Selective Disclosure in Overlapping Generations*

Nemanja Antić† Harry Pei‡

August 26, 2026

Abstract

We develop an overlapping-generations model with a constant state. Each agent observes a private signal about the state with some probability, and, independently with another probability, receives the set of signals disclosed by his predecessor. The agent then takes an action and selects a subset of signals to disclose to the next agent. Each agent's action has a positive externality on his immediate predecessor and his optimal action increases in his belief about the state. We show that when agents are more likely to observe private signals or receive disclosed signals, they become increasingly selective in their disclosure decisions. When the probability that disclosed signals are transmitted to the next generation exceeds a threshold, all signals except the most favorable one will be concealed.

Keywords: information disclosure, hard information, overlapping generations.

JEL Codes: C72, C73, D82, D83.

*We thank Nageeb Ali, Odilon Câmara, Eddie Dekel, Jeffrey Ely, Matteo Escudé, Mira Frick, Drew Fudenberg, Ying Gao, Dino Gerardi, Germán Gieczewski, Christian Hellwig, Navin Kartik, Peter Klibanoff, Bart Lipman, Qingmin Liu, Suraj Malladi, Moritz Meyer-ter-Vehn, Wojciech Olszewski, Juan Ortner, Marcin Pęski, Joao Ramos, Lara Sanchez, Sara Shahanaghi, Bruno Strulovici, and Joel Watson for helpful comments. Pei thanks the NSF Grant SES-2337566 for financial support.

†Northwestern University, nemanja.antic@kellogg.northwestern.edu

‡Northwestern University, harrydp@northwestern.edu

1 Introduction

In many settings, information is transmitted through a succession of agents. To fix ideas, consider an organization with overlapping cohorts. Senior members are managers who benefit from the effort of juniors (i.e., new recruits) but do not internalize its cost. Managers seek to motivate new recruits by shaping their beliefs about the returns to hard work. For example, before passing the organization's training materials to new recruits, a manager may revise them by adding newly available cases about past projects or employees or removing cases inherited from earlier managers.[1] In addition to deliberate omissions, cases may also be lost unintentionally over time due to negligence, data loss, or technical failures. Recruits who later become managers can then revise the materials they inherit before passing them on to the next cohort.

This paper studies how better access to, and more reliable transmission of, information affect agents' disclosure behavior in an overlapping-generations (OLG) setting. We show that agents become more selective in disclosing information when they are more likely to have private information or to receive signals disclosed by other agents, in contrast to Dye (1985). This difference arises because, in our setting, disclosing an additional signal is informative about the amount of evidence available for selection: it suggests a longer chain of successful information transmission and, consequently, more opportunities for unfavorable signals to have been concealed.

We study steady-state equilibria in an OLG model with a fixed binary state (either high or low). In each period, only one agent is active. The active agent observes a newly drawn private signal with some probability, and independently with a separate probability, receives the set of signals disclosed by his immediate predecessor. The agent then chooses his effort and discloses to the next agent a subset of *available signals*, which includes his private signal and any signals he inherited. Agents cannot directly observe others' actions, the exact times at which the disclosed signals arrived, and whether the lack of inherited signals is caused by communication failure or their predecessors' deliberate choice to disclose nothing.[2]

Each agent's optimal effort is strictly increasing in his belief that the state is high and higher

[1] As documented in Kipping and Clark (2012)'s handbook, consulting firms often motivate new recruits by transmitting the firm's culture through stories of successful projects, client impact, and exemplary consultants.

[2] In our example, it seems hard for managers to add fabricated cases to the organization's official training material since newly added cases are typically subject to internal reviews. The cases contained in the training material may not include precise information about the exact time at which each case occurred, since client confidentiality and past employees' privacy concerns may require the distributed files to omit identifying information.

effort benefits his immediate predecessor through a positive externality. Consequently, each agent's disclosure problem reduces to choosing a subset of available signals that maximizes his immediate successor's posterior belief that the state is high.

As in other disclosure games with multiple signals, our model admits many strict equilibria, some of which are sustained by unreasonable off-path beliefs. For example, with three signal realizations, there can exist strict equilibria where signal realizations with the highest likelihood ratio are concealed while those with the second-highest likelihood ratio are revealed. Such equilibria can be sustained by the off-path belief that the disclosure of any signal with the highest likelihood ratio indicates the sender concealed signals with the lowest likelihood ratio.

We therefore focus on strict equilibria in which agents use *monotone strategies*. Monotonicity requires that whenever a signal is disclosed at an information set, all available signals with strictly higher likelihood ratios are also disclosed at that information set.

We show that a signal is disclosed in some strict monotone equilibrium *if and only if* a measure of the informational environment falls below a signal-specific cutoff. The measure corresponds to the expected number of private signals observed since the latest communication failure and increases both when agents are more likely to observe their private signals and when disclosed signals are more likely to reach the next generation. The signal-specific cutoff is infinite only for the signal with the highest likelihood ratio, smaller for signals with lower likelihood ratios, and positive if and only if the signal's likelihood ratio exceeds 1. Consequently, each signal other than the strongest one is disclosed in equilibrium if and only if its likelihood ratio exceeds 1 and agents' chances of observing private and disclosed signals are sufficiently low.

Our result implies that when agents receive private signals or disclosed signals with higher probability, they become *more selective* in disclosing information in the sense that a smaller set of signals can be disclosed in equilibrium. It also implies that when communication frictions are small enough, all signals except for the strongest possible one will be concealed in equilibrium.

Our finding contrasts with Dye (1985)'s result that a higher probability of the sender having information leads to *less selective* disclosure. In Dye's model, when this probability is close to one, all signal realizations will be disclosed except for the one with the lowest likelihood ratio.[3]

To build intuition, consider a *constant-threshold equilibrium* where at all information sets,

[3] A similar logic implies that in Dye (1985)'s setting, if the receiver fails to receive the information disclosed with some probability, disclosure becomes *less selective* when this probability decreases. This mirrors the finding in Dye (1985)'s original model where the sender does not have information with positive probability.

agents disclose a signal if and only if its likelihood ratio is above some cutoff (i.e., *good signals*). The sequence of observed private signals after the latest communication failure has some good signals (the number of which can be 0), and in between each pair of adjacent good signals, a number of bad signals (could be anything from 0 to ∞). Hence, if an agent inherits one more good signal from his predecessor, he also infers from the equilibrium strategy that there are more bad signals in expectation. Such an inference lowers his belief about the state. This effect is absent in Dye (1985)'s model where any disclosure implies that nothing is hidden.

If agents are more likely to observe private signals or more likely to receive disclosed signals, then, in expectation, there will be more bad signals between each pair of adjacent good signals. As a result, the negative effect of disclosing a good signal becomes more pronounced. Disclosure therefore becomes more selective in equilibrium, since each disclosed signal needs to have a high enough likelihood ratio in order to overcome the stronger negative effect.

In general, agents' disclosure threshold can differ across information sets, so the marginal effect of disclosing another signal can depend in complex ways on what other signals are disclosed. Our proof exploits a *gradualism property* that distinguishes our OLG game from one-shot disclosure games: the number of signals disclosed can increase by at most one in any period.

Focusing on the most informative constant-threshold equilibrium, we also show that the amount of information contained in the disclosed signals is strictly increasing in the probabilities that agents observe private signals and receive disclosed signals, even though higher probabilities make disclosure more selective. Yet even as both probabilities approach 1, learning about the state remains incomplete. The reason is that observing a large number of disclosed signals may also indicate that the most recent communication failure occurred in the distant past.

Our work contributes to the literature on disclosure games where the sender has an unknown amount of signals.[4] Most of the existing literature on this topic studies one-shot games.[5] Shin

[4] Di Tillio, Ottaviani, and Sørensen (2021) study an agent who observes the realizations of the k ($\leq n$) highest draws among n conditionally i.i.d. signals, which differs from our model where the number of signals disclosed is endogenous. In Bardhi and Bobkova (2023), each sender decides whether to disclose his signal without observing its realization. In Antić and Chakraborty (2025), the state is n-dimensional and the sender selects which of k ($\leq n$) realized values of those dimensions to disclose (typically not the highest ones). Onuchic and Ramos (2025) consider multiple senders who decide whether to disclose a multi-dimensional state via a collective decision rule.

[5] Bull and Watson (2019) examine a one-shot disclosure game where the sender's evidence is correlated with an unverifiable signal, so disclosing good evidence may lead to unfavorable inference about the unverifiable signal.

(1994, 2003) and Dziuda (2011) focus on binary signals, which are not well-suited to studying the selectiveness of disclosure. Gao (2025) assumes that the sender has a continuum of signals and analyzes equilibria satisfying the *truth-leaning refinement* of Hart, Kremer, and Perry (2017). Rappoport (2025) focuses on receiver-optimal equilibria and shows that the receiver will take worse actions for the sender when the sender is expected to have more information.[6] We instead examine how the selectiveness of the sender's disclosure behavior changes with the informational environment. We also discuss in Section 5.1 the relationship between our monotonicity refinement, the truth-leaning refinement, and the receiver-optimality refinement.

A growing literature studies disclosure in other dynamic settings. In Guttman, Kremer, and Skrzypacz (2014), a sender may receive up to two signals over two periods and chooses both what to disclose and when. In Felgenhauer and Schulte (2014), Felgenhauer and Loerke (2017), Lou (2023), Arieli and Stewart (2025), and Dai, Fudenberg, and Pei (2026), a sender conducts a sequence of experiments before selecting a subset of outcomes to disclose to a receiver. In Gieczewski (2022) and Squintani (2026), some agents know the state and can disclose it to their neighbors. In Kremer, Schreiber, and Skrzypacz (2024), a potentially informed sender decides whether to disclose the current value of a random walk. Kolb and Zhou (2026) study settings in which the sender has at most one signal and compare the equilibrium dynamics when the sender can credibly disclose the timestamp of the signal with those when the sender cannot. Our paper complements this literature by examining how the informational environment affects disclosure selectiveness in settings where the sender may possess arbitrarily many signals.

Our work also contributes to the study of OLG models starting from Samuelson (1958). While many papers such as Bhaskar (1998) and Acemoglu and Jackson (2015) focus on the observability of agents' actions, we study settings where actions are not observed, so that agents can influence their successors' behaviors only by disclosing information. Anderlini, Gerardi, and Lagunoff (2012) analyze an OLG model with strategic communication and payoff externalities. Communication in their model is cheap talk rather than the disclosure of evidence. They examine whether agents learn the state in the long run, while we study the selectiveness of disclosure.

Our model is also related to the social learning literature, such as Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), and Smith and Sørensen (2000).[7] Unlike those models, our

[6] Gieczewski and Titova (2024) propose a coalitional proofness refinement and provide conditions under which their refinement coincides with the truth-leaning refinement and the receiver-optimality refinement.

[7] Niehaus (2011) studies the sharing of complementary or substitutable skills among a sequence of agents.

setting has payoff externalities and agents learn from signals disclosed by others rather than from their predecessors' actions. Bénabou and Vellodi (2025) incorporate strategic disclosure in a three-period social learning model. They show that disclosure becomes less selective as communication frictions vanish, which contrasts with our finding.

2 Baseline Model

Consider a game with a doubly infinite time horizon and a constant state of the world $\theta \in \Theta = \{\underline{\theta}, \overline{\theta}\}$. Let $\pi_0 \in (0,1)$ denote the prior probability that $\theta = \overline{\theta}$.

At the beginning of each period $t \in \mathbb{Z}$, a signal is drawn from a finite set $S = \{s_1, s_2, \ldots, s_l\}$ where $f_{\theta,i}$ is the probability of s_i conditional on θ. To simplify notation, we write $\overline{f}_i$ instead of $f_{\overline{\theta},i}$ and $\underline{f}_i$ instead of $f_{\underline{\theta},i}$. Let $\mathbf{f} = (\overline{f}_i, \underline{f}_i)_{i=1}^{l}$. Let $L_i = \overline{f}_i / \underline{f}_i$ denote the *likelihood ratio* of s_i. We assume that $l \geq 2$, $f_{\theta,i} > 0$ for every θ and i, $L_i \neq L_j$ for every $i \neq j$,[8] and that the signals drawn in different periods are i.i.d. conditional on θ. We index the signal realizations so that

$$\infty > L_1 > L_2 > \cdots > L_l > 0. \tag{2.1}$$

Agent t is the only active player in period $t \in \mathbb{Z}$. With probability $p \in (0,1]$, he observes the signal drawn in period t. Formally, we write $s^t = s_i$ if agent t observes the signal drawn in period t (i.e., his *private signal*) and its realization is s_i, and $s^t = \emptyset$ if he does not observe that signal. Independently of whether agent t observes the signal drawn in period t, with probability $\delta \in (0,1)$, he receives the (multi)set of signals disclosed to him by agent $t-1$. Agent t then takes an action $a^t \in \mathbb{R}$ and selects a set of signals to disclose to agent $t+1$.

We assume that agents can observe the count of signals disclosed to them but cannot directly observe their predecessors' actions, whether and when communication failures occurred, or the exact dates and sequence of the disclosed signals. Formally, in addition to his private signal, each agent only observes his *history* $\mathbf{h} = (h_1, h_2, \ldots, h_l) \in H = \mathbb{N}^l$. If agent $t-1$ reveals h_1 copies of s_1, h_2 copies of s_2, ..., and h_l copies of s_l, then agent t receives this set of signals with probability δ, in which case agent t's history is $(h_1, h_2, \ldots, h_l)$. Agent t fails to receive this set of signals with probability $1-\delta$, in which case agent t's history is $\mathbf{0} = (0, \ldots, 0)$. We interpret $1-\delta$ as a communication friction and refer to δ as the communication success rate.

[8] Analogs of our results can be derived when multiple signal realizations share the same likelihood ratio, or when there is a continuum of signal realizations. Details of these extensions are available upon request.

As in Grossman (1981), Milgrom (1981), and Dye (1985), we assume that signals can be concealed but not falsified. As a result, each agent can only reveal a subset of the *available signals*, including those he received from his immediate predecessor and his own private signal. Formally, let $\mathbf{1}_j \in \mathbb{N}^l$ denote a vector where the jth entry is 1 and all other entries are 0. If agent t's history is $\mathbf{h}$, then he can only disclose vectors no more than $\mathbf{h}+\mathbf{1}_j$ (in the vector order on $\mathbb{N}^l$) when $s^t = s_j$, and can only disclose vectors no more than $\mathbf{h}$ when $s^t = \emptyset$.

Agent t's payoff is $u(\theta, a^t) + v(\theta, a^{t+1})$, which depends on the state θ, his action a^t, and his immediate successor's action a^{t+1}. We assume that (i) v is strictly increasing in a^{t+1}, (ii) for each $\pi \in [0,1]$, there is a unique $a \in \mathbb{R}$ that maximizes $\pi u(\overline{\theta}, a) + (1-\pi) u(\underline{\theta}, a)$, and this unique maximizer is strictly increasing in π.

Our solution concept is *steady-state Bayesian equilibrium* (or *equilibrium*), which extends the notion of (strict) pure-strategy steady-state equilibrium in Clark, Fudenberg, and Wolitzky (2021) to an incomplete information game.[9]

An *equilibrium* of our game consists of (i) a pure strategy $\sigma : H \times \Big(S \cup \{\emptyset\}\Big) \to \mathbb{R} \times H$ that maps an agent's history and private signal to his action and a subset of available signals he discloses to his immediate successor, (ii) two distributions $\underline{\mu}, \overline{\mu} \in \Delta(H)$ over histories, and (iii) a belief map $\pi : H \to [0,1]$ which represents the probability that an agent assigns to state $\overline{\theta}$ after observing his history but before observing his private signal. These components satisfy (i) when all agents use strategy σ, the steady-state history distribution is $\underline{\mu}$ conditional on $\theta = \underline{\theta}$ and is $\overline{\mu}$ conditional on $\theta = \overline{\theta}$,[10] (ii) $\pi(\mathbf{h})$ is derived via Bayes rule from $(\underline{\mu}, \overline{\mu})$ at every $\mathbf{h}$ that occurs with positive probability,[11] and (iii) σ maximizes an agent's expected payoff when agents' beliefs after observing their histories are given by π.

A *strict equilibrium* is one where agents' equilibrium strategy is strictly optimal at every information set.[12] A strict equilibrium always exists, as will be implied by Proposition 1.

[9]We discuss mixed-strategy equilibria in Section 5.3. We allow $p = 1$ but require that $\delta < 1$ since steady-state equilibria are not well-defined when $\delta = 1$. Similarly, existing steady-state analysis such as Clark, Fudenberg, and Wolitzky (2021) and Pei (2026) assume that agents die with positive probability after each period.

[10]As in earlier works that study steady-state equilibria such as Clark et al. (2021) and Pei (2026), we require all agents to use the same mapping $\sigma : H \times \Big(S \cup \{\emptyset\}\Big) \to \mathbb{R} \times H$. We show in Appendix A.1 that for every pure strategy σ and conditional on every state $\theta \in \{\overline{\theta}, \underline{\theta}\}$, there is a unique steady-state distribution over histories.

[11]Our results extend if we additionally require that, for every $\mathbf{h} \in H$, the belief $\pi(\mathbf{h})$ be derived from some conditional probability system that is consistent with the agents' equilibrium strategy.

[12]Our theorem extends when we only require players to have strict incentives at on-path information sets.

Interpretation: In the organizational example from the introduction, agent t is a junior worker in period t who faces a strictly convex cost of effort, and whose effort benefits his manager, agent $t-1$. The state θ represents the returns to effort. The history $\mathbf{h}$ in period $t-1$ represents the stock of verified cases contained in the organization's official training materials, which could be about past consulting projects done by the firm or past employees. If $s^{t-1} = \emptyset$, the manager has no new case to add. If $s^{t-1} \neq \emptyset$, there is a new case based on the manager's own observations, such as stories about his peers and coworkers or projects undertaken by the company when he was a junior, that could be added to the training material. Finally, $1-\delta$ captures the accidental loss of the inherited record, for example, due to data loss, negligence, or other unforeseen disruptions.

The assumption that agents can disclose only available signals means that managers may omit cases but cannot fabricate them, since newly added materials are typically subject to internal review before being incorporated into the training materials presented to new recruits. Because training materials can be copied, summarized, and anonymized as they are passed down over time, juniors may be unable to know precisely when a case occurred or the identities of the individuals involved, even if they can reliably learn its substantive content.[13] In period $t+1$, some workers from period t become managers and can revise the organization's training materials by adding or removing cases before presenting them to the next cohort of workers.

Section 5 discusses our modeling choices in greater detail, along with several alternative formulations and extensions. These include settings in which agents observe the exact timestamps of disclosed signals or the exact time at which communication failed. We also consider settings in which the successful transmission of different signals is independent or positively, but not perfectly, correlated; the communication success rate depends on the number of signals disclosed; and agents face constraints on the number of signals they can disclose. Finally, we discuss the differences between our OLG model and an analogous one-shot disclosure game in which the receiver faces uncertainty about the number of signals the sender has.

[13]In consulting firms, cases are usually reviewed before being included in the case library or training materials. Client confidentiality and privacy concerns may require distributed materials to omit identifying information, including the clients and former employees involved, as well as the exact dates on which the cases occurred.

3 Main Result

Our game admits many strict equilibria because each agent may possess arbitrarily many signals, allowing his successor's belief to be arbitrarily pessimistic at off-path histories. This, in turn, discourages agents from inducing those off-path histories, making those beliefs self-fulfilling. We start with some preliminary observations that describe some of those strict equilibria:

Proposition 1. *For any $j \in \{1, \dots, l-1\}$ that satisfies $L_j > 1$, there is a strict equilibrium where at each history $\mathbf{h} = (h_1, \dots, h_l)$, agents disclose all available s_j and nothing else, i.e., they disclose $(h_j + 1)\mathbf{1}_j$ if they observe private signal s_j, and disclose $h_j \mathbf{1}_j$ otherwise.*

The proof is in Appendix A.2. Proposition 1 implies that an equilibrium always exists in which agents disclose all realized s_1 and nothing else. However, for any $s_j \neq s_1$ with a likelihood ratio greater than 1, there is also a strict equilibrium in which agents disclose only s_j and conceal the rest, including those with strictly higher likelihood ratios. Intuitively, these equilibria can be sustained by the off-path belief that whenever any signal other than s_j (including those with higher likelihood ratios) is disclosed, the agent who disclosed it must have concealed signals with low likelihood ratios, which justifies their successors' pessimistic off-path beliefs about θ.

Equilibria where agents conceal higher-likelihood-ratio signals while revealing lower-likelihood-ratio ones seem unreasonable. For example, compare two scenarios in which agent t receives only one signal from agent $t-1$. Intuitively, agent t should be more *optimistic* about the state if this signal is s_1 instead of s_2, since s_1 has a higher likelihood ratio. More generally, agent t should be more optimistic about the state when agent $t-1$ replaces a low-likelihood-ratio signal with a high-likelihood-ratio signal in the set of signals disclosed. If so, agent $t-1$ will have no incentive to disclose any s_j while concealing signals with likelihood ratios strictly greater than L_j.

The above logic motivates our *monotonicity refinement*.[14] Let $\sigma_D(\mathbf{h}, s) \in H = \mathbb{N}^l$ denote the set of signals disclosed to the next agent when the current agent observes history $\mathbf{h}$ and receives private signal $s \in S \cup \{\emptyset\}$. Notice that σ_D is a component of agents' strategy σ.

Monotone Strategy & Monotone Equilibrium. *A strategy σ is monotone if for every $\mathbf{h} = (h_1, \dots, h_l) \in H$, $s \in S \cup \{\emptyset\}$, and $j \in \{2, \dots, l\}$, the jth entry of $\sigma_D(\mathbf{h}, s)$ being strictly*

[14] In Section 5.1, we explain why standard refinements such as requiring agents' off-path beliefs to be derived from a conditional probability system or satisfying the intuitive criterion and D1, have little bite in our setting. We also relate our monotonicity refinement to other refinements in the voluntary disclosure literature that restrict agents' off-path beliefs, such as the belief monotonicity refinement in Guttman, Kremer, and Skrzypacz (2014).

positive implies that for every $i < j$, the ith entry of $\sigma_D(\mathbf{h}, s)$ equals h_i when $s \neq s_i$ and equals $h_i + 1$ when $s = s_i$. An equilibrium is monotone if agents use a monotone strategy.

Hence, a strategy is monotone if when at least one copy of s_j is disclosed at an information set, all available signals with likelihood ratios strictly greater than L_j are also disclosed at that information set. All equilibria in Dye (1985)'s model are monotone. In our dynamic setting, monotonicity allows, for example, agents to conceal s_1 at history-signal pair $(\mathbf{h}, s)$ and reveal s_2 at $(\mathbf{h}', s') \neq (\mathbf{h}, s)$, as long as all copies of s_1 contained in $(\mathbf{h}', s')$ are also revealed at $(\mathbf{h}', s')$.

A special class of monotone strategies is the class of *constant-threshold strategies*. Formally, for every $\mathbf{h} \in H$ and $j \in \{0, 1, \ldots, l\}$, let $\mathbf{h}[j] \in H$ denote an l-dimensional vector whose first j entries coincide with $\mathbf{h}$ and whose other entries are all 0. By definition, $\mathbf{h}[0] = \mathbf{0}$ and $\mathbf{h}[l] = \mathbf{h}$.

Constant-Threshold Strategy & Constant-Threshold Equilibrium. *A strategy σ is a constant-threshold strategy if there exists $j \in \{0, 1, 2, \ldots, l\}$ such that (i) $\sigma_D(\mathbf{h}, s_i) = (\mathbf{h} + \mathbf{1}_i)[j]$ for every $\mathbf{h} \in H$ and $i \in \{1, \ldots, l\}$ and (ii) $\sigma_D(\mathbf{h}, \emptyset) = \mathbf{h}[j]$. A constant-threshold equilibrium is an equilibrium in which agents use a constant-threshold strategy.*

Compared to monotone strategies, constant-threshold strategies require a fixed disclosure threshold that does not vary with agents' histories or realized private signals. They also preclude partial disclosure of identical signals, e.g., if an agent has two copies of s_2 available, a constant-threshold strategy must prescribe either disclosing both copies of s_2 or concealing both.[15]

Among the equilibria in Proposition 1, those in which agents disclose s_j for some $j \geq 2$ and conceal s_1 are not monotone. The equilibrium in which agents reveal only s_1 is a (strict) constant-threshold equilibrium, so it is also a (strict) monotone equilibrium.

Some natural follow-up questions arise. Are there other strict monotone equilibria or constant-threshold equilibria beyond the ones described in Proposition 1, such as equilibria in which agents reveal s_1, s_2, s_3 and conceal all other signals, or in which agents reveal all copies of s_1 when they have s_1 available and reveal all copies of s_2 when they have no s_1 available? More generally, how does agents' equilibrium disclosure behavior depend on the signal structure $\mathbf{f}$, the probability p with which they receive private signals, and the communication friction $1 - \delta$?

[15] Section 5.1 extends our result to equilibria where agents use *general threshold strategies*. Such strategies allow the disclosure threshold to depend on the history and private signal, but require that at every history-signal pair, either all available copies of a signal realization are disclosed or all available copies of it are concealed.

Our main theorem characterizes, for any given signal $s_j \in S$, conditions under which s_j is disclosed in some strict monotone equilibrium (or constant-threshold equilibrium). Let

$$\eta(\delta, p) = \frac{\delta p}{1-\delta}. \tag{3.1}$$

Intuitively, in the steady state, $\delta/(1-\delta)$ is the expected number of periods since the most recent communication failure. Therefore, $\eta(\delta, p)$ is the expected number of private signals observed along an uninterrupted sequence of successful communication. By definition, $\eta(\delta, p) > 0$ for every $(\delta, p) \in (0, 1) \times (0, 1]$, increases in both δ and p, and diverges to ∞ as $\delta \to 1$.

We set $\eta_1 = \infty$, and for every $j \in \{2, \ldots, l\}$, let

$$\eta_j = \frac{\overline{f}_j - \underline{f}_j}{\underline{f}_j \sum_{i=1}^{j-1} \overline{f}_i - \overline{f}_j \sum_{i=1}^{j-1} \underline{f}_i} \quad \text{or, equivalently,} \quad \eta_j = \frac{L_j - 1}{\sum_{i=1}^{j-1} \underline{f}_i (L_i - L_j)}. \tag{3.2}$$

For every $j \geq 2$, the denominators in (3.2) are strictly positive because $L_i > L_j$ for every $i < j$. Hence, $\eta_j > 0$ if and only if $L_j > 1$. Moreover, one can show that $\infty = \eta_1 > \eta_2 > \cdots > \eta_l = -1$.

Theorem 1. *For any $j \in \{1, \ldots, l\}$:*

1. *If $\eta(\delta, p) \geq \eta_j$, then in any strict monotone equilibrium, signal s_j is never disclosed at any history that occurs with positive probability.*

2. *Signal s_j is disclosed in some strict constant-threshold equilibrium if and only if $\eta(\delta, p) < \eta_j$. Signal s_j is disclosed in some constant-threshold equilibrium if and only if $\eta(\delta, p) \leq \eta_j$.*

The proof is in Section 4. Since every constant-threshold equilibrium is also a monotone equilibrium, Theorem 1 implies that signal s_j is disclosed in some strict monotone (or strict constant-threshold) equilibrium if and only if $\eta(\delta, p) < \eta_j$, and is disclosed in some constant-threshold equilibrium if and only if $\eta(\delta, p) \leq \eta_j$. Because $\eta_1 = \infty$, there always exists a strict constant-threshold equilibrium in which agents disclose s_1. Agents will never disclose any signal with a likelihood ratio weakly less than 1, because $\eta_j \leq 0$ for every j satisfying $L_j \leq 1$. Since η_j is strictly decreasing in j, an increase in $\eta(\delta, p)$ reduces the set of signal realizations that can be disclosed. Since $\eta(\delta, p) \to \infty$ as $\delta \to 1$, all signals except for the strongest one s_1 will be concealed in equilibrium when the communication friction $1 - \delta$ is sufficiently small.

Applying (3.1) and (3.2), Theorem 1 implies that signal $s_j \neq s_1$ is disclosed in some strict monotone equilibrium only if (i) communication frictions are large enough (i.e., δ is small), (ii) agents do not observe private signals too frequently (i.e., p is small), and (iii) signals with

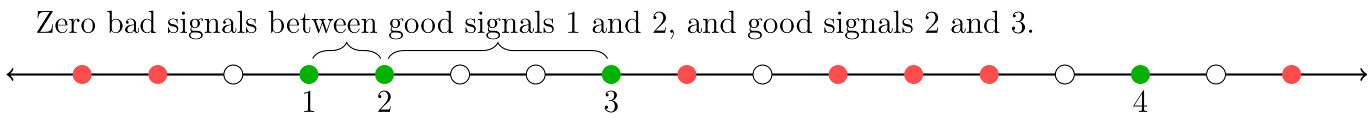

Figure 1: A sequence of realized signals after the most recent communication failure. Red circles represent observed bad signals, green circles represent observed good signals, and white circles represent unobserved signals. In this example, the number of bad signals is 0 between good signals 1 and 2 as well as between good signals 2 and 3, and is 4 between good signals 3 and 4.

likelihood ratios greater than L_j either arrive with low probability, or have likelihood ratios close to L_j. As p and δ become larger, i.e., agents are more likely to observe private signals and to inherit signals from their predecessors, $\eta(\delta, p)$ increases and by Theorem 1, disclosure becomes *more selective* in the sense that a smaller set of signals can be disclosed in equilibrium.

Our finding contrasts with that of Dye (1985), who shows that as the probability that the sender has information increases, disclosure becomes *less selective*, in the sense that a larger set of signals is disclosed in equilibrium. Furthermore, one can show that in Dye (1985)'s setting, when the receiver fails to receive the information disclosed by the sender with probability $1-\delta$, disclosure becomes *less selective* as δ increases. In the limiting case considered by Grossman (1981) and Milgrom (1981) where the sender has information and can successfully transmit it with probability close to 1, no disclosure leads the receiver to infer that the sender has the worst possible signal, which motivates the sender to disclose even the second-worst signal.

This difference between our result and Dye (1985)'s arises due to the nature of the uncertainty the receivers face about the senders' information structure. To see this, let us fix an arbitrary period $t \in \mathbb{Z}$ and consider the sequence of signals drawn from period τ, the last period before t in which communication failed, to period $t-1$. Some of these signals are observed and others are not. In Figure 1, the unobserved signals are represented by white dots. We pick an arbitrary s_i that satisfies $L_i > 1$. We say that an observed signal is *good* if $s^k \in \{s_1, \ldots, s_i\}$ (green dot) and is *bad* if $s^k \in \{s_{i+1}, \ldots, s_l\}$ (red dot). As shown in Figure 1, between each pair of adjacent good signals, there is a number of (observed) bad signals, which ranges from 0 to ∞.

Suppose that, in equilibrium, agents use a constant-threshold strategy of disclosing all good signals and concealing all bad ones. Each agent can then observe the number of copies of each good signal in the sequence since the latest failure but not the number or content of bad signals.

Recall that in between each pair of adjacent good signals, there is a number of bad signals, which, from the receiving agent's perspective, is a random variable $\widetilde{n}$. Since signals are conditionally i.i.d., the distribution of $\widetilde{n}$, conditional on state θ, does not depend on the good signals realized.

Hence, disclosing one more good signal has two effects on the next agent's belief. First, it reveals that another good signal is realized, which increases the belief as good signals have likelihood ratios more than 1. Second, the next agent can infer from the equilibrium strategy that $\widetilde{n}$ additional bad signals are realized and concealed, which has a negative effect on belief. This differs from Dye (1985) where the receiver believes that nothing is concealed after disclosure.

As (δ, p) increases, the distribution of $\widetilde{n}$, conditional on any θ, shifts to the right, so disclosing another good signal has a more pronounced negative effect. As a result, equilibrium disclosure becomes more selective, since only signals with a high enough likelihood ratio can offset the stronger negative effect. This logic is reflected in our result since

$$\eta(\delta, p) < \eta_j \quad \Leftrightarrow \quad L_j > \frac{1 + \eta(\delta, p) \sum_{i=1}^{j-1} \overline{f}_i}{1 + \eta(\delta, p) \sum_{i=1}^{j-1} \underline{f}_i}. \tag{3.3}$$

Nevertheless, the *net effect* of disclosing the strongest possible signal s_1 is always positive. To see why, consider an auxiliary scenario where agents observe only the total number of good signals $s_1, \ldots, s_i$ in the sequence, but not the identity of each one. If they learn that there is no good signal, their posterior is strictly less than their prior π_0. Suppose by way of contradiction that disclosing an additional good signal in this auxiliary game has a negative net effect on the next agent's belief. Then agents' posteriors would be strictly lower than their prior regardless of the number of good signals disclosed, which cannot happen since agents' belief is a martingale. Since the likelihood ratio of s_1 is no less than that of an "average" good signal in the set $\{s_1, \ldots, s_i\}$, the net effect of disclosing an additional s_1 on the next agent's belief must be positive.

The above logic does not generalize to all monotone strategies, since the set of signals disclosed can be history-dependent. In this case, even conditional on θ, the expected number of undisclosed signals between each pair of adjacent disclosed signals may depend on the set of signals that were already disclosed, and so too may the net effect of disclosing another copy of a given signal.

Our proof uses a *gradualism property* that distinguishes our overlapping-generations game from one-shot disclosure games in which receivers face uncertainty about the number of signals:

Gradualism Property. *For any $t \in \mathbb{Z}$ and any agent $t+1$ history $\mathbf{h}^{t+1} = (h_1^{t+1}, \ldots, h_l^{t+1})$ that comes after agent t history $\mathbf{h}^t = (h_1^t, \ldots, h_l^t)$, there exists at most one $j \in \{1, \ldots, l\}$ such that $h_j^{t+1} > h_j^t$. If such j exists, then $h_j^{t+1} = h_j^t + 1$.*

Intuitively, this property is driven by our modeling assumption that each agent receives at most one private signal in addition to the signals disclosed by his immediate predecessor.

Now, fix any strict monotone equilibrium and suppose s_j is the weakest signal disclosed on path. A consequence of our gradualism property is that there exist on-path histories with exactly one copy of s_j. We pick the shortest among those on-path histories, and if there are multiple, pick any one that minimizes the next agent's belief. We denote the chosen history by $\mathbf{h}^*$.

The key technical step is to show that when an agent's history is $\mathbf{h}^*$, his immediate predecessor's history is either $\mathbf{h}^*$ or $\mathbf{h}^* - \mathbf{1}_j$ (see Lemma 1).[16] That is, in the period right before history reaches $\mathbf{h}^*$ for the first time, the additional signal included must be s_j.

We then examine the marginal effect of disclosing s_j at history $\mathbf{h}^* - \mathbf{1}_j$ by comparing agents' beliefs at $\mathbf{h}^* - \mathbf{1}_j$ and at $\mathbf{h}^*$. As in constant-threshold equilibria, disclosing s_j reveals to the next agent that one copy of s_j occurred, which multiplies the likelihood ratio of their belief by L_j. However, when $s_j \neq s_1$, observing $\mathbf{h}^*$ can be bad news since it shows that no signal better than s_j has been realized after the disclosed s_j. This is because if an agent observes a private signal better than s_j at $\mathbf{h}^*$, he will no longer reveal $\mathbf{h}^*$ to the next agent as it requires disclosing s_j while hiding his private signal that has a higher likelihood ratio. This cannot happen when the agent uses a monotone strategy.

Hence, agents will disclose s_j at $\mathbf{h}^* - \mathbf{1}_j$ only if L_j is large enough to overcome the negative effect of no better signal arriving after s_j. This negative effect is stronger when signals better than s_j are more likely to arrive, such as when communication frictions are smaller (i.e., δ is larger) and when agents observe their private signals with higher probability (i.e., p is larger).

We conclude this section by discussing the implications of Theorem 1. First, when $\eta(\delta, p) \geq \eta_2$, no signal other than s_1 can be disclosed in any strict monotone equilibrium. This leads to a characterization of all strict monotone equilibria when $\eta(\delta, p) \geq \eta_2$:[17]

Corollary 1. *If $\eta(\delta, p) \geq \eta_2$, every strict monotone equilibrium is characterized by $n \in \mathbb{N} \cup \{\infty\}$ such that agents disclose up to n copies of s_1 and conceal all other signals, i.e., at*

[16] In the special case of constant-threshold equilibrium, history $\mathbf{h}^*$ is $\mathbf{1}_j$. However, our proof for *strict* monotone equilibrium does not apply to constant-threshold equilibrium as it requires players to have strict incentives (Lemma 1), whereas our result for constant-threshold equilibria does not require the equilibria to be strict.

[17] It seems hard to characterize *all* strict monotone equilibria when $\eta(\delta, p) < \eta_2$, since according to Theorem 1, multiple signals, including at least s_1 and s_2, can be disclosed on the equilibrium path, and a general monotone strategy allows agents' disclosure threshold to depend on the history and realized private signal.

on-path history $\mathbf{h} = (h_1, \ldots, h_l)$*, agents disclose* $\min\{n, h_1 + 1\}$ *copies of* s_1 *and nothing else when they observe private signal* s_1*, and disclose* $\min\{n, h_1\}$ *copies of* s_1 *otherwise.*

The proof is in Appendix A.2. Our theorem also leads to a characterization of all constant-threshold equilibria for any parameter configuration.

Corollary 2. *For any* $i \geq 2$*, if* $\eta(\delta, p) \in (\eta_i, \eta_{i-1}]$*, then for each* $j \in \{0, \ldots, i-1\}$*, there is a constant-threshold equilibrium in which agents conceal* $s_{j+1}, \ldots, s_l$ *and disclose all other signals, and for each* $j \geq i$*, there is no constant-threshold equilibrium in which agents disclose* $s_1, \ldots, s_j$*.*

We omit the proof of this corollary as it follows directly from the second part of Theorem 1. Focusing on the most informative constant-threshold equilibrium (i.e., the one in which the largest set of signals is disclosed), Corollary 2 seems to suggest a discontinuity at $\eta(\delta, p) = \eta_j$ since when $\eta(\delta, p) \leq \eta_j$, signal s_j is disclosed in the most informative equilibrium, whereas it can no longer be disclosed when $\eta(\delta, p) > \eta_j$. Proposition 2 shows that no such discontinuity exists.

Let $\mathbf{h}_{\delta,p}$ denote the steady-state history observed by an agent in the most informative constant-threshold equilibrium, which is a random variable that takes values in H. We measure the information contained in this history by its mutual information with the state. In particular, $I(\theta; \mathbf{h}_{\delta,p}) = \gamma(\pi_0) - \mathbb{E}\left[\gamma\left(\pi(\mathbf{h}_{\delta,p})\right)\right]$, where $\pi(\mathbf{h}) = \Pr(\theta = \overline{\theta} \mid \mathbf{h})$ is their posterior after observing $\mathbf{h}$ but before observing their private signal, and $\gamma(q) = -q \log q - (1-q) \log(1-q)$ is the entropy function. That is, $I(\theta; \mathbf{h}_{\delta,p})$ is the expected Kullback–Leibler divergence of the agent's posterior from his prior. Define $\mathcal{I}(\delta, p) = I(\theta; \mathbf{h}_{\delta,p})$.

Proposition 2. *There exists a continuous and strictly increasing function* $\widehat{\mathcal{I}} : (0, \infty) \to \mathbb{R}_+$ *such that* $\mathcal{I}(\delta, p) = \widehat{\mathcal{I}}\big(\eta(\delta, p)\big)$ *for every* $\delta \in (0, 1)$ *and* $p \in (0, 1]$*.*

The proof is in Appendix A.4. Proposition 2 shows that $\mathcal{I}$ is jointly continuous in (δ, p) and strictly increasing in each parameter, since $\eta(\delta, p)$ is strictly increasing in δ and p. The amount of information transmitted is continuous, since when $\eta(\delta, p) = \eta_j$, the number of copies of s_j disclosed does not affect the receiver's posterior belief about θ in the most informative constant-threshold equilibrium, as the positive effect of observing another s_j exactly cancels out the negative effect of expecting more bad signals.

In fact, our proof also implies that the information contained in the history improves in the Blackwell sense as δ or p increases, holding the other parameter fixed. Hence, young agents' expected welfare, $\mathbb{E}[u(\theta, a)]$, strictly increases in δ and p. However, because agents' actions generate a positive externality for their immediate predecessors, it is unclear whether an increase in δ or p improves social welfare. Indeed, coarsening young agents' information may induce them to take actions that increase old agents' expected welfare, $\mathbb{E}[v(\theta, a)]$.[18]

Since the amount of information contained in the disclosed history is shown to be strictly increasing in $\eta(\delta, p)$, it is natural to ask whether agents can learn the state with arbitrarily high precision as $\eta(\delta, p) \to \infty$. Proposition 3 shows that this is not the case. In particular, even when $p = 1$ and $\delta \to 1$, the limiting distribution over the agents' steady-state posterior beliefs in the most informative constant-threshold equilibrium is not the Dirac measure on 1 conditional on $\theta = \overline{\theta}$ and is not the Dirac measure on 0 conditional on $\theta = \underline{\theta}$.

Proposition 3. *Fix any $p \in (0, 1]$ and $\theta \in \Theta$. As $\delta \to 1$, the posterior belief $\pi(\mathbf{h}_{\delta,p})$ in the most informative constant-threshold equilibrium, conditional on θ, converges in distribution to a nondegenerate random variable that takes values strictly between 0 and 1 almost surely.*

The proof is in Appendix A.5. Proposition 3 implies that agents do not fully learn the state in the most informative constant-threshold equilibrium, even when $p = 1$ and $\delta \to 1$. Intuitively, although the expected number of disclosed copies of s_1 diverges to infinity as $\delta \to 1$, agents do not observe how many periods have elapsed since the latest communication failure. As a result, a large number of disclosed copies of s_1 may reflect either that s_1 is relatively likely, which suggests that the state is high, or that the time since the last communication failure is long. This uncertainty about the time since the last communication failure does not vanish as $\delta \to 1$. As a result, agents do not fully learn the state even in the limit.

4 Proof of Theorem 1

Section 4.1 establishes that, for any $j \geq 2$, $\eta(\delta, p) < \eta_j$ is necessary for s_j to be disclosed at any on-path history in any strict monotone equilibrium. Section 4.2 provides necessary and sufficient conditions for any $s_j \neq s_1$ to be disclosed in some (strict) constant-threshold equilibrium. We sometimes replace $\eta(\delta, p)$ with η in order to simplify notation.

[18]The logic is similar to that in Carrillo and Mariotti (2000), who study agents with time-inconsistent preferences. There, providing less information can induce agents to take actions that are better for their future selves.

4.1 Proof of Theorem 1: Part 1

We state two properties of strict equilibria. First, each agent tries to maximize his immediate successor's belief when deciding what to disclose. Since he can ignore his private signal and reveal only his history $\mathbf{h}$, from which he will induce the same belief when communication succeeds, the history he discloses in equilibrium must lead to a weakly higher belief than $\pi(\mathbf{h})$. This leads to an *improvement principle* similar to the one in Banerjee and Fudenberg (2004).

Improvement Principle. *For any equilibrium and any $t \in \mathbb{Z}$, if the realized history in period t is $\mathbf{h}^t$ and agent t chooses to disclose $\mathbf{h}^{t+1}$ when his history is $\mathbf{h}^t$, then $\pi(\mathbf{h}^{t+1}) \geq \pi(\mathbf{h}^t)$. If the equilibrium is strict and $\mathbf{h}^{t+1} \neq \mathbf{h}^t$, then $\pi(\mathbf{h}^{t+1}) > \pi(\mathbf{h}^t)$.*

An implication of our improvement principle is the following *no-turning-back property.*

No-Turning-Back Property. *In any strict equilibrium, for any $\mathbf{h}$ that occurs with positive probability and any $\mathbf{h}' < \mathbf{h}$, $\mathbf{h}'$ will not occur after $\mathbf{h}$ before the next communication failure.*

This is because whenever agents can disclose $\mathbf{h}$, they can disclose any $\mathbf{h}' < \mathbf{h}$, so their strict incentive to disclose $\mathbf{h}$ implies that $\pi(\mathbf{h}) > \pi(\mathbf{h}')$. Iterating this argument, we know that no agent will disclose $\mathbf{h}'$ after $\mathbf{h}$ until communication fails.

We now prove Theorem 1. Fix a strict monotone equilibrium $(\sigma, \underline{\mu}, \overline{\mu}, \pi)$ in which some signal is disclosed on the equilibrium path and let $\mu_\theta \in \Delta(H)$ denote the history distribution conditional on θ. Let $H_\sigma \subset H$ denote the set of histories that occur with positive probability under $\underline{\mu}$ (or, equivalently, under $\overline{\mu}$). Let j denote the largest integer in $\{1, \ldots, l\}$ such that there exists $\mathbf{h} = (h_1, \ldots, h_l) \in H_\sigma$ with $h_j \geq 1$ (if $H_\sigma = \{\mathbf{0}\}$, then set $j = 0$). By definition, s_j has the lowest likelihood ratio among the signals disclosed on the equilibrium path. Our gradualism property implies that there exists $\mathbf{h} \in H_\sigma$ with $h_j = 1$.

For any $j \geq 2$, we derive necessary conditions for strict monotone equilibria with s_j being the weakest signal disclosed on path to exist. Let $H^{**} \subset H_\sigma$ denote the set of positive-probability histories under σ such that the jth entry is 1, i.e., on-path histories that contain exactly one copy of s_j. Let $H^* \subset H^{**}$ denote the subset of H^{**} such that a history $\mathbf{h}$ belongs to H^* if and only if it minimizes $h_1 + \cdots + h_{j-1}$ among all histories in H^{**}. By definition, all histories in H^* have the same length, and hence, H^* is a finite set. Let $\mathbf{h}^* = (h_1^*, \ldots, h_l^*)$ denote an arbitrary solution to the minimization problem $\arg\min_{\mathbf{h} \in H^*} \pi(\mathbf{h})$. By construction, $h_{j+1}^* = \cdots = h_l^* = 0$ and $h_j^* = 1$. The following lemma shows that $\mathbf{h}^*$ can only originate from $\mathbf{h}^* - \mathbf{1}_j$.

Lemma 1. *In any strict monotone equilibrium, if an agent's history is* $\mathbf{h}^*$*, then his immediate predecessor's history is either* $\mathbf{h}^*$ *or* $\mathbf{h}^* - \mathbf{1}_j$*.*

The proof is in Appendix A.3. For any on-path history $\mathbf{h} \in H_\sigma$, let $D(\mathbf{h}) = \{i : \sigma_D(\mathbf{h}, s_i) \neq \mathbf{h}\} \subset \{1, 2, \ldots, l\}$. Intuitively, $i \in D(\mathbf{h})$ if and only if after an agent receives private signal s_i at history $\mathbf{h}$, he will disclose something different from $\mathbf{h}$, i.e., next period's history will not be $\mathbf{h}$ if communication succeeds. The no-turning-back property implies that for any $\mathbf{h} \in H_\sigma$, $i \in D(\mathbf{h})$ if and only if the disclosed vector after receiving private signal s_i at history $\mathbf{h}$ is not weakly less than $\mathbf{h}$, or equivalently, the number of copies of s_i disclosed is more than the ith entry of $\mathbf{h}$. Since $\mathbf{h}^* = (h_1^*, \ldots, h_{j-1}^*, 1, 0, \ldots, 0)$ and s_j is defined as the weakest signal disclosed on the equilibrium path, we have $\{1, \ldots, j-1\} \subset D(\mathbf{h}^*) \subset \{1, \ldots, j\}$. This is because if there exists $i < j$ such that $i \notin D(\mathbf{h}^*)$, agents disclose $\mathbf{h}^*$ upon receiving private signal s_i at history $\mathbf{h}^*$, in which case they disclose at least one copy of s_j while concealing at least one copy of s_i. This contradicts the monotonicity of agents' equilibrium strategy.

By Lemma 1, history $\mathbf{h}^*$ can occur only after $\mathbf{h}^*$ or $\mathbf{h}^* - \mathbf{1}_j$, which implies that

$$\mu_\theta(\mathbf{h}^*) = \mu_\theta(\mathbf{h}^* - \mathbf{1}_j) \cdot \delta p f_{\theta,j} \cdot \sum_{n=0}^{\infty} \delta^n (1 - p \sum_{i \in D(\mathbf{h}^*)} f_{\theta,i})^n = \mu_\theta(\mathbf{h}^* - \mathbf{1}_j) \cdot \frac{\delta p f_{\theta,j}}{1 - \delta(1 - p \sum_{i \in D(\mathbf{h}^*)} f_{\theta,i})}.$$

Hence, by Bayes rule, $\pi(\mathbf{h}^*) > \pi(\mathbf{h}^* - \mathbf{1}_j)$ if and only if

$$L_j \cdot \frac{1 - \delta(1 - p \sum_{i \in D(\mathbf{h}^*)} \underline{f}_i)}{1 - \delta(1 - p \sum_{i \in D(\mathbf{h}^*)} \overline{f}_i)} = L_j \cdot \frac{1 + \eta \sum_{i \in D(\mathbf{h}^*)} \underline{f}_i}{1 + \eta \sum_{i \in D(\mathbf{h}^*)} \overline{f}_i} > 1, \tag{4.1}$$

where the equality follows from the identity $1 - \delta(1 - px) = (1 - \delta)(1 + \eta x)$. Since $D(\mathbf{h}^*)$ is either $\{1, \ldots, j-1\}$ or $\{1, \ldots, j\}$, inequality (4.1) becomes either

$$L_j \cdot \frac{1 + \eta \sum_{i=1}^{j-1} \underline{f}_i}{1 + \eta \sum_{i=1}^{j-1} \overline{f}_i} > 1 \text{ or } L_j \cdot \frac{1 + \eta \sum_{i=1}^{j} \underline{f}_i}{1 + \eta \sum_{i=1}^{j} \overline{f}_i} > 1. \tag{4.2}$$

Applying (3.2) and (3.3), we know that either of the above two inequalities in (4.2) is satisfied if and only if $L_j > 1$ and $\eta < \eta_j$. Since η_j is strictly decreasing in j, it follows that when $L_j \leq 1$ or when $L_j > 1$ but $\eta \geq \eta_j$, there is also no strict monotone equilibrium in which the weakest signal disclosed has a likelihood ratio strictly less than L_j. This completes the proof of Part 1.

4.2 Proof of Theorem 1: Part 2

Suppose that in equilibrium, agents disclose $s_1, \ldots, s_j$ and conceal all other signals. Conditional on state θ, the probability of history $\mathbf{h} = (h_1, h_2, \ldots, h_j, 0, \ldots, 0)$ in the steady state is

$$(1-\delta)(\delta p)^{|\mathbf{h}|} \cdot X_\theta(\mathbf{h}) \cdot \prod_{i=1}^{j} (f_{\theta,i})^{h_i}, \tag{4.3}$$

where $|\mathbf{h}| = h_1 + \cdots + h_j$ denotes the length of history $\mathbf{h}$ and $X_\theta(\mathbf{h})$ denotes the coefficient in front of the term $\prod_{i=1}^{j}(q_i)^{h_i}$ in the following power series of $q_1, q_2, \ldots, q_j$:

$$\sum_{k=|\mathbf{h}|}^{\infty} \Big(q_1 + q_2 + \cdots + q_j + \delta(1 - p\sum_{i=1}^{j} f_{\theta,i})\Big)^k. \tag{4.4}$$

A series of algebraic manipulations reveals that

$$X_\theta(\mathbf{h}) = \Big(1 - \delta(1 - p\sum_{i=1}^{j} f_{\theta,i})\Big)^{-|\mathbf{h}|-1} \cdot \frac{|\mathbf{h}|!}{h_1! h_2! \ldots h_j!}. \tag{4.5}$$

Since $1 - \delta(1 - px) = (1-\delta)(1 + \eta x)$, the posterior likelihood ratio after observing $\mathbf{h}$ is

$$\frac{\pi(\mathbf{h})}{1 - \pi(\mathbf{h})} = \frac{\pi_0}{1 - \pi_0} \cdot \Big(\frac{1 + \eta \sum_{i=1}^{j} \underline{f}_i}{1 + \eta \sum_{i=1}^{j} \overline{f}_i}\Big)^{|\mathbf{h}|+1} \cdot \prod_{i=1}^{j} (L_i)^{h_i}. \tag{4.6}$$

Since $\eta > 0$ for any (δ, p) and $\sum_{i=1}^{j} \overline{f}_i > \sum_{i=1}^{j} \underline{f}_i$ for every $j < l$, we have $\frac{1+\eta\sum_{i=1}^{j} \underline{f}_i}{1+\eta\sum_{i=1}^{j} \overline{f}_i} < 1$. For a constant-threshold strategy in which agents only disclose $s_1, \ldots, s_j$ to be an equilibrium, a necessary condition is that agents weakly prefer to disclose $(h_1, \ldots, h_{j-1}, h_j, 0, \ldots, 0)$ to $(h_1, \ldots, h_{j-1}, h_j - 1, 0, \ldots, 0)$ for any $h_j \geq 1$ and $h_1, \ldots, h_{j-1} \geq 0$. By (4.6), this incentive constraint is satisfied if and only if

$$L_j \cdot \frac{1 + \eta \sum_{i=1}^{j} \underline{f}_i}{1 + \eta \sum_{i=1}^{j} \overline{f}_i} \geq 1. \tag{4.7}$$

For $j \geq 2$, expression (3.2) implies that (4.7) is equivalent to $\eta \leq \eta_j$. Similarly, for a constant-threshold strategy where agents only disclose $s_1, \ldots, s_j$ to be a strict equilibrium, we need

$$L_j \cdot \frac{1 + \eta \sum_{i=1}^{j} \underline{f}_i}{1 + \eta \sum_{i=1}^{j} \overline{f}_i} > 1 \quad \Leftrightarrow \quad \eta < \eta_j. \tag{4.8}$$

We next establish necessity for the existence of a constant-threshold equilibrium in which s_j is disclosed. Suppose $j \geq 2$ and that s_j is disclosed in a constant-threshold equilibrium with threshold k. Then $k \geq j$, and the preceding argument implies that $\eta_k \geq \eta$. Since $\eta > 0$, this requires $\eta_k > 0$, and therefore $L_k > 1$. Since $j \leq k$, we have $\eta_j \geq \eta_k \geq \eta$. If the constant-threshold equilibrium is strict, then $\eta_k > \eta$, and hence $\eta_j > \eta$.

We next show that disclosing only $s_1, \ldots, s_j$ is a constant-threshold equilibrium when $\eta \leq \eta_j$, and is a strict constant-threshold equilibrium when $\eta < \eta_j$. First, if (4.7) is satisfied,

$$L_k \cdot \frac{1 + \eta \sum_{i=1}^{j} \underline{f}_i}{1 + \eta \sum_{i=1}^{j} \overline{f}_i} > L_j \cdot \frac{1 + \eta \sum_{i=1}^{j} \underline{f}_i}{1 + \eta \sum_{i=1}^{j} \overline{f}_i} \geq 1 \text{ for every } k < j,$$

where the strict inequality follows from $L_k > L_j$. This implies that agents also have an incentive to disclose signals $s_1, \ldots, s_{j-1}$ when $\eta \leq \eta_j$. Similarly, when (4.8) is satisfied, agents will have strict incentives to disclose $s_1, \ldots, s_j$. Agents have no incentive to disclose $s_{j+1}, \ldots, s_l$ as long as their belief at any off-path history $\mathbf{h}$, i.e., those where $\mathbf{h} \neq \mathbf{h}[j]$, satisfies $\pi(\mathbf{h}) < \pi(\mathbf{0})$.[19]

5 Discussions and Extensions

Section 5.1 connects our monotonicity and constant-threshold refinements to those proposed in the disclosure literature. Section 5.2 discusses our modeling assumptions, such as the assumption that agents cannot observe the exact timestamps of the disclosed signals and that all disclosed signals are lost when communication fails. It also considers several extensions, such as when the successful transmission of different signals are not perfectly correlated, when agents face disclosure constraints, and when the communication success rate depends on the number of signals disclosed. Section 5.3 discusses non-strict equilibria and mixed-strategy equilibria. Section 5.4 explains the difference between our OLG game and an analogous one-shot disclosure game.

5.1 Alternative Equilibrium Refinements

All the constant-threshold equilibria in Corollary 2, as well as the equilibria that violate monotonicity in Proposition 1, are *Perfect Bayesian equilibria* as defined in Fudenberg and Tirole (1991), which requires players' beliefs to be derived from a conditional probability system induced by their equilibrium strategies.[20] This is because, when agents only disclose signals in

[19]Similarly, a strict constant-threshold equilibrium in which no signal is disclosed always exists, as long as we set agents' beliefs at any history $\mathbf{h} \neq \mathbf{0}$ to be strictly less than their prior belief about the state.

[20]Part 1 of Theorem 1 characterizes a common property of all strict monotone equilibria, so it trivially extends once we impose the conditional probability system refinement. Part 2 of Theorem 1 establishes the existence of

$S' \subset S$, any history that contains signals in $S \backslash S'$ occurs with zero probability conditional on any $(\mathbf{h}, s) \in H \times (S \cup \{\emptyset\})$, so the conditional probability system requirement has no bite.

Standard forward-induction refinements, such as the intuitive criterion and D1 in Cho and Kreps (1987), also have little bite, e.g., both monotone and non-monotone equilibria described in Proposition 1 satisfy those refinements. Intuitively, the sender's payoff is independent of the information he observes, so these refinements cannot distinguish among sender-types who disclose the same set of signals in equilibrium and differ only in the signals they conceal.

In a dynamic disclosure game with two periods and the sender having at most two signals, Guttman, Kremer, and Skrzypacz (2014) introduced a refinement on off-path beliefs called *belief monotonicity*. Their refinement requires that, at histories where only one signal is disclosed, the receiver's posterior belief is strictly higher when the disclosed signal has a higher likelihood ratio. We extend their refinement to our setting where senders can have arbitrarily many signals.

Monotone Belief. *A belief map $\pi : H \to [0,1]$ is monotone if $\pi(\mathbf{h}) > \pi(\mathbf{h}^*)$ for every pair of histories $\mathbf{h} = (h_1, \ldots, h_l)$ and $\mathbf{h}^* = (h_1^*, \ldots, h_l^*)$ that satisfy $\sum_{j=1}^{l} h_j = \sum_{j=1}^{l} h_j^*$ and*

$$\sum_{j=1}^{n} h_j \geq \sum_{j=1}^{n} h_j^* \text{ for every } n < l \text{ with strict inequality for some.} \tag{5.1}$$

Intuitively, a belief map is monotone if agents have a strictly higher belief about the state when their predecessors disclose one more higher-likelihood-ratio signal and one fewer lower-likelihood-ratio signal. As in Guttman et al. (2014), it only imposes restrictions on the receiver's posterior beliefs at pairs of histories with the same length, i.e., $\sum_{j=1}^{l} h_j = \sum_{j=1}^{l} h_j^*$.

Theorem 1 extends once we replace *monotone strategy* with *monotone belief*. This is because when the belief map is monotone, any non-monotone strategy has a strictly profitable deviation, by replacing a disclosed low-likelihood-ratio signal with a concealed high-likelihood-ratio signal. This implies that if s_j is never disclosed in any equilibrium with a monotone strategy, then it is also never disclosed in any equilibrium with a monotone belief map. Moreover, for each constant-threshold equilibrium we constructed in the proof of Theorem 1, there exists an off-path belief that satisfies agents' incentive constraints as well as the belief monotonicity requirement.[21]

a (strict) constant-threshold equilibrium, so imposing the conditional probability system refinement makes the result stronger. This is why our discussion in this paragraph focuses on constant-threshold equilibria.

[21]Similar to the off-path beliefs in Section 4.2, for every $\mathbf{h}$ that occurs with zero probability in equilibrium, let $\pi(\mathbf{h}) < \pi(\mathbf{0})$. For any $n \in \mathbb{N}$, there is a finite number of off-path histories with length n, which means that there exists a mapping π from length-n off-path histories to the open interval $(0, \pi(\mathbf{0}))$ that satisfies (5.1).

Guttman et al. (2014) also proposed a *threshold strategy* refinement, which requires the sender to disclose a signal if and only if its likelihood ratio is above some cutoff. Extending it to our setting leads to a refinement that is more restrictive than monotone equilibrium but is more permissive than constant-threshold equilibrium as the threshold can be history-dependent.

Threshold Strategy & Threshold Equilibrium. *A strategy σ is a threshold strategy if there exists a function $J : H \times (S \cup \{\emptyset\}) \to \{0, 1, \ldots, l\}$ such that for every $\mathbf{h} \in H$ and $i \in \{1, 2, \ldots, l\}$, we have $\sigma_D(\mathbf{h}, s_i) = (\mathbf{h}+\mathbf{1}_i)[J(\mathbf{h}, s_i)]$ and $\sigma_D(\mathbf{h}, \emptyset) = \mathbf{h}[J(\mathbf{h}, \emptyset)]$. An equilibrium is a threshold equilibrium if agents use a threshold strategy.*

Our results for strict monotone equilibria extend to strict threshold equilibria. Part 1 of Theorem 1 extends since the set of strict monotone equilibria includes the set of strict threshold equilibria. Part 2 of Theorem 1 extends because, when $\eta(\delta, p) < \eta_j$, there exists a constant-threshold equilibrium where s_j is disclosed, which by definition, must also be a threshold equilibrium.

Hart, Kremer, and Perry (2017) introduced a truth-leaning refinement, which, in our setting, requires that (A0) an agent will disclose all available signals if doing so maximizes his payoff and (P0) agents believe that their predecessors did not hide any information upon observing any $\mathbf{h} \in H$ that occurs with zero probability in equilibrium. See page 700 of their paper for details.

We explain why there is no strict equilibrium that satisfies both our monotonicity refinement and their (P0) requirement when $L_2 > 1$ and $\eta(\delta, p) > \eta_2$. Suppose by way of contradiction that there exists an equilibrium that satisfies both refinements. Our theorem implies that in every strict monotone equilibrium, agents will disclose $\mathbf{0}$ when their history is $\mathbf{0}$ and their private signal is s_2, so $\pi(\mathbf{0}) \geq \pi(\mathbf{1}_2)$. Since agents can always disclose $\mathbf{0}$, the probability that agents' belief assigns to $\overline{\theta}$ upon observing history $\mathbf{0}$ is no more than their prior, so $\pi_0 \geq \pi(\mathbf{0})$. If such an equilibrium also satisfies (P0), then given that history $\mathbf{1}_2$ occurs with zero probability, $\pi(\mathbf{1}_2)$ must equal the agent's posterior belief when they know that their predecessor's history is $\mathbf{0}$ and their private signal is s_2. Since $L_2 > 1$, we have $\pi(\mathbf{1}_2) > \pi(\mathbf{0})$, which contradicts $\pi(\mathbf{0}) \geq \pi(\mathbf{1}_2)$.

Lastly, our refinements are neither stronger nor weaker than the receiver-optimal equilibrium used in Rappoport (2025). First, equilibria where the sender discloses nothing satisfy both of our refinements, but they are dominated by the one where s_1 is disclosed in terms of receiver welfare. Second, suppose for example $l = 3$, $L_1 > L_2 > 1 > L_3$, and $\overline{f}_1$ and $\underline{f}_1$ are much smaller than $\overline{f}_2$ and $\underline{f}_2$. When $\eta(\delta, p) > \eta_2$, signals s_2 and s_3 are never disclosed in any strict monotone equilibrium. However, the receiver will be strictly better off in the non-monotone equilibrium where only signal s_2 is disclosed, because signal s_2 occurs with much higher probability than s_1. According to our Proposition 1, the latter equilibrium exists regardless of the parameters.

5.2 Discussion of Modeling Assumptions and Extensions

Imperfect Correlation in Signal Transmission: Our baseline model assumes that an agent receives either all signals disclosed by his predecessor, or none of them. Alternatively, one can also consider settings where each disclosed signal is transmitted independently. For example, if an agent discloses two signals, then his successor receives both with probability δ^2, exactly one with probability $2\delta(1-\delta)$, and neither with probability $(1-\delta)^2$.

In this alternative setting, for every $\delta \in (0,1)$ and $p \in (0,1]$, there exists a constant-threshold equilibrium in which all signals with likelihood ratios greater than 1 are disclosed, which contrasts with our setting where only the strongest possible signal can be disclosed when δ and p are large.

To see why, suppose the constant disclosure threshold is j, where $L_j > 1$. Conditional on state θ, agent $t \in \mathbb{Z}$ fails to observe the signal realized in period $\tau(< t)$ with probability $1 - \delta^{t-\tau} p \sum_{i=1}^{j} f_{\theta,i}$. This can occur for three reasons: the signal is lost during transmission, agent τ fails to observe his private signal, or the realized signal has a likelihood ratio below L_j and is therefore concealed. Because signals are transmitted independently, disclosing s_j both reveals one additional realization of s_j and lowers the expected number of concealed bad signals. When $L_j > 1$, both effects increase the next agent's posterior belief.

By contrast, in our model, where either all disclosed signals are successfully transmitted or all are lost, agents can infer only the sequence of realized signals following the most recent communication failure, which occurred in some random period τ. Observing one additional disclosed good signal leads agents to infer that τ is more likely to be earlier, implying that the sequence of signals following τ contains more bad signals in expectation. Thus, disclosing good signals has a negative effect that is absent in the model with independent transmission.

We now extend our main result to a hybrid setting that allows both signal-specific and aggregate transmission failures. For each sufficiently small $\varepsilon > 0$, let $\delta_\varepsilon = 1 - a\varepsilon$ and $\rho_\varepsilon = 1 - b\varepsilon$, where $a, b \geq 0$ are constants. Suppose that, after each agent discloses a set of signals, an aggregate shock occurs with probability $1 - \delta_\varepsilon$, in which case all disclosed signals are lost. Conditional on no aggregate shock, each disclosed signal is independently transmitted to the successor with probability ρ_ε. Our baseline model corresponds to $a > 0$ and $b = 0$, while the independent transmission setting discussed above corresponds to $a = 0$ and $b > 0$.

Proposition 4. *Suppose $a > 0$ and $b \in [0, a]$. There exists $\varepsilon^* > 0$ such that, when $\varepsilon \in (0, \varepsilon^*)$:*

1. *Signals $s_2, ..., s_l$ are never disclosed in any constant-threshold equilibrium.*
2. *There exists a constant-threshold equilibrium in which only s_1 is disclosed.*

The proof is in Appendix A.6. Proposition 4 partially extends Theorem 1 by showing that, as long as the probability of an aggregate shock is at least as large as that of a signal-specific shock, agents become extremely selective in disclosing information as communication frictions vanish (i.e., $\varepsilon \to 0$), in the sense that only the best possible signal s_1 can be disclosed in equilibrium. This result suggests that the mechanism underlying our main result remains relevant when the successful transmission of different signals has enough positive correlation.

Timestamps: We assume that agents can only observe the number of each signal disclosed but cannot directly observe (i) the exact date at which each disclosed signal was realized and (ii) the time at which communication failed in the past. Our assumption fits when senders (e.g., managers who can edit training materials presented to new recruits) can credibly disclose the substantive content of the signals (e.g., fabricating cases is hard due to organizations' internal review processes), but find it hard to reveal the exact time at which the signals arrived due to, for example, privacy regulations that prevent them from revealing too much identifying information.

Alternatively, one may consider settings in which agents cannot credibly disclose hard information, so communication takes the form of cheap talk. Such an environment is studied by Anderlini, Gerardi, and Lagunoff (2012) in an OLG model where the sender's payoff is single-peaked in the receiver's action. By contrast, we focus on environments where the sender's payoff is strictly increasing in the receiver's action, in which case cheap talk is uninformative. Our assumption on the sender's payoff seems reasonable when the receiver's action is interpreted as his costly effort, which can strictly benefit the sender (e.g., the sender is the receiver's manager).

There are also situations where agents can credibly disclose the signals' arrival times in addition to their substantive content. In the special case where each agent observes his private signal for sure ($p = 1$), observing the exact time at which communication failed in the past leads to unraveling, since each agent can infer the exact number of bad signals concealed.

If agents can observe the timestamps of disclosed signals but not the time at which communication failed in the past, unraveling may not happen since there is still uncertainty regarding the number of realized signals after the latest communication breakdown, which, depending on the equilibrium strategies, could be correlated with the set of signals disclosed.

In the more general case where $p < 1$, the absence of a disclosed signal in a given period may also be caused by the agent in that period failing to observe his private signal. Consequently, agents remain uncertain about the number of signals observed since the most recent communication breakdown, even when they can observe the timestamps of the disclosed signals or the time at which communication failed in the past. So, depending on the equilibrium strategies, disclosing an additional good signal may still lead the subsequent agent to infer a larger expected number of bad signals, so the forces identified in our analysis may continue to operate. A related question is: What if agents can voluntarily disclose the timestamps of the disclosed signals? It is unclear whether agents have incentives to disclose the timestamps, which may depend on the equilibrium being played. Characterizing how changes in δ, p, and $\mathbf{f}$ affect the selectiveness of disclosure in richer environments is an interesting direction for future research.

General payoffs: Theorem 1 extends when each agent t receives s^t after choosing his action a^t, as it only hinges on agents' beliefs after observing the set of disclosed signals. Similarly, our result extends when agents have heterogeneous payoffs, as long as they are supermodular with respect to (θ, a^t), are separable between a^t and a^{t+1}, and are strictly increasing in a^{t+1}.

By contrast, new complications arise when agents' payoffs are not separable between a^t and a^{t+1}. In this case, an agent's action depends not only on his belief about the state, but also on his belief about his successor's action. As a result, when agent t decides which set of signals to disclose, his objective is not only to maximize agent $t+1$'s belief, but also to equip agent $t+1$ with signals to increase (or decrease) agent $t+2$'s belief, depending on whether agents' actions are complements or substitutes. A similar complication arises when agent t's payoff depends also on $\{a^{t+2}, a^{t+3}, \dots\}$,[22] or when the distribution of agent t's private signal s^t depends on agent t's action. Whether our theorem extends to these more general cases remains an open question.

[22] The classic OLG model of Samuelson (1958), and the subsequent works of Bhaskar (1998) and Acemoglu and Jackson (2015) all assume that each agent's payoff depends only on his own action and that of his immediate successor. Allowing agents' payoffs to depend on their predecessors' actions does not change our results. This is because agents cannot affect their predecessors' actions, so such dependence does not affect their incentives.

Disclosure Constraints: First, Part 1 of Theorem 1 extends when agents face disclosure constraints. For example, suppose each agent can disclose at most $N \in \mathbb{N}$ signals.[23] That is, agents can disclose $(h_1, \ldots, h_l)$ only if $\sum_{j=1}^{l} h_j \leq N$. This is because the gradualism property, the improvement principle, and the no-turning-back property all continue to apply. Lemma 1 also extends since it only rules out the possibility that the current-period history is $\mathbf{h}^*$ yet the previous-period history is neither $\mathbf{h}^*$ nor $\mathbf{h}^* - \mathbf{1}_j$, and such a possibility cannot re-appear due to disclosure constraints. The calculation after Lemma 1 compares agents' posterior beliefs at $\mathbf{h}^*$ and at $\mathbf{h}^* - \mathbf{1}_j$, which is also unaffected.

Heterogeneous Communication Success Rate: We consider an extension in which the communication success rate depends on the number of signals disclosed. Suppose when agent t discloses a set of signals $\mathbf{h}$, the probability that agent $t+1$ receives $\mathbf{h}$ is $\delta(\mathbf{h}) = \exp(-\Delta \cdot q(\mathbf{h}))$ where $\Delta > 0$ is a parameter and $q : H \to [\underline{q}, \overline{q}]$ for some $0 < \underline{q} < \overline{q} < \infty$ such that $q(\mathbf{h})$ depends on $\mathbf{h}$ only through its length and $q(\mathbf{h})$ weakly increases as the length of $\mathbf{h}$ increases. In this setting, communication is weakly more likely to fail when more signals are disclosed. As $\Delta \to 0$, the communication success rate converges to 1.

When Δ is below some cutoff, we show that no signal other than s_1 is ever disclosed in any strict monotone equilibrium. We explain at the end of Appendix A.3 why Lemma 1 extends to this more general setting. Now suppose by way of contradiction that no matter how small Δ is, there exists a strict monotone equilibrium in which the weakest signal disclosed on path is $s_j \neq s_1$. By Lemma 1, $\mathbf{h}^*$ can only occur after $\mathbf{h}^*$ or $\mathbf{h}^* - \mathbf{1}_j$, so we have the following equation that relates the probability of history $\mathbf{h}^*$ conditional on θ to the probability of history $\mathbf{h}^* - \mathbf{1}_j$ conditional on θ:

$$\mu_\theta(\mathbf{h}^*) = \mu_\theta(\mathbf{h}^* - \mathbf{1}_j) \cdot \delta(\mathbf{h}^*) \cdot p f_{\theta,j} \cdot \sum_{n=0}^{\infty} \delta(\mathbf{h}^*)^n \cdot (1 - p \sum_{i \in D(\mathbf{h}^*)} f_{\theta,i})^n, \tag{5.2}$$

where $D(\mathbf{h}^*) \subset \{1, 2, \ldots, l\}$ is the index set of private signals such that the next period history is no longer $\mathbf{h}^*$ once an agent receives such a private signal at $\mathbf{h}^*$. As in the baseline model, $D(\mathbf{h}^*)$ is either $\{1, \ldots, j-1\}$ or $\{1, \ldots, j\}$. Equation (5.2) implies that $\pi(\mathbf{h}^*) \geq \pi(\mathbf{h}^* - \mathbf{1}_j)$ if and only if

$$L_j \cdot \frac{1 - \delta(\mathbf{h}^*)(1 - p \sum_{i \in D(\mathbf{h}^*)} \underline{f}_i)}{1 - \delta(\mathbf{h}^*)(1 - p \sum_{i \in D(\mathbf{h}^*)} \overline{f}_i)} \geq 1. \tag{5.3}$$

[23] Constant-threshold equilibria with threshold $j \geq 1$ are not well-defined under disclosure constraints, since agents cannot disclose $N+1$ copies of s_1 after receiving s_1 at history $N\mathbf{1}_1$. So Part 2 does not extend as stated.

Since $\delta(\mathbf{h}^*) = \exp(-\Delta \cdot q(\mathbf{h}^*))$ and the function $q(\cdot)$ is uniformly bounded above and below, we know that there exists $\Delta^* > 0$ such that when $\Delta < \Delta^*$, inequality (5.3) fails for all j and $D(\mathbf{h}^*)$ that satisfy $j \geq 2$ and $j \geq \max D(\mathbf{h}^*)$. When $\pi(\mathbf{h}^*) < \pi(\mathbf{h}^* - \mathbf{1}_j)$, agents strictly prefer to disclose $\mathbf{h}^* - \mathbf{1}_j$ to $\mathbf{h}^*$. This is because disclosing $\mathbf{h}^* - \mathbf{1}_j$ results in a weakly higher probability of successful communication and induces a higher belief conditional on successful communication.

5.3 Non-Strict Equilibria & Mixed-Strategy Equilibria

As in Clark et al. (2021), Part 1 of Theorem 1 focuses on *strict* equilibria. We do not know whether this result extends to all pure-strategy equilibria, in particular, those where agents do not have strict incentives. This is because the *no-turning-back property* stated in Section 4.1 breaks down, so Lemma 1 and the calculations that follow no longer apply.

To see intuitively where our argument breaks down, let s_j denote the weakest signal disclosed on path and recall the definition of $\mathbf{h}^*$ in Section 3. Suppose $\pi(\mathbf{h}^*) = \pi(\mathbf{h}^* - \mathbf{1}_j)$, so that agents are indifferent between revealing and concealing s_j upon receiving it at $\mathbf{h}^* - \mathbf{1}_j$. They may then disclose $\mathbf{h}^*$ upon receiving s_j at $\mathbf{h}^* - \mathbf{1}_j$, while disclosing $\mathbf{h}^* - \mathbf{1}_j$ upon receiving some bad signal s_i, with $i > j$, at $\mathbf{h}^*$. This, in turn, raises the next agent's belief upon observing history $\mathbf{h}^*$, since observing $\mathbf{h}^*$ implies that the bad signal s_i did not arrive after the disclosure of s_j. Consequently, the argument leading to the contradiction $\pi(\mathbf{h}^*) < \pi(\mathbf{h}^* - \mathbf{1}_j)$ no longer goes through.

Nevertheless, in an extension we considered in Section 5.2 where each agent can disclose at most $N \in \mathbb{N}$ signals, all pure-strategy equilibria are outcome-equivalent to some strict equilibria under generic parameter values. This is because for any finite N, the set of histories is finite, as is the set of signal profiles that can be disclosed. If two histories lead to the same expected payoff for the sender, then the indifference condition defines a polynomial equation in δ. Since there is a finite number of signal profiles, there can be at most a finite number of values of δ for every $(u, v, \mathbf{f}, \pi_0, p)$ under which the sender is indifferent. Hence, the set of parameters under which non-strict pure-strategy equilibria exist has Lebesgue measure zero.

We now turn to mixed-strategy equilibria. Since we assumed that there is a unique $a \in \mathbb{R}$ that maximizes $\pi u(\overline{\theta}, a) + (1 - \pi) u(\underline{\theta}, a)$ for every $\pi \in [0, 1]$, agents always have strict incentives when choosing their actions. Consequently, the only source of randomization is their disclosure decisions. As discussed earlier, we do not know whether our result on monotone equilibria extends to mixed-strategy equilibria. We focus on a class of mixed-strategy equilibria that are tractable to analyze and are natural generalizations of constant-threshold equilibria:

Constant-Threshold Mixed Strategy. *A constant-threshold mixed strategy is characterized by $j \in \{1, \ldots, l\}$ and $q \in (0,1)$ such that agents disclose all available copies of $s_1, \ldots, s_{j-1}$, disclose all copies of s_j in their histories, disclose their private signal s_j with probability q, and never disclose any copy of $s_{j+1}, \ldots, s_l$.*

Intuitively, this class of strategies requires agents to disclose all signals with likelihood ratio strictly above a cutoff as well as all inherited signals with likelihood ratio at the cutoff, to conceal all signals with likelihood ratio strictly below the cutoff, and to randomize, with probability strictly between 0 and 1, over disclosing private signals with likelihood ratio at the cutoff.

One can show that for any $j \in \{1, \ldots, l\}$ and $q \in (0,1)$, a constant-threshold mixed strategy equilibrium parameterized by (j, q) exists if and only if (i) $j \geq 2$, (ii) $L_j > 1$, and (iii) $\eta(\delta, p) = \eta_j$. This is because agents are indifferent between disclosing private signal s_j and concealing it if and only if

$$L_j = \frac{1 - \delta(1 - p(q\overline{f_j} + \sum_{i=1}^{j-1} \overline{f_i}))}{1 - \delta(1 - p(q\underline{f_j} + \sum_{i=1}^{j-1} \underline{f_i}))}. \tag{5.4}$$

Regardless of $q \in (0,1)$, the RHS of (5.4) is strictly greater than 1 and strictly less than L_1, which leads to the requirements that $j \geq 2$ and $L_j > 1$. Furthermore, regardless of q, equation (5.4) holds if and only if $\eta(\delta, p) = \eta_j$. Hence, the class of mixed-strategy equilibria we considered exists only under knife-edge parameters. When agents also randomize over the cutoff signals contained in their histories, or when their disclosure rule depends non-trivially on their information set, the marginal effect of disclosing another s_j depends on the entire set of signals disclosed. Characterizing such mixed-strategy equilibria is left for future research.

5.4 OLG Game versus One-Shot Disclosure Game

To understand the role of our OLG structure, we consider a one-shot disclosure game between a sender and a receiver in which the receiver does not know the number of signals the sender has.

Suppose the state is binary $\theta \in \{\overline{\theta}, \underline{\theta}\}$. The sender observes $\widetilde{n}$ conditionally i.i.d. signals $\{s^t\}_{t=1}^{\widetilde{n}}$, where $s^t \in \{s_1, \ldots, s_l\}$ and $\Pr(\widetilde{n} = n) = (1-\delta)\delta^n$ for every $n \geq 0$. We assume that signal realizations with lower indices have higher likelihood ratios. The sender discloses a subset of the signals in order to maximize the probability that the receiver's posterior belief assigns to $\overline{\theta}$. The receiver only observes the number of each signal the sender discloses. Hence, both players' information sets belong to $H = \mathbb{N}^l$. This is analogous to our OLG game where $p = 1$.

The sender's strategy is $\sigma_D : \mathbb{N}^l \to \mathbb{N}^l$ subject to $\sigma_D(\mathbf{h}) \leq \mathbf{h}$ for every $\mathbf{h} \in \mathbb{N}^l$. We extend the notions of monotone and constant-threshold strategies to this setting. The sender's strategy is *monotone* if for every $\mathbf{h} = (h_1, \ldots, h_l)$ and $j \in \{2, \ldots, l\}$, if the jth entry of $\sigma_D(\mathbf{h})$ is not 0, then the ith entry of $\sigma_D(\mathbf{h})$ equals h_i for every $i < j$.[24] The sender's strategy is a *constant-threshold strategy* if there exists $j \in \{0, 1, \ldots, l\}$ such that $\sigma_D(\mathbf{h}) = \mathbf{h}[j]$ for every $\mathbf{h} \in H$.

Our results for (strict) constant-threshold equilibria, Part 2 of Theorem 1 and Corollary 2, extend to this one-shot game once we set $p = 1$. The intuition remains the same: suppose in equilibrium, the sender uses a constant-threshold strategy of disclosing $s_1, \ldots, s_i$ (good signals) and concealing $s_{i+1}, \ldots, s_l$ (bad signals). Disclosing another good signal shows to the receiver not only that there is one more good signal, but also there are more bad signals in expectation, so the sender's disclosure incentives depend on the relative magnitudes of the two effects.

However, Part 1 of Theorem 1, which concerns strict monotone equilibria, does not extend to this one-shot game, thereby distinguishing this game from our OLG game. To highlight the difference, we introduce a class of strategies parameterized by a positive integer $k \in \mathbb{N}$ where s_2 is sometimes disclosed. We show that no matter how close δ is to 1, there always exists $k \in \mathbb{N}$ such that Strategy k is part of a strict monotone equilibrium in the one-shot game:

- Strategy k: For every $j \neq k + 1$, if the sender observes j copies of signal s_1, then he discloses all the s_1 he has and nothing else. If the sender has $k + 1$ copies of s_1 but no s_2, then he discloses k copies of s_1 and nothing else. If the sender has $k + 1$ copies of s_1 and at least one copy of s_2, then he discloses $k + 1$ copies of s_1, one copy of s_2, and nothing else.

Proposition 5. *For any $\delta \in (0, 1)$, there exists $k \in \mathbb{N}$ such that there exists a strict monotone equilibrium in the one-shot disclosure game where the sender uses Strategy k.*[25]

The proof is in Appendix A.7. Compared to our OLG game, the gradualism property no longer holds in the one-shot disclosure game, so Lemma 1 is no longer true. To see intuitively

[24] Our proposition in this section extends if we replace strategy monotonicity with belief monotonicity, that the receiver's belief increases when the sender replaces a lower-likelihood-ratio signal with a higher one.

[25] Strategy k is *not* a threshold strategy since only one copy of s_2 is disclosed when the sender has $k + 1$ copies of s_1 and 2 copies of s_2 available. In fact, one can show that there exists $\delta^* \in (0, 1)$ such that when $\delta > \delta^*$, the sender will never disclose $s_2, \ldots, s_l$ in any strict threshold equilibrium in the one-shot disclosure game.

where the argument for Theorem 1 breaks down, note that under Strategy k, s_2 is the weakest signal disclosed on the equilibrium path, so history $\mathbf{h}^*$ is $(k+1)\mathbf{1}_1 + \mathbf{1}_2$, and history $\mathbf{h}^* - \mathbf{1}_j$ is $(k+1)\mathbf{1}_1$. Since the sender never discloses $(k+1)\mathbf{1}_1$ on the equilibrium path, when he has $k+1$ copies of s_1 and one copy of s_2 available, his decision regarding whether to disclose s_2 depends on the comparison between the receiver's posterior beliefs at $k\mathbf{1}_1$ and at $(k+1)\mathbf{1}_1 + \mathbf{1}_2$, instead of between the receiver's beliefs at $(k+1)\mathbf{1}_1$ and at $(k+1)\mathbf{1}_1 + \mathbf{1}_2$. If the sender discloses $(k+1)\mathbf{1}_1 + \mathbf{1}_2$ instead of $k\mathbf{1}_1$, he shows to the receiver that (i) there is one more s_1, (ii) there is one more s_2, and (iii) there is no more s_1 beyond the $k+1$ copies disclosed. In our OLG game, the first effect, which is in favor of disclosing s_2, is absent. Instead of disclosing $(k+1)\mathbf{1}_1 + \mathbf{1}_2$, a more attractive option is to disclose $(k+1)\mathbf{1}_1$, which happens on the equilibrium path thanks to the gradualism property. It shows to the receiver that there are $k+1$ copies of s_1 without showing her that there is also an s_2 as well as some bad signals between the disclosed signals.

6 Conclusion

We studied a model in which disclosing favorable evidence triggers skepticism, because more disclosure implies more opportunities to conceal unfavorable evidence. In our overlapping-generations environment, greater opportunities to observe private signals and more reliable information transmission expand the set of evidence that could be hidden behind a favorable disclosed record. This amplifies the adverse inference associated with disclosing good signals and, consequently, reduces the set of signals that can be revealed in equilibrium. As communication frictions vanish, only the strongest possible signal survives, since its direct evidentiary effect always dominates the adverse inference it triggers.

A Appendix: Omitted Proofs

A.1 Uniqueness of the Steady State

We show that conditional on state θ, any pure strategy σ induces a unique steady-state distribution over the agents' histories.[26] Let $\sigma_D : H \times (S \cup \{\emptyset\}) \to H$ denote the component of σ that maps an agent's history and private signal (if any) to the set of signals disclosed to the next agent. For any $\mathbf{h}, \mathbf{h}' \in H$, define the probability of moving from history $\mathbf{h}$ to history $\mathbf{h}'$ conditional on the selected history being successfully observed by the next generation as $Q^\sigma_\theta(\mathbf{h}, \mathbf{h}') = (1-p)\mathbb{1}_{\sigma_D(\mathbf{h},\emptyset)=\mathbf{h}'} + p\sum_{i=1}^{l} f_{\theta,i}\mathbb{1}_{\sigma_D(\mathbf{h},s_i)=\mathbf{h}'}$ and observe that it defines a Markov process. We then define the probability of transitioning from history $\mathbf{h}$ to history $\mathbf{h}'$ as

$$P^\sigma_\theta(\mathbf{h}, \mathbf{h}') = (1-\delta)\mathbb{1}_{\mathbf{h}'=\mathbf{0}} + \delta Q^\sigma_\theta(\mathbf{h}, \mathbf{h}'). \tag{A.1}$$

We show existence and uniqueness of a steady-state distribution using the Banach fixed point theorem. Let $T^\sigma_\theta : \Delta(H) \to \Delta(H)$ be the transition operator implied by P^σ_θ, so that $(T^\sigma_\theta(\mu))(\mathbf{h}') = \sum_{\mathbf{h}\in H}\mu(\mathbf{h})P^\sigma_\theta(\mathbf{h}, \mathbf{h}')$. We show that this is a contraction mapping when equipped with the total variation distance.

To see that applying only the Q^σ_θ transition does not increase total variation distance, take any $\mu, \mu' \in \Delta(H)$, and observe that

$$\begin{aligned}
&\Big\|\sum_{\mathbf{h}}\mu(\mathbf{h})Q^\sigma_\theta(\mathbf{h},\cdot) - \sum_{\mathbf{h}}\mu'(\mathbf{h})Q^\sigma_\theta(\mathbf{h},\cdot)\Big\|_{\mathrm{TV}} = \frac{1}{2}\sum_{\mathbf{h}'\in H}\Big|\sum_{\mathbf{h}\in H}\big(\mu(\mathbf{h}) - \mu'(\mathbf{h})\big)Q^\sigma_\theta(\mathbf{h},\mathbf{h}')\Big| \\
&\le \frac{1}{2}\sum_{\mathbf{h}'\in H}\sum_{\mathbf{h}\in H}\big|\mu(\mathbf{h}) - \mu'(\mathbf{h})\big|\,Q^\sigma_\theta(\mathbf{h},\mathbf{h}') = \frac{1}{2}\sum_{\mathbf{h}\in H}\big|\mu(\mathbf{h}) - \mu'(\mathbf{h})\big|\sum_{\mathbf{h}'\in H}Q^\sigma_\theta(\mathbf{h},\mathbf{h}') \\
&= \frac{1}{2}\sum_{\mathbf{h}\in H}\big|\mu(\mathbf{h}) - \mu'(\mathbf{h})\big| = \|\mu - \mu'\|_{\mathrm{TV}}.
\end{aligned}$$

By (A.1) and the above inequality, we have that for any $\mu, \mu' \in \Delta(H)$, $\|T^\sigma_\theta(\mu) - T^\sigma_\theta(\mu')\|_{\mathrm{TV}} \le \delta\,\|\mu - \mu'\|_{\mathrm{TV}}$. Since $\delta \in (0,1)$, T^σ_θ is a contraction on $(\Delta(H), \|\cdot\|_{\mathrm{TV}})$. By the Banach fixed point theorem, there exists a unique steady-state distribution over histories $\mu_\theta \in \Delta(H)$.

A.2 Proof of Proposition 1 and Corollary 1

We show that for any $n \in \mathbb{N} \cup \{\infty\}$ and $i \in \{1, \ldots, l-1\}$ that satisfies $L_i > 1$, there exists a strict equilibrium where agents disclose up to n copies of s_i and nothing else.

[26] The proof also applies to mixed strategies by defining $P^\sigma_\theta(\mathbf{h}, \mathbf{h}') = (1-\delta)\mathbb{1}_{\mathbf{h}'=\mathbf{0}} + \delta\big((1-p)\sigma_D(\mathbf{h},\emptyset)(\mathbf{h}') + p\sum_{i=1}^{l} f_{\theta,i}\sigma_D(\mathbf{h},s_i)(\mathbf{h}')\big)$, where $\sigma_D(\mathbf{h},s)(\mathbf{h}')$ denotes the probability of disclosing $\mathbf{h}'$ at information set $(\mathbf{h}, s)$.

Suppose in equilibrium, all agents use the strategy of disclosing no signal other than s_i and disclosing up to n copies of s_i. For any $\mathbf{h}$ that occurs with zero probability under this strategy, let $\pi(\mathbf{h}) < \pi(\mathbf{0})$. Then conditional on state θ, for every $m < n$, the probability of history $m\mathbf{1}_i$ is $(1-\delta)\delta^m \cdot (1-\delta(1-pf_{\theta,i}))^{-(m+1)} \cdot (pf_{\theta,i})^m$. Therefore,

$$\frac{\pi(m\mathbf{1}_i)}{1-\pi(m\mathbf{1}_i)} = \frac{\pi_0}{1-\pi_0} \cdot \Big(\frac{1-\delta(1-p\underline{f}_i)}{1-\delta(1-p\overline{f}_i)}\Big)^{m+1} \cdot L_i^m. \tag{A.2}$$

If $L_i > 1$, then the value of (A.2) is strictly increasing in m regardless of δ. The probability of history $n\mathbf{1}_i$ is $\sum_{k=n}^{\infty}(1-\delta)\delta^k \cdot (1-\delta(1-pf_{\theta,i}))^{-(k+1)} \cdot (pf_{\theta,i})^k$, so

$$\begin{aligned}
\frac{\pi(n\mathbf{1}_i)}{1-\pi(n\mathbf{1}_i)} &\geq \frac{\pi_0}{1-\pi_0} \cdot \min_{k\geq n}\Big\{\frac{(1-\delta)\,\delta^k\big(1-\delta(1-p\overline{f}_i)\big)^{-(k+1)}\overline{f}_i^{\,k}}{(1-\delta)\,\delta^k\big(1-\delta(1-p\underline{f}_i)\big)^{-(k+1)}\underline{f}_i^{\,k}}\Big\} \\
&\geq \frac{\pi_0}{1-\pi_0} \cdot \frac{(1-\delta)\,\delta^n\big(1-\delta(1-p\overline{f}_i)\big)^{-(n+1)}\overline{f}_i^{\,n}}{(1-\delta)\,\delta^n\big(1-\delta(1-p\underline{f}_i)\big)^{-(n+1)}\underline{f}_i^{\,n}} \\
&= \frac{\pi_0}{1-\pi_0} \cdot \left(\frac{1-\delta(1-p\underline{f}_i)}{1-\delta(1-p\overline{f}_i)}\right)^{n+1} L_i^n.
\end{aligned} \tag{A.3}$$

The RHS of (A.3) is strictly greater than the value of (A.2) for any $m < n$. These calculations together with our earlier specifications of off-path beliefs verify that agents have strict incentives to disclose up to n copies of s_i and nothing else, which establishes Proposition 1.

For Corollary 1, let $i = 1$ and note that when $\eta(\delta, p) \geq \eta_2$, Theorem 1 implies that only s_1 can be disclosed on the equilibrium path. Therefore, every strict monotone equilibrium is characterized by $n \in \mathbb{N} \cup \{\infty\}$ such that agents disclose up to n copies of s_1 and nothing else.

A.3 Proof of Lemma 1

Suppose by way of contradiction that the realized history in period $t-1$ was $\hat{\mathbf{h}} = (\hat{h}_1, \ldots, \hat{h}_l)$, which is neither $\mathbf{h}^*$ nor $\mathbf{h}^* - \mathbf{1}_j$. We obtain a contradiction in each of the following three cases:

First, if $\hat{h}_j = 0$, then $h_j^* - \hat{h}_j = 1$, so the gradualism property implies that $\sum_{i=1}^{j-1} \hat{h}_i \geq \sum_{i=1}^{j-1} h_i^*$. Since $\hat{\mathbf{h}} \neq \mathbf{h}^* - \mathbf{1}_j$, we have $(\hat{h}_1, \ldots, \hat{h}_{j-1}) \neq (h_1^*, \ldots, h_{j-1}^*)$. Hence, there exists $1 \leq i \leq j-1$ such that $\hat{h}_i > h_i^*$. This implies that after observing s_j at history $\hat{\mathbf{h}}$, the agent discloses s_j yet conceals at least one realization of s_i, which contradicts the monotonicity of their strategy.

Second, if $\hat{h}_j = 1$, then the construction of $\mathbf{h}^*$ implies that $\sum_{i=1}^{j-1} \hat{h}_i \geq \sum_{i=1}^{j-1} h_i^*$. Since $\hat{\mathbf{h}} \neq \mathbf{h}^*$, we have $(\hat{h}_1, \ldots, \hat{h}_{j-1}) \neq (h_1^*, \ldots, h_{j-1}^*)$, so there exists $1 \leq i \leq j-1$ such that $\hat{h}_i > h_i^*$. This again implies that from history $\hat{\mathbf{h}}$ to $\mathbf{h}^*$, at least one copy of s_j is disclosed while at least one copy of s_i is concealed, which contradicts the monotonicity of their strategy.

Third, if $\hat{h}_j \geq 2$, then consider the sequence of agents' realized histories starting from the latest communication failure, which we normalize as period 0. Let $\mathbf{h}(\tau)$ denote the realized history in period τ. By definition, $\mathbf{h}(0) = (0, 0, \ldots, 0)$, $\mathbf{h}(t-1) = \hat{\mathbf{h}}$, and $\mathbf{h}(t) = \mathbf{h}^*$. The gradualism property implies the existence of $0 < \tau < t-1$ such that (i) the jth entry of $\mathbf{h}(\tau)$ is 1, and (ii) the jth entry of $\mathbf{h}(\hat{\tau})$ is weakly greater than 1 for every $\hat{\tau} \in [\tau, t-1]$. Let $\mathbf{h}(\tau) = \mathbf{h}^{**} = (h_1^{**}, \ldots, h_{j-1}^{**}, 1, 0, \ldots, 0)$. The definition of vector $\mathbf{h}^*$ implies that $\sum_{i=1}^{j-1} h_i^{**} \geq \sum_{i=1}^{j-1} h_i^*$.

The first subcase is $\sum_{i=1}^{j-1} h_i^{**} = \sum_{i=1}^{j-1} h_i^*$, which by the construction of history $\mathbf{h}^*$, implies that $\pi(\mathbf{h}^*) \leq \pi(\mathbf{h}^{**})$. The construction of τ implies that there is no communication breakdown between period τ to t, so the conclusion that $\pi(\mathbf{h}^*) \leq \pi(\mathbf{h}^{**})$ contradicts the improvement principle, which requires that $\pi(\mathbf{h}^*) = \pi(\mathbf{h}(t)) > \pi(\mathbf{h}(t-1)) > \pi(\mathbf{h}(\tau)) = \pi(\mathbf{h}^{**})$, as $\mathbf{h}(t) \neq \mathbf{h}(t-1)$ and $\mathbf{h}(t-1) \neq \mathbf{h}(\tau)$.

The second subcase is $\sum_{i=1}^{j-1} h_i^{**} > \sum_{i=1}^{j-1} h_i^*$, i.e., the total number of signals $s_1, \ldots, s_{j-1}$ is strictly more in $\mathbf{h}(\tau)$ than in $\mathbf{h}(t)$. By construction of τ, the jth entry of $\mathbf{h}(\hat{\tau})$ is weakly greater than 1 for every $\hat{\tau} \in [\tau, t-1]$. This together with $\hat{h}_j \geq 2$ implies that there exists at least one period from τ to $t-1$ in which signal s_j is disclosed while some signals in $\{s_1, \ldots, s_{j-1}\}$ are concealed. This contradicts agents using a monotone strategy.

Remark: Lemma 1 extends when the communication success rate weakly decreases with the number of signals disclosed. We revisit the three cases considered in the proof. The first two cases derive a contradiction with the monotonicity of agents' strategy, where the communication success rate is irrelevant. For the third case, in the subcase where $\sum_{i=1}^{j-1} h_i^{**} = \sum_{i=1}^{j-1} h_i^*$, histories $\mathbf{h}(\tau) = \mathbf{h}^{**}$ and $\mathbf{h}(t) = \mathbf{h}^*$ have the same length, the definition of $\mathbf{h}^*$ implies that $\pi(\mathbf{h}^*) \leq \pi(\mathbf{h}^{**})$, yet $\mathbf{h}^*$ occurs after $\mathbf{h}^{**}$ with a history in between them, $\mathbf{h}(t-1)$, that differs from both. Therefore, each agent's expected payoff from disclosing $\mathbf{h}^{**}$ is weakly greater than his expected payoff from disclosing $\mathbf{h}^*$, which violates the extended version of improvement principle in this general setting. The subcase where $\sum_{i=1}^{j-1} h_i^{**} > \sum_{i=1}^{j-1} h_i^*$ is ruled out via a contradiction with the monotonicity of agents' strategy, where the communication success rate is irrelevant.

A.4 Proof of Proposition 2

For every $n \geq 1$ and $\theta \in \{\underline{\theta}, \overline{\theta}\}$, let $q_{\theta,n} = \sum_{i=1}^{n} f_{\theta,i}$, and write $\underline{q}_n = q_{\underline{\theta},n}$, $\overline{q}_n = q_{\overline{\theta},n}$. Fix $\eta > 0$ and choose any $\delta \in (0,1)$ and $p \in (0,1]$ such that $\eta(\delta,p) = \eta$. Let $\mathbf{h}_{\eta}^{(n)} = (h_1, \ldots, h_n, 0, \ldots, 0)$ denote the random steady-state history when agents use the constant-threshold strategy with threshold n. We now show that the distribution of $\mathbf{h}_{\eta}^{(n)}$ is independent of the particular pair (δ, p) chosen to represent η. Let K be the number of periods since last communication failure. Then $\Pr(K = k) = (1-\delta)\delta^k$ for every $k \in \mathbb{N}$. Fixing $K = k$, we have a multinomial experiment with $n+1$ possible outcomes (no signal, as well as each of the n possible signals). Because of independence across periods, conditional on $K = k$, the probability-generating function becomes $(1 - pq_{\theta,n} + p\sum_{i=1}^{n} f_{\theta,i} z_i)^k$. Thus, the probability-generating function of $\mathbf{h}_{\eta}^{(n)}$ is

$$\begin{aligned} G_{\theta,n}^{\eta}(z_1, \ldots, z_n) &= \mathbb{E}\left[\prod_{i=1}^{n} z_i^{h_i} \mid \theta\right] = \sum_{k=0}^{\infty}(1-\delta)\delta^k \left(1 - pq_{\theta,n} + p\sum_{i=1}^{n} f_{\theta,i} z_i\right)^k \\ &= \frac{1-\delta}{1-\delta\left(1 - pq_{\theta,n} + p\sum_{i=1}^{n} f_{\theta,i} z_i\right)} = \frac{1}{1 + \eta(\delta,p)\sum_{i=1}^{n} f_{\theta,i}(1-z_i)}. \end{aligned} \tag{A.4}$$

So the distribution $\mathbf{h}_{\eta}^{(n)}$ is well-defined.

To show that $I\left(\theta; \mathbf{h}_{\eta}^{(n)}\right)$ is strictly increasing in η for every fixed threshold n, fix $\eta' < \eta$. Generate random history $\mathbf{h}$ according to the distribution of $\mathbf{h}_{\eta}^{(n)}$. Intuitively, in the steady-state, reducing the effective information from η to η' is equivalent to independently deleting each disclosed signal with probability $1 - \eta'/\eta$. We now verify that this indeed reproduces the entire history distribution at η'. Conditional on $\mathbf{h}$, construct random history $\mathbf{h}'$ by retaining each signal in $\mathbf{h}$ with probability $r = \frac{\eta'}{\eta}$, independent of other signals and the state. By (A.4),

$$\mathbb{E}_{\theta}\left[\prod_{i=1}^{n} z_i^{h'_i} \mid \theta\right] = G_{\theta,n}^{\eta}\left(1 - r + rz_1, \ldots, 1 - r + rz_n\right) = \frac{1}{1 + \eta r \sum_{i=1}^{n} f_{\theta,i}(1-z_i)} = G_{\theta,n}^{\eta'}(z_1, \ldots, z_n).$$

Hence, conditional on either state, $\mathbf{h}'$ has the same distribution as $\mathbf{h}_{\eta'}^{(n)}$. Since the history at η' can be obtained from the history at η through state-independent deletion, what we have constructed is a Blackwell garbling. Let $\Pi = \Pr(\theta = \overline{\theta} \mid \mathbf{h})$ and $\Pi' = \Pr(\theta = \overline{\theta} \mid \mathbf{h}')$ and observe that $\Pi' = \mathbb{E}[\Pi \mid \mathbf{h}']$ because the deletion is state-independent conditional on $\mathbf{h}$. As γ is strictly concave, Jensen's inequality implies $\mathbb{E}[\gamma(\Pi) \mid \mathbf{h}'] \leq \gamma\left(\mathbb{E}[\Pi \mid \mathbf{h}']\right) = \gamma(\Pi')$. Taking expectations gives $I\left(\theta; \mathbf{h}_{\eta'}^{(n)}\right) \leq I\left(\theta; \mathbf{h}_{\eta}^{(n)}\right)$. We now show that this inequality is strict.

Conditional on $\mathbf{h}' = \mathbf{0}$, both $\mathbf{h} = \mathbf{0}$ and $\mathbf{h} = \mathbf{1}_1$ occur with positive probability. Moreover, the posterior beliefs following histories $\mathbf{h}$ and $\mathbf{h}'$ differ. By (4.6), their posterior odds satisfy

$$\frac{\Pr(\overline{\theta} \mid \mathbf{h}=\mathbf{1}_1)/\Pr(\underline{\theta} \mid \mathbf{h}=\mathbf{1}_1)}{\Pr(\overline{\theta} \mid \mathbf{h}=\mathbf{0})/\Pr(\underline{\theta} \mid \mathbf{h}=\mathbf{0})} = L_1 \frac{1+\eta \underline{q}_n}{1+\eta \overline{q}_n} > 1,$$

where the inequality holds since $L_1(1+\eta\underline{q}_n)-(1+\eta\overline{q}_n) = (L_1-1)+\eta\sum_{i=1}^n \underline{f}_i(L_1-L_i) > 0$. Thus, Π is not constant conditional on $\mathbf{h}'=\mathbf{0}$. Strict concavity of γ implies $I\left(\theta; \mathbf{h}_{\eta'}^{(n)}\right) < I\left(\theta; \mathbf{h}_{\eta}^{(n)}\right)$.

We now show that $I\left(\theta; \mathbf{h}_{\eta}^{(n)}\right)$ is continuous in η. Let $\varepsilon = \Pr(\mathbf{h} \neq \mathbf{h}')$. By (A.4), conditional on θ, the expected number of deleted signals is $\left(1-\frac{\eta'}{\eta}\right)\mathbb{E}_\theta[|\mathbf{h}|] = (\eta-\eta')q_{\theta,n}$. So, $\varepsilon \leq (\eta - \eta')\max_\theta q_{\theta,n} \leq \eta-\eta'$. Since $\mathbf{h}'$ is constructed from $\mathbf{h}$ using a state-independent randomization, by the chain rule for mutual information $I(\theta;\mathbf{h}) - I(\theta;\mathbf{h}') = I(\theta;\mathbf{h} \mid \mathbf{h}')$.

Let $E = \mathbb{1}_{\{\mathbf{h}\neq\mathbf{h}'\}}$. Observe that E is a deterministic function of $(\mathbf{h},\mathbf{h}')$, so that the chain rule gives $I(\theta;\mathbf{h}\mid\mathbf{h}') = I(\theta;\mathbf{h},E\mid\mathbf{h}') = I(\theta;E\mid\mathbf{h}') + I(\theta;\mathbf{h}\mid\mathbf{h}',E)$. For the first term, we have $I(\theta;E\mid\mathbf{h}') \leq \mathbb{E}\left[\gamma\left(\Pr(E=1\mid\mathbf{h}')\right)\right] \leq \gamma(\varepsilon)$, where the first inequality is because the mutual information between θ and E, conditional on $\mathbf{h}'$, is bounded by the binary entropy of E, and the second follows from Jensen's inequality. We also have that

$$I(\theta;\mathbf{h}\mid\mathbf{h}',E) = (1-\varepsilon)I(\theta;\mathbf{h}\mid\mathbf{h}',E=0) + \varepsilon I(\theta;\mathbf{h}\mid\mathbf{h}',E=1).$$

When $E=0$, we have $\mathbf{h}=\mathbf{h}'$, so the first term is 0. When $E=1$, the mutual information is at most $\log 2$ because θ is binary. Hence, $I(\theta;\mathbf{h}\mid\mathbf{h}',E) \leq \varepsilon\log 2$. Combining these two bounds, $I(\theta;\mathbf{h}\mid\mathbf{h}') \leq \gamma(\varepsilon)+\varepsilon\log 2$. Since $\varepsilon\to 0$ whenever $\eta-\eta'\to 0$, the preceding bound applies to every pair $0<\eta'<\eta$, and so, $I\left(\theta;\mathbf{h}_\eta^{(n)}\right)$ is continuous in η for fixed n.

We now show that mutual information is continuous when n changes. Fix $n\geq 2$ and consider $\eta=\eta_n>0$. Let $\pi_{\eta_n}^{(m)}(\mathbf{h})$ denote the posterior belief after history $\mathbf{h}$ under threshold m. By the definition of η_n, (4.7) gives $L_n \cdot \frac{1+\eta_n\underline{q}_n}{1+\eta_n\overline{q}_n} = 1$. Since $\overline{f}_n = L_n\underline{f}_n$, this implies $L_n\cdot\frac{1+\eta_n\underline{q}_{n-1}}{1+\eta_n\overline{q}_{n-1}} = 1$. So,

$$\frac{1+\eta_n\underline{q}_n}{1+\eta_n\overline{q}_n} = \frac{1+\eta_n\underline{q}_{n-1}}{1+\eta_n\overline{q}_{n-1}} = \frac{1}{L_n}. \tag{A.5}$$

Let $\mathbf{h}=(h_1,\ldots,h_n,0,\ldots,0)$ and let $\widehat{\mathbf{h}}=(h_1,\ldots,h_{n-1},0,\ldots,0)$. Using (4.6) and (A.5), the posterior odds under threshold n satisfy

$$\frac{\pi_{\eta_n}^{(n)}(\mathbf{h})}{1-\pi_{\eta_n}^{(n)}(\mathbf{h})} = \frac{\pi_0}{1-\pi_0}\left(\frac{1}{L_n}\right)^{|\mathbf{h}|+1}\prod_{i=1}^n L_i^{h_i} = \frac{\pi_0}{1-\pi_0}\left(\frac{1}{L_n}\right)^{\sum_{i=1}^{n-1}h_i+1}\prod_{i=1}^{n-1} L_i^{h_i} = \frac{\pi_{\eta_n}^{(n-1)}(\widehat{\mathbf{h}})}{1-\pi_{\eta_n}^{(n-1)}(\widehat{\mathbf{h}})}.$$

Therefore, $\pi_{\eta_n}^{(n)}(\mathbf{h}) = \pi_{\eta_n}^{(n-1)}(\widehat{\mathbf{h}})$. That is, at the cutoff, the number of disclosed copies of s_n does not affect the posterior belief. This implies that the information structures under thresholds n and $n-1$ are Blackwell equivalent at η_n.

Putting $z_n = 1$ in (A.4) gives $G_{\theta,n}^{\eta_n}(z_1, \ldots, z_{n-1}, 1) = \left(1 + \eta_n \sum_{i=1}^{n-1} f_{\theta,i}(1 - z_i)\right)^{-1} = G_{\theta,n-1}^{\eta_n}(z_1, \ldots, z_{n-1})$. Thus, conditional on either state, the first $n-1$ coordinates of $\mathbf{h}_{\eta_n}^{(n)}$ have the same distribution as $\mathbf{h}_{\eta_n}^{(n-1)}$. Together with $\pi^{(n)}(\mathbf{h}) = \pi^{(n-1)}(\widehat{\mathbf{h}})$ this implies $I\left(\theta; \mathbf{h}_{\eta_n}^{(n)}\right) = I\left(\theta; \mathbf{h}_{\eta_n}^{(n-1)}\right)$.

For every fixed threshold n, $I\left(\theta; \mathbf{h}_{\eta}^{(n)}\right)$ is continuous and strictly increasing in η. Moreover, the mutual information associated with thresholds n and $n-1$ coincides at every positive cutoff η_n. By Theorem 1, the largest threshold in a constant-threshold equilibrium depends on (δ, p) only through $\eta(\delta, p)$. Since there are only finitely many thresholds, there exists a continuous and strictly increasing function $\widehat{\mathcal{I}} : (0, \infty) \to \mathbb{R}_+$ such that $\mathcal{I}(\delta, p) = \widehat{\mathcal{I}}\big(\eta(\delta, p)\big)$.

A.5 Proof of Proposition 3

Fix $p \in (0, 1]$ and write $\eta = \eta(\delta, p)$. We know that $\eta \to \infty$ as $\delta \to 1$. By Corollary 2, when δ is sufficiently close to 1, the most informative constant-threshold equilibrium is the one in which agents disclose all available copies of s_1 and conceal all other signals. Hence, every on-path history takes the form $m\mathbf{1}_1$ for some $m \in \mathbb{N}$ (and note that the length of the history is $|\mathbf{h}_{\delta,p}| = m$). Setting $n = 1$ in equation (A.4), conditional on state θ, the probability-generating function of the length of $\mathbf{h}_{\delta,p}$ is $\mathbb{E}\left[z^{|\mathbf{h}_{\delta,p}|} \mid \theta\right] = (1 + \eta f_{\theta,1}(1 - z))^{-1}$. It follows that for every $m \in \mathbb{N}$, we have

$$\Pr\left(|\mathbf{h}_{\delta,p}| = m \mid \theta\right) = \frac{1}{1 + \eta f_{\theta,1}} \left(\frac{\eta f_{\theta,1}}{1 + \eta f_{\theta,1}}\right)^m, \tag{A.6}$$

Thus, conditional on any state θ, the number of disclosed copies of s_1 has a geometric distribution with mean $\eta f_{\theta,1}$. For every $x \geq 0$, equation (A.6) implies

$$\Pr\left(\frac{|\mathbf{h}_{\delta,p}|}{\eta} > x \,\middle|\, \theta\right) = \sum_{k=\lfloor \eta x \rfloor + 1}^{\infty} \frac{1}{1 + \eta f_{\theta,1}} \left(\frac{\eta f_{\theta,1}}{1 + \eta f_{\theta,1}}\right)^k = \left(\frac{\eta f_{\theta,1}}{1 + \eta f_{\theta,1}}\right)^{\lfloor \eta x \rfloor + 1}.$$

Since $\eta \log\left(\frac{\eta f_{\theta,1}}{1 + \eta f_{\theta,1}}\right) = -\eta \log\left(1 + \frac{1}{\eta f_{\theta,1}}\right) \to -\frac{1}{f_{\theta,1}}$, and $\frac{\lfloor \eta x \rfloor + 1}{\eta} \to x$, we have

$$\log \Pr\left(\frac{|\mathbf{h}_{\delta,p}|}{\eta} > x \,\middle|\, \theta\right) = \frac{\lfloor \eta x \rfloor + 1}{\eta} \eta \log\left(\frac{\eta f_{\theta,1}}{1 + \eta f_{\theta,1}}\right) \to -\frac{x}{f_{\theta,1}}.$$

Thus, the preceding calculations imply that, for every $x \geq 0$, $\Pr\left(\frac{|\mathbf{h}_{\delta,p}|}{\eta} > x \,\middle|\, \theta\right) \to \exp\left(-\frac{x}{f_{\theta,1}}\right)$, or that $\Pr\left(\frac{|\mathbf{h}_{\delta,p}|}{\eta} \leq x \,\middle|\, \theta\right) \to 1 - \exp\left(-\frac{x}{f_{\theta,1}}\right)$, which is the CDF of an exponential random variable with mean $f_{\theta,1}$. Therefore, $\frac{|\mathbf{h}_{\delta,p}|}{\eta} \Rightarrow Y_\theta$, where, conditional on state θ, Y_θ is exponentially distributed with mean $f_{\theta,1}$.

We now derive the limiting distribution of agents' posterior belief in the steady state. From equation (4.6), the posterior odds after observing history $m\mathbf{1}_1$ are

$$\frac{\pi(m\mathbf{1}_1)}{1-\pi(m\mathbf{1}_1)} = \frac{\pi_0}{1-\pi_0}\left(\frac{1+\eta\underline{f}_1}{1+\eta\overline{f}_1}\right)^{m+1} L_1^m. \tag{A.7}$$

Taking logarithms, we can write the posterior log likelihood ratio as

$$\log\frac{\pi(m\mathbf{1}_1)}{1-\pi(m\mathbf{1}_1)} = \log\frac{\pi_0}{1-\pi_0} + \log\left(\frac{1+\eta\underline{f}_1}{1+\eta\overline{f}_1}\right) + \frac{m}{\eta}\eta\log\left(L_1\frac{1+\eta\underline{f}_1}{1+\eta\overline{f}_1}\right).$$

Note that for the second term we have $\log\left(\frac{1+\eta\underline{f}_1}{1+\eta\overline{f}_1}\right) \to -\log L_1$. Using $\overline{f}_1 = L_1\underline{f}_1$, the expression inside the logarithm in the last term can be rewritten as $L_1\frac{1+\eta\underline{f}_1}{1+\eta\overline{f}_1} = \frac{L_1+\eta\overline{f}_1}{1+\eta\overline{f}_1} = 1+\frac{L_1-1}{1+\eta\overline{f}_1}$, so that

$$\eta\log\left(L_1\frac{1+\eta\underline{f}_1}{1+\eta\overline{f}_1}\right) \to \frac{L_1-1}{\overline{f}_1} = \frac{1}{\underline{f}_1} - \frac{1}{\overline{f}_1}.$$

Let $Z_\theta = \log\frac{\pi_0}{1-\pi_0} - \log L_1 + \left(\frac{1}{\underline{f}_1} - \frac{1}{\overline{f}_1}\right)Y_\theta$. The preceding limits and Slutsky's theorem imply that, conditional on state θ, $\log\frac{\pi(\mathbf{h}_{\delta,p})}{1-\pi(\mathbf{h}_{\delta,p})} \Rightarrow Z_\theta$.

Since $L_1 > 1$, $\frac{1}{\underline{f}_1} - \frac{1}{\overline{f}_1} > 0$. Conditional on any state θ, Y_θ is finite with probability one and has a nondegenerate exponential distribution. Therefore, Z_θ is also finite with probability one and nondegenerate. Because the function $z \mapsto \frac{\exp(z)}{1+\exp(z)}$ is continuous and strictly increasing from $\mathbb{R}$ into $(0,1)$, the continuous mapping theorem implies that, conditional on either state, $\pi(\mathbf{h}_{\delta,p}) \Rightarrow \frac{\exp(Z_\theta)}{1+\exp(Z_\theta)}$, which is a nondegenerate random variable taking values in $(0,1)$. Full learning would require $\pi(\mathbf{h}_{\delta,p})$ to converge in probability to 1 conditional on $\theta = \overline{\theta}$ and to 0 conditional on $\theta = \underline{\theta}$. Convergence in probability to either constant would imply convergence in distribution to that same constant. This is impossible because the limiting posterior distribution is nondegenerate and takes values in $(0,1)$.

A.6 Proof of Proposition 4

Suppose in equilibrium, agents use a constant-threshold strategy with threshold $n \geq 2$, and let $\mu_{\theta,\varepsilon}^{(n)}$ denote the steady-state distribution of observed histories conditional on state θ. Let $q_{\theta,n} = \sum_{i=1}^n f_{\theta,i}$. Let K denote the number of consecutive periods since the most recent total communication failure. Then $\Pr(K=k) = a\varepsilon(1-a\varepsilon)^k$ for $k \in \mathbb{N}$. Therefore, conditional on θ, the probability-generating function of the observed history is

$$G_{\theta,n}^{\varepsilon}(z_1,\ldots,z_n) = a\varepsilon\sum_{k=0}^{\infty}(1-a\varepsilon)^k\prod_{t=1}^{k}\left[1-p\rho_\varepsilon^t\sum_{i=1}^n f_{\theta,i}(1-z_i)\right]. \tag{A.8}$$

From (A.8) we have

$$\frac{\mu_{\theta,\varepsilon}^{(n)}(\mathbf{0})}{a\varepsilon} = \sum_{k=0}^{\infty}(1-a\varepsilon)^k \prod_{t=1}^{k}(1-pq_{\theta,n}\rho_\varepsilon^t), \tag{A.9}$$

$$\frac{\mu_{\theta,\varepsilon}^{(n)}(\mathbf{1}_n)}{a\varepsilon} = \sum_{k=1}^{\infty}(1-a\varepsilon)^k \sum_{t=1}^{k} pf_{\theta,n}\rho_\varepsilon^t \prod_{\substack{s=1\\ s\neq t}}^{k}(1-pq_{\theta,n}\rho_\varepsilon^s). \tag{A.10}$$

For every fixed k, $(1-a\varepsilon)^k \to 1$ and $\rho_\varepsilon^t \to 1$ for every $t \leq k$. Hence, the kth summand in (A.9) converges to $(1-pq_{\theta,n})^k$. Similarly, each of the k terms in the inner sum in (A.10) converges to $pf_{\theta,n}(1-pq_{\theta,n})^{k-1}$, so the kth summand converges to $kpf_{\theta,n}(1-pq_{\theta,n})^{k-1}$.

To show that $\lim_{\varepsilon\to 0}\frac{\mu_{\theta,\varepsilon}^{(n)}(\mathbf{0})}{a\varepsilon} = \sum_{k=0}^{\infty}\lim_{\varepsilon\to 0}\Big((1-a\varepsilon)^k\prod_{t=1}^{k}(1-pq_{\theta,n}\rho_\varepsilon^t)\Big) = \sum_{k=0}^{\infty}(1-pq_{\theta,n})^k = 1/(pq_{\theta,n})$, we need to justify taking the limit inside the sum. To this end, let $M_\varepsilon = \lfloor \varepsilon^{-1/2}\rfloor$. For all sufficiently small ε, if $t \leq M_\varepsilon$, Bernoulli's inequality implies $\rho_\varepsilon^t = (1-b\varepsilon)^t \geq 1-b\varepsilon t \geq 1-b\sqrt{\varepsilon} \geq \frac{1}{2}$. Thus, for $k \leq M_\varepsilon$, the summands in (A.9) are bounded by $(1-pq_{\theta,n}/2)^k$.

If $k \geq M_\varepsilon$, the product contains at least the first M_ε factors. As such, the summand is bounded above by $(1-a\varepsilon)^k\left(1-\frac{pq_{\theta,n}}{2}\right)^{M_\varepsilon}$. Therefore,

$$\sum_{k=M_\varepsilon}^{\infty}(1-a\varepsilon)^k\prod_{t=1}^{k}\left(1-pq_{\theta,n}\rho_\varepsilon^t\right) \leq \frac{1}{a\varepsilon}\left(1-\frac{pq_{\theta,n}}{2}\right)^{M_\varepsilon},$$

which converges to zero as $\varepsilon \to 0$. For any M and small enough ε such that $M_\varepsilon > M$, we have

$$\left|\frac{\mu_{\theta,\varepsilon}^{(n)}(\mathbf{0})}{a\varepsilon} - \sum_{k=0}^{\infty}(1-pq_{\theta,n})^k\right| \leq \sum_{k=0}^{M}\left|(1-a\varepsilon)^k\prod_{t=1}^{k}\left(1-pq_{\theta,n}\rho_\varepsilon^t\right) - (1-pq_{\theta,n})^k\right| + \sum_{k=M+1}^{M_\varepsilon-1}\left(1-\frac{pq_{\theta,n}}{2}\right)^k$$
$$+\sum_{k=M_\varepsilon}^{\infty}(1-a\varepsilon)^k\prod_{t=1}^{k}\left(1-pq_{\theta,n}\rho_\varepsilon^t\right) + \sum_{k=M+1}^{\infty}(1-pq_{\theta,n})^k.$$

Fixing M, the first term converges to zero as $\varepsilon \to 0$, because it is a finite sum and each summand converges. The third term also converges to zero by the bound established above. Therefore,

$$\limsup_{\varepsilon\to 0}\left|\frac{\mu_{\theta,\varepsilon}^{(n)}(\mathbf{0})}{a\varepsilon} - \sum_{k=0}^{\infty}(1-pq_{\theta,n})^k\right| \leq \sum_{k=M+1}^{\infty}\left(1-\frac{pq_{\theta,n}}{2}\right)^k + \sum_{k=M+1}^{\infty}(1-pq_{\theta,n})^k.$$

Since both sums on the right converge to zero as $M \to \infty$, we have $\lim_{\varepsilon\to 0}\frac{\mu_{\theta,\varepsilon}^{(n)}(\mathbf{0})}{a\varepsilon} = \sum_{k=0}^{\infty}(1-pq_{\theta,n})^k = \frac{1}{pq_{\theta,n}}$. The argument for taking the limit inside equation (A.10) is analogous and establishes that $\lim_{\varepsilon\to 0}\frac{\mu_{\theta,\varepsilon}^{(n)}(\mathbf{1}_n)}{a\varepsilon} = \frac{f_{\theta,n}}{pq_{\theta,n}^2}$.

Write $\underline{q}_n = \sum_{i=1}^{n}\underline{f}_i$ and $\overline{q}_n = \sum_{i=1}^{n}\overline{f}_i$. By what we just showed, the ratio of posterior odds after histories $\mathbf{1}_n$ and $\mathbf{0}$ converges, as $\varepsilon \to 0$, to

$$\frac{\mu_{\overline{\theta},\varepsilon}^{(n)}(\mathbf{1}_n)/\mu_{\underline{\theta},\varepsilon}^{(n)}(\mathbf{1}_n)}{\mu_{\overline{\theta},\varepsilon}^{(n)}(\mathbf{0})/\mu_{\underline{\theta},\varepsilon}^{(n)}(\mathbf{0})} \to \frac{\overline{f}_n/(p\overline{q}_n^2)}{\underline{f}_n/(p\underline{q}_n^2)}\frac{1/(p\underline{q}_n)}{1/(p\overline{q}_n)} = L_n\frac{\underline{q}_n}{\overline{q}_n}.$$

For every $n \geq 2$, $\overline{q}_n - L_n \underline{q}_n = \sum_{i=1}^{n} \underline{f}_i (L_i - L_n) = \sum_{i=1}^{n-1} \underline{f}_i (L_i - L_n) > 0$, since $L_i > L_n$ for every $i < n$. Therefore, $L_n \frac{\underline{q}_n}{\overline{q}_n} < 1$ and for every $n \geq 2$ and all sufficiently small ε, posterior odds are strictly lower after $\mathbf{1}_n$ than after $\mathbf{0}$. Thus $\pi_\varepsilon^{(n)}(\mathbf{1}_n) < \pi_\varepsilon^{(n)}(\mathbf{0})$.

Consider an agent with inherited history $\mathbf{0}$ and private signal s_n. Under the threshold-n strategy, he discloses s_n. If he instead conceals it, his successor observes $\mathbf{0}$, whereas disclosure leads to $\mathbf{1}_n$ with probability $\delta_\varepsilon \rho_\varepsilon > 0$ and to $\mathbf{0}$ otherwise. Since $\pi_\varepsilon^{(n)}(\mathbf{1}_n) < \pi_\varepsilon^{(n)}(\mathbf{0})$, and this ordering is preserved after every realization of the successor's private signal, concealment induces a strictly higher expected successor action. Thus, the agent strictly prefers to conceal s_n, hence this cannot be an equilibrium. Since there are finitely many $n \geq 2$, we can choose a common $\varepsilon^* > 0$ for which this argument applies to every such n.

We next construct a constant-threshold equilibrium where only s_1 is disclosed. For each state θ, let $\{Y_{\theta,t}\}_{t \geq 1}$ be independent Bernoulli random variables with $\Pr(Y_{\theta,t} = 1) = p f_{\theta,1} \rho_\varepsilon^t$. If $a > b$ so that $\delta_\varepsilon < \rho_\varepsilon$, let N be independent of these random variables and satisfy $\Pr(N = r) = \left(1 - \frac{\delta_\varepsilon}{\rho_\varepsilon}\right)\left(\frac{\delta_\varepsilon}{\rho_\varepsilon}\right)^{r-1}$ for $r \geq 1$. If $a = b$, so that $\delta_\varepsilon = \rho_\varepsilon$, set $N = \infty$.

Let $S_{\theta,k} = \sum_{t=1}^{k} Y_{\theta,t}$, $S_{\theta,0} = 0$, and let $M_\theta = S_{\theta,N}$. When $N = \infty$, this sum is finite almost surely because $\sum_{t=1}^{\infty} p f_{\theta,1} \rho_\varepsilon^t < \infty$. Under the constant-threshold strategy of disclosing only s_1, conditional on the most recent total communication failure having occurred k periods ago, the number of surviving copies of s_1 has the same distribution as $S_{\theta,k}$. Hence,

$$\mu_{\theta,\varepsilon}^{(1)}(m\mathbf{1}_1) = (1 - \delta_\varepsilon) \sum_{k=0}^{\infty} \delta_\varepsilon^k \Pr(S_{\theta,k} = m). \tag{A.11}$$

Moreover, since $S_{\theta,k+1} = S_{\theta,k} + Y_{\theta,k+1}$, we have $\Pr(S_{\theta,k+1} > m) - \Pr(S_{\theta,k} > m) = p f_{\theta,1} \rho_\varepsilon^{k+1} \Pr(S_{\theta,k} = m)$. The event $M_\theta > m$ can be decomposed according to the stage at which the number of successes first becomes greater than m. Using $\Pr(N \geq k+1) = \left(\frac{\delta_\varepsilon}{\rho_\varepsilon}\right)^k$, we get

$$\begin{aligned} \Pr(M_\theta > m) &= \sum_{k=0}^{\infty} \Pr(N \geq k+1) \left[\Pr(S_{\theta,k+1} > m) - \Pr(S_{\theta,k} > m)\right] \\ &= p f_{\theta,1} \rho_\varepsilon \sum_{k=0}^{\infty} \delta_\varepsilon^k \Pr(S_{\theta,k} = m). \end{aligned}$$

Combining this with (A.11) gives

$$\mu_{\theta,\varepsilon}^{(1)}(m\mathbf{1}_1) = \frac{1 - \delta_\varepsilon}{p f_{\theta,1} \rho_\varepsilon} \Pr(M_\theta > m). \tag{A.12}$$

We next show that the sequence $\{\Pr(M_{\overline{\theta}} \geq m)\}_{m \geq 0}$ is log-concave. Let $c = \frac{\delta_\varepsilon}{\rho_\varepsilon} \leq 1$. First truncate the construction after an arbitrary finite stage R. Let $T_{t,m}^{(R)}$ denote the probability of obtaining at least m successes from stages $t, \ldots, R$ when the process is observed immediately before the continuation decision at stage t. We show by backward induction on t that $\left(T_{t,m}^{(R)}\right)^2 \geq T_{t,m-1}^{(R)} T_{t,m+1}^{(R)}$ for every $m \geq 1$.

The terminal sequence satisfies $T_{R+1,0}^{(R)} = 1$ and $T_{R+1,m}^{(R)} = 0$ for $m \geq 1$, so it is log-concave. Suppose $U_m = T_{t+1,m}^{(R)}$ is log-concave, and write $q = p\overline{f}_1 \rho_\varepsilon^t$. For every $m \geq 1$, $T_{t,m}^{(R)} = c\left[qU_{m-1} + (1-q)U_m\right]$. For $m \geq 2$, we have

$$\begin{aligned}
&[qU_{m-1} + (1-q)U_m]^2 - [qU_{m-2} + (1-q)U_{m-1}]\,[qU_m + (1-q)U_{m+1}] \\
&= q^2\left(U_{m-1}^2 - U_{m-2}U_m\right) + (1-q)^2\left(U_m^2 - U_{m-1}U_{m+1}\right) + q(1-q)\left(U_{m-1}U_m - U_{m-2}U_{m+1}\right) \geq 0,
\end{aligned}$$

since the first two terms are nonnegative by log-concavity, while the third is nonnegative because the successive ratios of a nonnegative log-concave sequence are weakly decreasing.

We now consider $m = 1$. Let $x = U_1$. Since every subsequent success probability is weakly below q, the probability of at least one subsequent success is no greater than it would be in an infinite process with success probability q at every stage. The latter probability solves $u = c\left[q + (1-q)u\right]$, and therefore $x \leq \frac{cq}{1-c+cq}$. Log-concavity of U and $U_0 = 1$ imply $U_2 \leq x^2$. So, $\left(T_{t,1}^{(R)}\right)^2 - T_{t,2}^{(R)} \geq c\left\{c[q + (1-q)x]^2 - qx - (1-q)x^2\right\} = c(qx - q - x)\left[x(1-c+cq) - cq\right] \geq 0$, because both factors are nonpositive. This completes the backward induction.

Letting $R \to \infty$ preserves these inequalities. When $c < 1$, the distribution of $M_{\overline{\theta}}$ is obtained by conditioning the preceding process on continuation at the first stage. This conditioning also preserves log-concavity. For $m \geq 2$, it divides each survival probability in the relevant inequality by the same factor c. For $m = 1$, the desired conditioned inequality follows from the unconditioned inequality and $c \leq 1$. When $c = 1$, no conditioning is required. Thus, the sequence $\Pr(M_{\overline{\theta}} \geq m)$ is log-concave.

Now independently retain each success counted by $M_{\overline{\theta}}$ with probability $r = \frac{\underline{f}_1}{\overline{f}_1} \in (0,1)$. At stage t, a retained success occurs with probability $p\overline{f}_1\rho_\varepsilon^t \frac{\underline{f}_1}{\overline{f}_1} = p\underline{f}_1\rho_\varepsilon^t$. The retained count therefore has the same distribution as $M_{\underline{\theta}}$. For $k \geq 1$, let $D_k = M_{\overline{\theta}} - k \big| M_{\overline{\theta}} \geq k$. Log-concavity of the survival probabilities implies $D_{k+1} \leq_{\text{st}} D_k$.

Indeed, for every $s \geq 0$, $\Pr(D_k \geq s) = \frac{\Pr(M_{\overline{\theta}} \geq k+s)}{\Pr(M_{\overline{\theta}} \geq k)}$, and this expression is weakly decreasing in k when the survival probabilities are log-concave. For $s \geq 0$, let $g_k(s) = \Pr\left(\text{Binomial}(k+s, r) \geq k\right)$. The function g_k is increasing in s, while $g_{k+1}(s) < g_k(s)$ for every s. To see the latter inequality, let $Z \sim \text{Binomial}(k+s, r)$ and let $C \sim \text{Bernoulli}(r)$ be independent. Then $g_k(s) - g_{k+1}(s) = \Pr(Z \geq k) - \Pr(Z + C \geq k+1) = (1-r)\Pr(Z = k) > 0$.

Moreover, $\frac{\Pr(M_{\underline{\theta}} \geq k)}{\Pr(M_{\overline{\theta}} \geq k)} = \mathbb{E}[g_k(D_k)]$. It follows that

$$\frac{\Pr(M_{\underline{\theta}} \geq k+1)}{\Pr(M_{\overline{\theta}} \geq k+1)} = \mathbb{E}[g_{k+1}(D_{k+1})] < \mathbb{E}[g_k(D_{k+1})] \leq \mathbb{E}[g_k(D_k)] = \frac{\Pr(M_{\underline{\theta}} \geq k)}{\Pr(M_{\overline{\theta}} \geq k)}.$$

Therefore, $\frac{\Pr(M_{\overline{\theta}} > m)}{\Pr(M_{\underline{\theta}} > m)}$ is strictly increasing in m. By (A.12), we have

$$\frac{\mu^{(1)}_{\overline{\theta},\varepsilon}(m\mathbf{1}_1)}{\mu^{(1)}_{\underline{\theta},\varepsilon}(m\mathbf{1}_1)} = \frac{\underline{f}_1}{\overline{f}_1} \cdot \frac{\Pr(M_{\overline{\theta}} > m)}{\Pr(M_{\underline{\theta}} > m)},$$

so the posterior belief is strictly increasing in the number of surviving copies of s_1.

Off the equilibrium path, let us assign belief zero to every history containing a signal other than s_1. Since the posterior is strictly increasing in the number of surviving copies of s_1, disclosing all available copies of s_1 weakly dominates any disclosure containing fewer copies. Any disclosure containing another signal is also weakly worse, since with positive probability that signal survives and induces belief zero. Whenever the alternative disclosure differs from the prescribed one, one of these comparisons is strict with positive probability. Hence, disclosing all available copies of s_1 and no other signal is strictly optimal at every information set. Thus, there exists a strict constant-threshold equilibrium with threshold 1.

A.7 Proof of Proposition 5

If the sender uses Strategy k, then the receiver will only observe $\mathbf{h} \equiv (k+1)\mathbf{1}_1 + \mathbf{1}_2$ and $j\mathbf{1}_1$ for every $j \neq k+1$ with positive probability. The receiver's posterior beliefs $\pi : H \to [0,1]$ at these histories are pinned down by Bayes rule. Strategy k is part of a strict equilibrium if and only if (i) $\pi(j\mathbf{1}_1)$ is strictly increasing in j for every $j \neq k+1$, (ii) $\pi((k+2)\mathbf{1}_1) > \pi(\mathbf{h})$, and (iii) $\pi(\mathbf{h}) > \pi(k\mathbf{1}_1)$. The first two conditions are satisfied for all large enough $k \in \mathbb{N}$.

Next, we show that for any $\delta \in (0,1)$, there exists a large enough $k \in \mathbb{N}$ such that the third condition is also satisfied. If the sender uses Strategy k in the one-shot game, then conditional on state θ, the receiver will observe $k\mathbf{1}_1$ with probability

$$(1-\delta)\delta^k \Big(1 - \delta(1 - f_{\theta,1})\Big)^{-(k+1)} f_{\theta,1}^k + (1-\delta)\delta^{k+1} \Big(1 - \delta(1 - f_{\theta,1} - f_{\theta,2})\Big)^{-(k+2)} f_{\theta,1}^{k+1},$$

where the first term is the probability that the sender observes k copies of s_1 and an unknown number of copies of s_2 to s_l, and the second term is the probability that the sender observes $k+1$ copies of s_1, no signal s_2, as well as an unknown number of copies of s_3 to s_l. Still under

Strategy k and conditional on state θ, the receiver will observe $\mathbf{h}$ with probability

$$(1-\delta)\delta^{k+1}\Big(1-\delta(1-f_{\theta,1})\Big)^{-(k+2)} f_{\theta,1}^{k+1} - (1-\delta)\delta^{k+1}\Big(1-\delta(1-f_{\theta,1}-f_{\theta,2})\Big)^{-(k+2)} f_{\theta,1}^{k+1},$$

where the first term is the probability that the sender observes $k+1$ copies of s_1 and an unknown number of signals from s_2 to s_l and the second term coincides with the second term in the previous expression except with the opposite sign. According to Bayes rule, we know that $\pi(\mathbf{h}) > \pi(k\mathbf{1}_1)$ if and only if the value of

$$\frac{\Big(1-\delta(1-f_{\theta,1})\Big)^{-(k+1)} + \delta f_{\theta,1}\Big(1-\delta(1-f_{\theta,1}-f_{\theta,2})\Big)^{-(k+2)}}{f_{\theta,1}\Big(1-\delta(1-f_{\theta,1})\Big)^{-(k+2)} - f_{\theta,1}\Big(1-\delta(1-f_{\theta,1}-f_{\theta,2})\Big)^{-(k+2)}} \tag{A.13}$$

strictly decreases once θ increases from $\underline{\theta}$ to $\overline{\theta}$. This is the case if and only if

$$\frac{\Big(\overline{X}^{-(k+1)} + \delta\overline{f}_1 \cdot \overline{Y}^{-(k+2)}\Big) \cdot \Big(\underline{X}^{-(k+2)} - \underline{Y}^{-(k+2)}\Big)}{\Big(\underline{X}^{-(k+1)} + \delta\underline{f}_1 \cdot \underline{Y}^{-(k+2)}\Big) \cdot \Big(\overline{X}^{-(k+2)} - \overline{Y}^{-(k+2)}\Big)} < \frac{\overline{f}_1}{\underline{f}_1} = L_1, \tag{A.14}$$

where $\overline{X} = 1-\delta(1-\overline{f}_1)$, $\underline{X} = 1-\delta(1-\underline{f}_1)$, $\overline{Y} = 1-\delta(1-\overline{f}_1-\overline{f}_2)$, and $\underline{Y} = 1-\delta(1-\underline{f}_1-\underline{f}_2)$. Let $\overline{R} = \overline{X}/\overline{Y}$ and $\underline{R} = \underline{X}/\underline{Y}$. Inequality (A.14) can then be rewritten as

$$\frac{(\overline{X} + \delta\overline{f}_1 \cdot \overline{R}^{k+2})(1-\underline{R}^{k+2})}{(\underline{X} + \delta\underline{f}_1 \cdot \underline{R}^{k+2})(1-\overline{R}^{k+2})} < L_1 \tag{A.15}$$

For every $\delta \in (0,1)$, we have $\overline{X} < \overline{Y}$ and $\underline{X} < \underline{Y}$, so $\overline{R}, \underline{R} \in (0,1)$. Fix any $\delta \in (0,1)$, we know that both $\overline{R}^{k+2}$ and $\underline{R}^{k+2}$ converge to 0 as $k \to \infty$. Hence, the numerator of the LHS of (A.15) converges to $\overline{X}$ and its denominator converges to $\underline{X}$, so the LHS of (A.15) converges to $\overline{X}/\underline{X}$ as $k \to \infty$. Moreover, since $\overline{f}_1 > \underline{f}_1$, we have

$$L_1 - \frac{\overline{X}}{\underline{X}} = \frac{\overline{f}_1\underline{X} - \underline{f}_1\overline{X}}{\underline{f}_1\underline{X}} = \frac{(1-\delta)\Big(\overline{f}_1 - \underline{f}_1\Big)}{\underline{f}_1\underline{X}} > 0 \quad \text{for every } \delta \in (0,1).$$

Therefore, for every $\delta \in (0,1)$, the LHS of (A.15), converges to $\overline{X}/\underline{X}$ as $k \to \infty$, which is strictly less than L_1. It follows that there exists a large enough $K \in \mathbb{N}$ such that inequality (A.15) is satisfied for every $k \geq K$. Hence, for any $\delta \in (0,1)$, there exists a large enough $k \in \mathbb{N}$ such that Strategy k is part of a strict monotone equilibrium.

References

Acemoglu, D. and M. O. Jackson (2015): "History, Expectations, and Leadership in the Evolution of Social Norms," *Review of Economic Studies*, 82, 423–456.

Anderlini, L., D. Gerardi, and R. Lagunoff (2012): "Communication and Learning," *Review of Economic Studies*, 79, 419–450.

Antić, N. and A. Chakraborty (2025): "Selected Facts," *Working Paper*.

Arieli, I. and C. Stewart (2025): "Bayesian Persuasion without Commitment," *Working Paper*.

Banerjee, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.

Banerjee, A. and D. Fudenberg (2004): "Word-of-Mouth Learning," *Games and Economic Behavior*, 46, 1–22.

Bardhi, A. and N. Bobkova (2023): "Local Evidence and Diversity in Minipublics," *Journal of Political Economy*, 131, 2451–2508.

Bénabou, R. and N. Vellodi (2025): "(Pro-)Social Learning and Strategic Disclosure," *American Economic Journal: Microeconomics*, 17, 102–125.

Bhaskar, V. (1998): "Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems," *Review of Economic Studies*, 65, 135–149.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

Bull, J. and J. Watson (2019): "Statistical evidence and the problem of robust litigation," *RAND Journal of Economics*, 50, 974–1003.

Carrillo, J. D. and T. Mariotti (2000): "Strategic Ignorance as a Self-Disciplining Device," *Review of Economic Studies*, 67, 529–544.

Cho, I.-K. and D. Kreps (1987): “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 102, 179–221.

Clark, D., D. Fudenberg, and A. Wolitzky (2021): “Record-Keeping and Cooperation in Large Societies,” *Review of Economic Studies*, 88, 2179–2209.

Dai, Y., D. Fudenberg, and H. Pei (2026): “Bayesian Persuasion with Selective Disclosure,” *Working Paper*.

Di Tillio, A., M. Ottaviani, and P. N. Sørensen (2021): “Strategic Sample Selection,” *Econometrica*, 89, 911–953.

Dye, R. A. (1985): “Disclosure of Nonproprietary Information,” *Journal of Accounting Research*, 123–145.

Dziuda, W. (2011): “Strategic Argumentation,” *Journal of Economic Theory*, 146, 1362–1397.

Felgenhauer, M. and P. Loerke (2017): “Bayesian Persuasion with Private Experimentation,” *International Economic Review*, 58, 829–856.

Felgenhauer, M. and E. Schulte (2014): “Strategic Private Experimentation,” *American Economic Journal: Microeconomics*, 6, 74–105.

Fudenberg, D. and J. Tirole (1991): “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory*, 53, 236–260.

Gao, Y. (2025): “Inference from Selectively Disclosed Data,” *Working Paper*.

Gieczewski, G. (2022): “Verifiable Communication on Networks,” *Journal of Economic Theory*, 204, 105494.

Gieczewski, G. and M. Titova (2024): “Coalition-Proof Disclosure,” *Working Paper*.

Grossman, S. (1981): “The Informational Role of Warranties and Private Disclosure about Product Quality,” *Journal of Law and Economics*, 24, 461–483.

Guttman, I., I. Kremer, and A. Skrzypacz (2014): “Not Only What But Also When: A Theory of Dynamic Voluntary Disclosure,” *American Economic Review*, 104, 2400–2420.

Hart, S., I. Kremer, and M. Perry (2017): "Evidence Games: Truth and Commitment," *American Economic Review*, 107, 690–713.

Kipping, M. and T. Clark (2012): *The Oxford Handbook of Management Consulting*, Oxford University Press.

Kolb, A. and B. Zhou (2026): "Dynamic Disclosure with(out) Timestamps," *Working Paper*.

Kremer, I., A. Schreiber, and A. Skrzypacz (2024): "Disclosing a Random Walk," *Journal of Finance*, 79, 1123–1146.

Lou, Y. (2023): "Private Experimentation, Data Truncation, and Verifiable Disclosure," *Working Paper*.

Milgrom, P. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 380–391.

Niehaus, P. (2011): "Filtered Social Learning," *Journal of Political Economy*, 119, 686–720.

Onuchic, P. and J. Ramos (2025): "Disclosure by Groups," *Working Paper*.

Pei, H. (2026): "Community Enforcement with Endogenous Records," *Review of Economic Studies*, 93, 1241–1260.

Rappoport, D. (2025): "Evidence and Skepticism in Verifiable Disclosure Games," *Theoretical Economics*, 20, 1213–1246.

Samuelson, P. (1958): "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money," *Journal of Political Economy*, 66, 467–482.

Shin, H. S. (1994): "News Management and the Value of Firms," *RAND Journal of Economics*, 58–71.

——— (2003): "Disclosures and Asset Returns," *Econometrica*, 71, 105–133.

Smith, L. and P. Sørensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.

Squintani, F. (2026): "Strategic Disclosure in Networks," *Theoretical Economics*, 21, 467–499.