 0$.

In an equilibrium with age-independent strategies, if disclosure is on path at time t, then $q_t^B = \bar{q}_{t-}^B = \frac{B_{1,t-}^{un}}{B_{t-}^{un}}$ at all times, with $q_0^B = 0$. We also have $p_t \geq p_{t-}$ by Lemma 19, where

$$p_{t-} = \frac{N_t^{un}\phi_t + B_{1,t-}^{un}}{N_t^{un} + B_{t-}^{un}}. \tag{54}$$

Hence,

$$p_t - q_t^B = p_t - \bar{q}_{t-}^B \geq p_{t-} - \bar{q}_{t-}^B = \frac{N_t^{un}}{N_t^{un} + B_{t-}^{un}}(\phi_t - \bar{q}_{t-}^B) \geq N_t^{un}(\phi_t - \bar{q}_{t-}^B) = e^{-\mu t}(\phi_t - \bar{q}_{t-}^B).$$

We also have $\bar{q}_{t-}^B \leq q_t^{B,0} = p^*(1 - e^{-\Lambda t})$ and $\phi_t = p^* - (p^* - p_0)e^{-\Lambda t}$, so

$$\phi_t - \bar{q}_{t-}^B \geq p_0 e^{-\Lambda t}.$$

Together, for all $t \in [0, T]$, these inequalities imply

$$p_t - q_t^B \geq e^{-\mu t} p_0 e^{-\Lambda t} \geq p_0 e^{-(\Lambda+\mu)T} > 0,$$

concluding the proof. □

Proof of Lemma 1. We prove the second part of the proposition first. More specifically, we state and prove the following lemma.

Lemma 20. *Assume $p_0 \in \left(\frac{\lambda_0}{\lambda_0+\lambda_1+r}, \frac{\lambda_0+r}{\lambda_0+\lambda_1+r}\right)$ and that $\mu < \left(\frac{\lambda_0+\lambda_1+r}{\lambda_1+r}\right)r$. Then there exists an equilibrium in which good evidence is disclosed immediately and bad evidence is never* *disclosed.*

Proof of Lemma 20. We verify optimality of the sender's strategy assuming the receiver correctly conjectures immediate disclosure of good evidence and no disclosure of bad evidence, and that following off-path disclosure of bad evidence at time t, the receiver assumes the evidence arrived at time t.

Let $q_{B,t}$ (identically 0 under our off-path belief assumption) denote the receiver's posterior immediately after disclosure of bad evidence.

Some algebra shows that delaying bad evidence disclosure is locally incentive compatible if and only if

$$\dot{q}_{B,t} \geq (\lambda_0 + \lambda_1 + r)(q_{B,t} - p_t) + (\lambda_0 + \lambda_1)(p^* - q_{B,t}).$$

With $q_B \equiv 0$, this reduces to

$$p_t \geq \frac{\lambda_0}{\lambda_0 + \lambda_1 + r}, \tag{55}$$

and hence to establish that it is optimal to never disclose bad evidence, it suffices to show that the receiver's belief conditional always satisfies this inequality. Note that by assumption, the inequality is satisfied strictly at time 0.

Let $q_{G,t}$ (identically 1 under the conjectured strategy) denote the receiver's belief immediately after disclosure of good evidence. Delaying disclosure of good evidence is not locally incentive compatible if

$$\begin{aligned}\dot{q}_{G,t} &< (\lambda_0 + \lambda_1 + r)(q_{G,t} - p_t) + (\lambda_0 + \lambda_1)(p^* - q_{G,t}) \\ &= r q_{G,t} - (\lambda_0 + \lambda_1 + r)p_t + \lambda_0.\end{aligned}$$

With $q_G \equiv 1$, this reduces to

$$p_t < \frac{\lambda_0 + r}{\lambda_0 + \lambda_1 + r}. \tag{56}$$

Hence, whenever this inequality is satisfied, the sender prefers to disclose good evidence immediately rather than at any future time. Recall that this inequality is satisfied at time 0 by assumption.

We now show that under the condition on μ in the proposition, $p_t \in \left(\frac{\lambda_0}{\lambda_0+\lambda_1+r}, \frac{\lambda_0+r}{\lambda_0+\lambda_1+r}\right)$ for all $t \geq 0$ conditional on no disclosure.

Let $B_{x,t}$ denote the probability that the sender has bad evidence and $\theta_t = x \in \{0, 1\}$, conditional on no disclosure yet. The laws of motion for (B_1, B_0, p) under the receiver's conjecture are

$$\dot{B}_{1,t} = B_{0,t}\lambda_0 - B_{1,t}\lambda_1 + B_{1,t}(p_t - B_{1,t})\mu. \tag{57}$$

$$\dot{B}_{0,t} = -B_{0,t}\lambda_0 + B_{1,t}\lambda_1 + (1 - p_t - B_{0,t})\mu + B_{0,t}(p_t - B_{1,t})\mu \tag{58}$$

$$\dot{p}_t = \lambda_0(1 - p_t) - \lambda_1 p_t - \mu(p_t - B_{1,t})(1 - p_t). \tag{59}$$

with $B_{0,0} = B_{1,0} = 0$.

Standard arguments using the comparison theorem show that there is a unique solution (B_1, B_0, p) to this system for all time, all variables lie in $(0, 1)$, and $p_t > B_{1,t}$. By inspection, the right hand side of (59) is negative when p_t exceeds $p^* = \frac{\lambda_0}{\lambda_0+\lambda_1}$, which lies in $\left(\frac{\lambda_0}{\lambda_0+\lambda_1+r}, \frac{\lambda_0+r}{\lambda_0+\lambda_1+r}\right)$. Hence, p_t cannot hit the upper end of this interval in finite time.

Next, note that the right hand side of (59) is bounded below by

$$f(p_t) := \lambda_0(1 - p_t) - \lambda_1 p_t - \mu(p_t - 0)(1 - p_t).$$

Let $\tilde{p}$ denote the solution to $\dot{\tilde{p}}_t = f(\tilde{p}_t)$ with initial condition $\tilde{p}_0 = p_0$. It is easy to show that $p_t \geq \tilde{p}_t$ at all times, and that $\tilde{p}_t$ converges monotonically to a steady state $\tilde{p}^* \in (0, 1)$. The condition $\mu < \left(\frac{\lambda_0+\lambda_1+r}{\lambda_1+r}\right) r$ ensures that this steady state is above $\frac{\lambda_0}{\lambda_0+\lambda_1+r}$. Therefore, p_t does not cross this threshold. □

For the first part of Lemma 1

Suppose the receiver conjectures immediate disclosure of good evidence and no disclosure of bad evidence. We state the explicit formulas for belief stock variables, which we refer to in later proofs as well:

$$B_{1,t} = \frac{\int_0^t \mu e^{-\mu s}(1 - \phi_s) q_t^{B,s}\, ds}{e^{-\mu t} + \int_0^t \mu e^{-\mu s}(1 - \phi_s)\, ds} \tag{60}$$

$$B_{0,t} = \frac{\int_0^t \mu e^{-\mu s}(1 - \phi_s)(1 - q_t^{B,s})\, ds}{e^{-\mu t} + \int_0^t \mu e^{-\mu s}(1 - \phi_s)\, ds} \tag{61}$$

$$B_{1,t}^{un} = \int_0^t \mu e^{-\mu s}(1 - \phi_s) q_t^{B,s}\, ds \tag{62}$$

$$B_{0,t}^{un} = \int_0^t \mu e^{-\mu s}(1 - \phi_s)(1 - q_t^{B,s})\, ds \tag{63}$$

$$B_t^{un} = \int_0^t \mu e^{-\mu s}(1 - \phi_s)\, ds \tag{64}$$

$$N_t^{un} = e^{-\mu t}. \tag{65}$$

Under passive beliefs, the belief immediately after disclosure of bad evidence at time t is $q_t^B = \bar{q}_t^B = \frac{B_{1,t}}{B_{0,t}+B_{1,t}}$, where $B_{1,t}$ and $B_{0,t}$ are given by (60) and (61).

We have

$$p_t = \frac{\phi_t N_t^{un} + B_{1,t}^{un}}{N_t^{un} + B_t^{un}} \tag{66}$$

$$\bar{q}_t^B = \frac{B_{1,t}}{B_t} = \frac{B_{1,t}^{un}}{B_t^{un}}. \tag{67}$$

A useful identity is

$$p_t - \bar{q}_t^B = \frac{N_t^{un}}{N_t^{un} + B_t^{un}}(\phi_t - \bar{q}_t^B). \tag{68}$$

Note that this is positive, since $\bar{q}_t^B \leq q_t^0 < \phi_t$. Moreover, (B_1^{un}, B^{un}) satisfy the system

$$\dot{B}_{1,t}^{un} = \lambda_0(B_t^{un} - B_{1,t}^{un}) - \lambda_1 B_{1,t}^{un} \tag{69}$$

$$\dot{B}_t^{un} = \mu e^{-\mu t}(1 - \phi_t). \tag{70}$$

Thus $\bar{q}^B$ satisfies

$$\dot{\bar{q}}_t^B = \frac{\dot{B}_{1,t}^{un}}{B_t^{un}} - \bar{q}_t^B \frac{\dot{B}_t^{un}}{B_t^{un}} = \lambda_0(1 - \bar{q}_t^B) - \lambda_1 \bar{q}_t^B - \bar{q}_t^B \frac{\dot{B}_t^{un}}{B_t^{un}}. \tag{71}$$

The local IC constraint for delaying disclosure is

$$\dot{q}_t^B \geq (\lambda_0 + \lambda_1 + r)(q_t^B - p_t) + (\lambda_0 + \lambda_1)(p^* - q_t^B) = r q_t^B - (\Lambda + r)p_t + \lambda_0. \tag{72}$$

Note that if this condition fails for all $s \geq t$, then disclosure at t is optimal. We now show that there exists such t.

Under passive off-path beliefs, we have $q_t^B = \bar{q}_t^B$ at all times. Using this and substituting (71), (72) becomes

$$\begin{aligned} \lambda_0(1 - \bar{q}_t^B) - \lambda_1 \bar{q}_t^B - \bar{q}_t^B \frac{\dot{B}_t^{un}}{B_t^{un}} &\geq r\bar{q}_t^B - (\Lambda + r)p_t + \lambda_0 \\ \iff \bar{q}_t^B \frac{\dot{B}_t^{un}}{B_t^{un}} &\leq (\Lambda + r)(p_t - \bar{q}_t^B). \end{aligned} \tag{73}$$

By (68) and rearranging, this is equivalent to

$$\Lambda + r \geq \bar{q}_t^B \frac{\dot{B}_t^{un}}{B_t^{un}} \frac{N_t^{un} + B_t^{un}}{N_t^{un}(\phi_t - \bar{q}_t^B)} = \bar{q}_t^B \frac{\mu(1-\phi_t)(N_t^{un} + B_t^{un})}{B_t^{un}(\phi_t - \bar{q}_t^B)}. \tag{74}$$

On the right hand side, assuming $\lambda_1, \lambda_0 > 0$, we have the following limits. First, $\phi_t \to p^* \in (0,1)$, $N_t^{un} \to 0$, $B_t^{un} \to B_\infty^{un} := \int_0^\infty \mu e^{-\mu s}(1-\phi_s)\, ds \in (0,1)$, and for each s, $\lim_{t\to\infty} q_t^{B,s} = p^*$. Next, by dominated convergence, $B_{1,t}^{un} = \int_0^\infty \mathbb{1}_{s\leq t} \mu e^{-\mu s}(1-\phi_s) q_t^{B,s}\, ds \to p^* \int_0^\infty \mu e^{-\mu s}(1-\phi_s)\, ds = p^* B_\infty^{un}$. Hence, $\bar{q}_t^B \to p^*$, which implies $\phi_t - \bar{q}_t^B \to 0$ from above. Hence, (74) fails for sufficiently large t, as desired. □

Proof of Theorem 3. We show that such an equilibrium exists for large r, supported by passive off-path beliefs satisfying the support restriction in the event of off-path bad evidence disclosure strictly between bursts.

We first establish optimality of immediately disclosing good evidence. Immediately after a burst, the receiver believes that the sender does not possess evidence, and thus beliefs coincide with $\phi_t = p^* = p_0$ given the stationary prior. In between bursts, the possibility that the agent possesses undisclosed bad evidence means that $p_t < p^*$. Hence, at all times we have $p_t \leq p^* < \frac{\lambda_0 + r}{\Lambda + r}$, and by the argument in the proof of Lemma 20, disclosing good evidence immediately is indeed optimal when it is conjectured.

We now show that for large r, there exists Δ such that the bad evidence disclosure policy is optimal when it is conjectured.

Suppose the bursts are conjectured to be spaced Δ units apart for some $\Delta > 0$. Immediately after any burst, normalizing the current time to 0, the belief variables evolve exactly as in (60)-(65) and (66)-(67).

The first optimality condition is that delay must be optimal between bursts. Under the passive off-path beliefs with support restriction, disclosing at time $t \in (0, \Delta)$ or at the next burst, $t = \Delta$, results in a posterior of $\bar{q}_t^B$. Hence, the local delay IC condition is still (73). At time 0, we have $\dot{q}_0^B > 0$ while the RHS of the IC condition reduces to $(\lambda_0 + \lambda_1 + r)p^* - (\lambda_0 + \lambda_1)p^* = -rp^* < 0$. Hence, by continuity, delay is strictly optimal initially. Moreover, we have already seen that (73) fails for large t; define $\bar{\Delta}_r$ as the infimum of such t:

$$\overline{\Delta}_r := \inf\{t > 0 : (73) \text{ fails}\}.$$

By definition, delay is locally incentive compatible in between bursts at regular intervals of length Δ if and only if $\Delta \leq \overline{\Delta}_r$.

The second optimality condition is that disclosure must be optimal at each burst. Suppose the time is $n\Delta$ and the agent has bad evidence. By disclosing it immediately, the sender obtains payoff calculated earlier as

$$\Pi(q^B_\Delta) := \frac{p^*}{r} - \frac{p^* - q^B_\Delta}{r + \lambda_0 + \lambda_1},$$

where we have used that $q^B_{n\Delta} = q^B_\Delta$ by the cyclicazl property of beliefs.

By waiting until $(n+1)\Delta$ to disclose, the sender obtains

$$\int_{n\Delta}^{(n+1)\Delta} e^{-r(s-n\Delta)} p_s \, ds + e^{-r\Delta}\Pi(q^B_\Delta).$$

Disclosure at time Δ is therefore optimal iff

$$\Gamma(\Delta) := \Pi(q^B_\Delta)(1 - e^{-r\Delta}) - \int_0^\Delta e^{-rs} p_s \, ds \geq 0, \tag{75}$$

where we have used that beliefs are cyclical with period Δ. Note that if the agent is willing to forego disclosure at a burst, he is willing to forego disclosure at all bursts, so the condition is equivalent to checking whether the agent prefers to disclose at a burst versus never disclose.

The condition (75) fails for all sufficiently small Δ. In particular, we have equality at $\Delta = 0$ while the derivatives of the left and right hand sides with respect to Δ at $\Delta = 0$ are $r\Pi$ and p^*, respectively, where $\Pi < p^*/r$. However, (75) holds for all sufficiently large Δ; taking $\Delta \to \infty$, we have $q^B_\Delta \to p^*$ and therefore $\Pi(q^B_\Delta) \to \Pi(p^*) = p^*/r$, while the value of concealing forever is $\int_0^\infty e^{-rs} p_s \, ds < \int_0^\infty e^{-rs} p^* = p^*/r$.

The rest of the proof consists of showing that for sufficiently large r, $\Gamma(\overline{\Delta}_r) > 0$, so that our conjectured strategy profile with $\Delta = \overline{\Delta}_r$ is an equilibrium.

Let $y_t := p^* - p_t$. For later use, note that $y_0 = 0, \dot{y}_0 = \mu p^*(1 - p^*)$, and $\ddot{y}_0 = \mu p^*(1 - p^*)(\mu(2p^* - 1) - \Lambda)$. Then $\Gamma(\overline{\Delta}_r)$ rearranges to

$$\int_0^{\overline{\Delta}_r} e^{-rs} y_s \, ds - (1 - e^{-r\overline{\Delta}_r}) \frac{p^* - \bar{q}^B_{\overline{\Delta}_r}}{\Lambda + r}$$

Lemma 21. *As $r \to \infty$, we have $\overline{\Delta}_r \to +\infty$, $r(p^* - \bar{q}^B_{\overline{\Delta}_r}) \to \dot{y}_0$, and $r^2\Gamma(\overline{\Delta}_r) \to 0$.*

In light of Lemma 21, to show eventual positivity of $\Gamma(\overline{\Delta}_r)$, we show that

$$\lim_{r\to\infty} r^3\Gamma(\overline{\Delta}_r) > 0.$$

Now $r^3\Gamma(\overline{\Delta}_r)$ expands as

$$r^3\Gamma(\overline{\Delta}_r) = r^3 \int_0^{\overline{\Delta}_r} e^{-rs} y_s \, ds - r^3(1 - e^{-r\overline{\Delta}_r}) \frac{p^* - \bar{q}^B_{\overline{\Delta}_r}}{\Lambda + r}. \tag{76}$$

Since the limit $r^2\Gamma_{\overline{\Delta}_r}$ had the form $\dot{y}_0 - \dot{y}_0$, taking limits on each term alone would yield $\infty - \infty$; so instead, we recenter by $\dot{y}_0$:

$$r^3\Gamma_{\overline{\Delta}_r} = \underbrace{r^3\left(\int_0^{\overline{\Delta}_r} e^{-rs} y_s \, ds - \frac{\dot{y}_0}{r^2}\right)}_{:=X(r)} - \underbrace{r^3\left((1 - e^{-r\overline{\Delta}_r}) \frac{p^* - \bar{q}^B_{\overline{\Delta}_r}}{\Lambda + r} - \frac{\dot{y}_0}{r^2}\right)}_{=:Y(r)}. \tag{77}$$

Lemma 22. *As $r \to \infty$, $X(r) \to \ddot{y}_0 = \dot{y}_0(\mu(2p^* - 1) - \Lambda)$ and $Y(r) \to \dot{y}_0(\max\{\Lambda - \mu, 0\} - 2\Lambda - \mu(1 - p^*)) < \ddot{y}_0$. Hence, $\lim_{r\to\infty} r^3\Gamma_{\overline{\Delta}_r} > 0$.*

Given Lemma 22, $\Gamma(\overline{\Delta}_r) > 0$ for sufficiently large $r > 0$, so the construction is complete. □

Proof of Lemma 2. Since bad evidence is disclosed immediately at all $t \geq t_0$, all bad evidence disclosed strictly after t_0 is assumed to be fresh: $q^B_t \equiv 0$ for all $t > t_0$. This implies that after disclosing bad evidence at any $t > t_0$, the sender's flow payoff at each $s \geq t$ is $\phi_{s-t}(0)$. Moreover, conditional on nondisclosure after t_0, the receiver believes that with probability 1, the sender does not possess evidence of either type. Hence $p_t \equiv \phi_t$ for all $t > t_0$. It follows that by never disclosing, the sender's flow payoff at each $s \geq t$ is $\phi_s = \phi_{s-t}(\phi_t) > \phi_{s-t}(0)$. Thus, disclosing bad evidence at t is not optimal for any $t > t_0$. □

A.9 Proof of Corollary 2

Proof. Consider any three-phase equilibrium with stockpiling phase ending at $\underline{T}_0$. Since there is no disclosure before $\underline{T}_0$, the on-path variables $(p_{\underline{T}_0}, S_{\underline{T}_0}, G_{1,\underline{T}_0})$ at time $\underline{T}_0$ are independent of the off-path beliefs before $\underline{T}_0$, and these pin down q_{T_0}. Hence

$$(p_{\underline{T}_0}, S_{\underline{T}_0}, q_{\underline{T}_0}) = (p^0_{\underline{T}_0}, S^0_{\underline{T}_0}, q^0_{\underline{T}_0}).$$

Additionally, the law of motion for q depends only on (p, q, S). Local indifference in the purging phase therefore implies

$$\dot{q}^0_{\underline{T}_0} = r q^0_{\underline{T}_0} - (\lambda_1 + r) p^0_{\underline{T}_0}.$$

By Lemma 13, this equation has the unique solution $\underline{T}$. Thus $\underline{T}_0 = \underline{T}$. □

A.10 Proofs of supporting results

Proof of Lemma 19. Let π_N, π_G, and π_B denote the probability of no disclosure, good evidence disclosure, and bad evidence disclosure, respectively, on $(t, t+s]$ conditional on no disclosure on $[0,t]$. Let $\phi_s(p_t)$ denote the belief at time $t+s$ based on the belief p_t at time t and no information other than Markov switching dynamics in $(t, t+s]$. By the law of iterated expectations,

$$\phi_s(p_t) = \pi_N p_{t+s} + \pi_B \, \mathbb{E}[\theta_{t+s} | \sigma^B \in (t, t+s]] + \pi_G \, \mathbb{E}[\theta_{t+s} | \sigma^G \in (t, t+s]].$$

We claim that $\mathbb{E}[\theta_{t+s}|\sigma^B \in (t,t+s]] \leq \phi_s(\bar{q}_t^B) \leq \phi_s(p_t)$. To understand the first inequality, write

$$\begin{aligned}\mathbb{E}[\theta_{t+s}|\sigma^B \in (t,t+s]] = \Pr(\tau \leq t | \sigma^B \in (t,t+s])\mathbb{E}[\theta_{t+s}|\tau \leq t, \sigma^B \in (t,t+s]] \\ + \Pr(\tau > t | \sigma^B \in (t,t+s])\mathbb{E}[\theta_{t+s}|\tau > t, \sigma^B \in (t,t+s]].\end{aligned}$$

For the first term, since bad evidence disclosure is age-independent, we have $\mathbb{E}[\theta_{t+s}|\tau \leq t, \sigma^B \in (t,t+s]] = \phi_s(\bar{q}_t^B)$. For the second term, if evidence arrives at some time $u \in (t, t+s]$, then $\theta_u = 0$, and the conditional belief about θ_{t+s} is $\phi_{t+s-u}(0) \leq \phi_s(0) \leq \phi_s(\bar{q}_t^B)$. Hence, averaging over u, we have $\mathbb{E}[\theta_{t+s}|\tau > t, \sigma^B \in (t,t+s]] \leq \phi_s(\bar{q}_t^B)$ and overall $\mathbb{E}[\theta_{t+s}|\sigma^B \in (t,t+s]] \leq \phi_s(\bar{q}_t^B) \leq \phi_s(p_t)$, where the last inequality is due to $\bar{q}_t^B \leq p_t$.

Also, $\mathbb{E}[\theta_{t+s}|\sigma^G \in (t,t+s]] \leq 1$. Rearranging the equation and using these inequalities, we have

$$\begin{aligned}\pi_N p_{t+s} &\geq \phi_s(p_t) - \pi_B \phi_s(p_t) - \pi_G \\ \iff \pi_N(p_{t+s} - p_t) &\geq (\pi_N + \pi_G)(\phi_s(p_t) - p_t) - \pi_G(1 - p_t).\end{aligned}$$

Now $\pi_N > 0$ and we have

$$\begin{aligned}p_{t+s} - p_t &\geq \left(1 + \frac{\pi_G}{\pi_N}\right)(\phi_s(p_t) - p_t) - \frac{\pi_G}{\pi_N}(1 - p_t) \\ &= (\phi_s(p_t) - p_t) - \frac{\pi_G}{\pi_N}(1 - \phi_s(p_t)). \qquad (78)\end{aligned}$$

Expanding the first term yields

$$\begin{aligned}\phi_s(p_t) - p_t &= p^* - (p^* - p_t)e^{-(\lambda_0+\lambda_1)s} - p_t = (p^* - p_t)(1 - e^{-(\lambda_0+\lambda_1)s}) \\ &\geq -(1-p^*)(1 - e^{-(\lambda_0+\lambda_1)s}) = -\frac{\lambda_1}{\lambda_0+\lambda_1}(1 - e^{-(\lambda_0+\lambda_1)s}) \geq -\lambda_1 s.\end{aligned}$$

Let M be the probability that no evidence has arrived in $[0, t]$ conditional on no disclosure yet in $[0, t]$. We have $\pi_G \le M(1 - e^{-\mu s})$ and $\pi_N \ge M e^{-\mu s}$, and therefore $\frac{\pi_G}{\pi_N} \le \frac{1-e^{-\mu s}}{e^{-\mu s}} = e^{\mu s} - 1$. Hence, returning to (78),

$$p_{t+s} - p_t \ge -\lambda_1 s - (e^{\mu s} - 1).$$

Since we can partition $[t, t+s]$ into N intervals of width s/N and apply the same bound on each interval, summing, and taking N arbitrarily large, we have

$$p_{t+s} - p_t \ge -(\lambda_1 + \mu)s.$$

To derive the implication $p_t \ge p_{t-}$, fix $u > 0$, set $t = u - s$, and take $s \to 0$. Then the previous result and existence of a left limit imply $p_u \ge p_{u-s} - (\lambda_1 + \mu)s \to p_{u-}$. □

The following lemma states a useful upper bound on p_t. Intuitively, for p_t to be close to 1, it must be that the agent possesses very recent evidence with very high probability, which is impossible.

Lemma 23. *Fix any $\delta > 0$. Then in all equilibria, p_t is uniformly bounded away from 1: $p_t < p(\delta) := e^{-\mu\delta} \max\{p_0, p^* + (1 - p^*)e^{-\Lambda\delta}\} + (1 - e^{-\mu\delta}) \cdot 1 < 1$.*

Proof. Suppose there has been no disclosure up to time t. By definition, the belief is p_t. Now fix any $\delta > 0$. Let $\underline{t}(\delta) := \max\{t - \delta, 0\}$. There are four possibilities: (i) no evidence has arrived by t, (ii) good evidence arrived by $\underline{t}(\delta)$, (iii) good evidence arrived in $(\underline{t}(\delta), t]$, and (iv) bad evidence arrived. In case (i), the updated belief is $\phi_t \le \max\{p_0, p^*\}$. In case (ii), it is at most $\phi_\delta(1) = p^* + (1 - p^*)e^{-\Lambda\delta}$. In case (iii), it is at most 1. In case (iv), it is at most $\phi_t(0) = p^*(1 - e^{-\Lambda t}) \le p^*$. Let A be the probability of (iii). Then we have

$$p_t \le (1 - A) \max\{p_0, p^* + (1 - p^*)e^{-\Lambda\delta}\} + A \cdot 1. \tag{79}$$

Also, by Bayes' rule we have

$$\begin{aligned}
A &= \frac{\Pr(\tau^G \in (\underline{t}(\delta), t]), \sigma > t)}{\Pr(\sigma > t)} \\
&\leq \frac{\Pr(\tau \in (\underline{t}(\delta), t]), \sigma > t)}{\Pr(\sigma > t)} \\
&= \frac{\Pr(\tau \in (\underline{t}(\delta), t]), \sigma > t)}{\Pr(\tau \in (\underline{t}(\delta), t], \sigma > t) + \Pr(\tau \leq t - \delta, \sigma > t) + \Pr(\tau > t)} \\
&\leq \frac{\Pr(\tau \in (\underline{t}(\delta), t]), \sigma > t)}{\Pr(\tau \in (\underline{t}(\delta), t], \sigma > t) + \Pr(\tau > t)} \\
&= \frac{e^{-\mu(t-\delta)} - e^{-\mu t}}{e^{-\mu(t-\delta)} - e^{-\mu t} + e^{-\mu t}} \\
&= 1 - e^{-\mu\delta}.
\end{aligned}$$

Using this in the inequality for p_t yields the uniform upper bound in the lemma. □

Proof of Lemma 21. The fact that $\overline{\Delta}_r \to \infty$ can be seen from (73), where r only enters the term $(\Lambda + r)$, and hence on any finite interval, the inequality holds for all sufficiently large r.

Next, we show that $\lim_{r\to\infty} r^2 \Gamma_{\overline{\Delta}_r} = 0$. We have

$$r^2 \Gamma_{\overline{\Delta}_r} = r^2 \int_0^{\epsilon} e^{-rs} y_s \, ds + r^2 \int_{\epsilon}^{\overline{\Delta}_r} e^{-rs} y_s \, ds - r^2 (1 - e^{-r\overline{\Delta}_r}) \frac{p^* - \bar{q}^B_{\overline{\Delta}_r}}{\Lambda + r}. \tag{80}$$

On the right hand side of (80), the second term clearly tends to 0 as $r \to \infty$. The first term Taylor-expands as

$$\begin{aligned}
&r^2 \int_0^{\epsilon} e^{-rs} [y_0 + s\dot{y}_0 + \frac{s^2}{2}\ddot{y}_0 + o(s^2)] \, ds \\
&= \dot{y}_0 [1 - e^{-r\epsilon}(1 + r\epsilon)] + \frac{\ddot{y}_0}{2r}[2 - e^{-r\epsilon}(2 + r\epsilon(2 + r\epsilon))] + r^2 \int_0^{\epsilon} e^{-rs} o(s^2) \, ds \\
&\to \dot{y}_0 + 0 + 0 = \dot{y}_0.
\end{aligned}$$

The last term in (80) has the same limit as $r(p^* - \bar{q}^B_{\overline{\Delta}_r})$ as $r \to \infty$ because $\frac{r}{\Lambda + r}(1 - e^{-r\overline{\Delta}_r}) \to 1$ (due to $\overline{\Delta}_r \to \infty$). Because of the stationary prior, $p^* = \phi_{\overline{\Delta}_r}$, and then since local IC (74) binds at $\overline{\Delta}_r$, we have

$$r(p^* - \bar{q}^B_{\overline{\Delta}_r}) = \frac{r}{\Lambda + r} \bar{q}^B_{\overline{\Delta}_r} \frac{\mu(1 - p^*)(N^{un}_{\overline{\Delta}_r} + B^{un}_{\overline{\Delta}_r})}{B^{un}_{\overline{\Delta}_r}} \to \mu p^*(1 - p^*) = \dot{y}_0, \tag{81}$$

where we have used that $N^{un}_{\overline{\Delta}_r} = e^{-\mu\overline{\Delta}_r}$, that $\overline{\Delta}_r \to \infty$, and the limits derived in the proof of Lemma 1. Thus, the last term in (80) tends to $\dot{y}_0$, so $\lim_{r\to\infty} r^2\Gamma_{\overline{\Delta}_r} = \dot{y}_0 + 0 - \dot{y}_0 = 0$. $\square$

Proof of Lemma 22. For $X(r)$, we split the integral into $[0,\epsilon)$ and $[\epsilon,\Delta)$ and repeat the Taylor expansion from the proof of Lemma 21:

$$\int_0^{\overline{\Delta}_r} e^{-rs} y_s\, ds = \int_0^{\epsilon} e^{-rs}[y_0 + \dot{y}_0 s + \frac{1}{2}\ddot{y}_0 s^2 + o(s^2)]\, ds + \int_{\epsilon}^{\overline{\Delta}_r} e^{-rs} y_s\, ds. \tag{82}$$

We have $r^3 \int_{\epsilon}^{\overline{\Delta}_r} e^{-rs} y_s\, ds \to 0$ as $r \to \infty$, so $X(r)$ has the same limit as

$$r^3 \left(\int_0^{\epsilon} e^{-rs}[y_0 + \dot{y}_0 s + \frac{1}{2}\ddot{y}_0 s^2 + o(s^2)]\, ds - \frac{\dot{y}_0}{r^2} \right). \tag{83}$$

Recall that $y_0 = 0$ and $\dot{y}_0 = -\dot{p}_0 = A$. Hence (83) simplifies to

$$r^3 \dot{y}_0 \left(\int_0^{\epsilon} e^{-rs} s\, ds - \frac{1}{r^2} \right) + r^3 \int_0^{\epsilon} e^{-rs} \frac{1}{2}\ddot{y}_0 s^2\, ds + r^3 \int_0^{\epsilon} e^{-rs} o(s^2)\, ds. \tag{84}$$

We address each term in (84) individually. The first term reduces to

$$r^3 \dot{y}_0 \left(\frac{1}{r^2}(1 - (1 + r\epsilon)e^{-r\epsilon}) - \frac{1}{r^2} \right) = -\dot{y}_0 r(1 + r\epsilon)e^{-r\epsilon} \xrightarrow{r\to\infty} 0. \tag{85}$$

The second term in (84) gives

$$\frac{1}{2}\ddot{y}_0[2 - (2 + 2r\epsilon + r^2\epsilon^2)e^{-r\epsilon}] \to \ddot{y}_0. \tag{86}$$

Finally, the third limit can be made arbitrarily small by taking ϵ sufficiently small. We conclude that $\lim_{r\to\infty} X(r) = \ddot{y}_0 = \dot{y}_0(\mu(2p^* - 1) - \Lambda)$.

For $Y(r)$, rearrange it to

$$r^3(1 - e^{-r\overline{\Delta}_r}) \left(\frac{p^* - \bar{q}^B_{\overline{\Delta}_r}}{\Lambda + r} - \frac{\dot{y}_0}{r^2} \right) - r^3 e^{-r\overline{\Delta}_r} \frac{\dot{y}_0}{r^2}.$$

The final term tends to 0 and $1 - e^{-r\overline{\Delta}_r} \to 1$, leaving us to calculate the limit of

$$\begin{aligned}
r^3\left(\frac{p^* - \bar{q}^B_{\overline{\Delta}_r}}{\Lambda + r} - \frac{\dot{y}_0}{r^2}\right) &= \dot{y}_0 r^3\left(\frac{\bar{q}^B_{\overline{\Delta}_r}(N^{un}_{\overline{\Delta}_r} + B^{un}_{\overline{\Delta}_r})}{(\Lambda + r)^2 B^{un}_{\overline{\Delta}_r} p^*} - \frac{1}{r^2}\right) \\
&= \dot{y}_0 r \frac{r^2 \bar{q}^B_{\overline{\Delta}_r}(N^{un}_{\overline{\Delta}_r} + B^{un}_{\overline{\Delta}_r}) - (\Lambda + r)^2 B^{un}_{\overline{\Delta}_r} p^*}{(\Lambda + r)^2 B^{un}_{\overline{\Delta}_r} p^*} \\
&= \dot{y}_0\left[\frac{r^3(\bar{q}^B_{\overline{\Delta}_r}(N^{un}_{\overline{\Delta}_r} + B^{un}_{\overline{\Delta}_r}) - B^{un}_{\overline{\Delta}_r} p^*)}{(\Lambda + r)^2 B^{un}_{\overline{\Delta}_r} p^*} - \frac{2\Lambda r^2}{(\Lambda + r)^2} - \frac{r\Lambda^2}{(\Lambda + r)^2}\right] \qquad (87)
\end{aligned}$$

where the first equality uses the expression for $r(p^* - \bar{q}^B_{\overline{\Delta}_r})$ in (81) and that $\mu(1 - p^*) = \frac{\dot{y}_0}{p^*}$, the second simply combines fractions and simplifies, and the third splits by degree of r. Now the second and third terms in the square brackets in (87) contribute $-2\Lambda\dot{y}_0 - 0$ to the limit as $r \to \infty$. The first term in (87) contributes

$$\dot{y}_0 \frac{r^3 B^{un}_{\overline{\Delta}_r}(\bar{q}^B_{\overline{\Delta}_r} - p^*)}{(\Lambda + r)^2 B^{un}_{\overline{\Delta}_r} p^*} + \dot{y}_0 \frac{r^3 \bar{q}^B_{\overline{\Delta}_r} N^{un}_{\overline{\Delta}_r}}{(\Lambda + r)^2 B^{un}_{\overline{\Delta}_r} p^*}. \qquad (88)$$

Since $r(\bar{q}^B_{\overline{\Delta}_r} - p^*) \to -\dot{y}_0$ by Lemma 21, the first term of (88) tends to $-\frac{(\dot{y}_0)^2}{p^*}$. The second term of (88) tends to the same limit as $\dot{y}_0 r \frac{N^{un}_{\overline{\Delta}_r}}{B^{un}_{\overline{\Delta}_r}}$, because $\bar{q}^B_{\overline{\Delta}_r} \to p^*$. We now calculate the limit of

$$r\frac{N^{un}_{\overline{\Delta}_r}}{B^{un}_{\overline{\Delta}_r}} = r(p^* - \bar{q}^B_{\overline{\Delta}_r})\frac{\frac{N^{un}_{\overline{\Delta}_r}}{B^{un}_{\overline{\Delta}_r}}}{p^* - \bar{q}^B_{\overline{\Delta}_r}},$$

where again $r(p^* - \bar{q}^B_{\overline{\Delta}_r}) \to \dot{y}_0$. Since r only enters the remaining term via $\overline{\Delta}_r$, we are only left with calculating

$$L := \lim_{t\to\infty} \frac{N^{un}_t}{B^{un}_t(p^* - \bar{q}^B_t)}, \qquad (89)$$

for which we turn to closed forms. For $\Lambda \neq \mu$, under the stationary prior, we have

$$q^B_t = p^*\left(1 + \frac{\mu}{\Lambda - \mu}\frac{e^{-\Lambda t} - e^{-\mu t}}{1 - e^{-\mu t}}\right) \qquad \text{and} \qquad \frac{N^{un}_t}{B^{un}_t} = \frac{e^{-\mu t}}{(1 - p^*)(1 - e^{-\mu t})}.$$

Thus

$$\frac{N^{un}_t}{B^{un}_t(p^* - \bar{q}^B_t)} = \frac{e^{-\mu t}}{-(1 - p^*)(1 - e^{-\mu t})p^* \frac{\mu}{\Lambda - \mu}\frac{e^{-\Lambda t} - e^{-\mu t}}{1 - e^{-\mu t}}} = -\frac{\Lambda - \mu}{\dot{y}_0(e^{-(\Lambda - \mu)t} - 1)}. \qquad (90)$$

If $\Lambda > \mu$, we have $L = \frac{\Lambda - \mu}{\dot{y}_0}$. If $\Lambda < \mu$, $L = 0$. For the knife-edge case $\Lambda = \mu$, we have $\frac{N_t^{un}}{B_t^{un}(p^* - \bar{q}_t^B)} = \frac{1}{\dot{y}_0 t}$, so again $L = 0$. Hence, in general, $L = \frac{\max\{\Lambda - \mu, 0\}}{\dot{y}_0}$.

Retracing our steps, we have $\lim_{r\to\infty} Y(r) = -\frac{(\dot{y}_0)^2}{p^*} + (\dot{y}_0)^2 \frac{\max\{\Lambda-\mu,0\}}{\dot{y}_0} - 2\Lambda \dot{y}_0$, where

$$\begin{aligned}
\lim_{r\to\infty}(X(r) - Y(r)) &= \dot{y}_0(\mu(2p^* - 1) - \Lambda) - \left[-\frac{(\dot{y}_0)^2}{p^*} + \dot{y}_0 \max\{\Lambda - \mu, 0\} - 2\Lambda \dot{y}_0 \right] \\
&= \dot{y}_0 \left[\mu p^* + \Lambda - \max\{\Lambda - \mu, 0\} \right] > 0.
\end{aligned}$$

□

References

Acharya, Viral V, Peter DeMarzo, and Ilan Kremer, "Endogenous information flows and the clustering of announcements," *American Economic Review*, 2011, *101* (7), 2955–2979.

Antić, Nemanja and Harry Pei, "Selective Disclosure in Overlapping Generations," 2026. Working paper.

Dye, Ronald A, "Disclosure of nonproprietary information," *Journal of accounting research*, 1985, pp. 123–145.

Ely, Jeffrey C, "Beeps," *American Economic Review*, 2017, *107* (1), 31–53.

_ , George Georgiadis, and Luis Rayo, "Feedback design in dynamic moral hazard," *Econometrica*, 2025, *93* (2), 597–621.

Grossman, S J and D Hart, "Disclosure Laws and Takeover Bids," *The Journal of Finance*, 1980, *35* (2), 323–334.

Grossman, Sanford J., "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 1981, *24* (3,), 461–483.

Guttman, Ilan, Ilan Kremer, and Andrzej Skrzypacz, "Not only what but also when: A theory of dynamic voluntary disclosure," *American Economic Review*, 2014, *104* (8), 2400–2420.

Kremer, Ilan, Amnon Schreiber, and Andrzej Skrzypacz, "Disclosing a random walk," *The Journal of Finance*, 2024, *79* (2), 1123–1146.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, *12* (2), 380.