

Persuasion via Sequentially Acquired Evidence

SHAOFEI JIANG*

December 17, 2024

I study a sender who can privately acquire and partially disclose hard evidence to persuade a receiver about a binary state of the world. The sender privately and sequentially acquires noisy binary signals. Signals are time-stamped and costly to acquire. When she stops, the sender discloses a left truncation of the signals. That is, it is possible to omit some most dated signals. The receiver, uncertain of how many signals the sender acquires, takes a binary action based on the difference between the number of good and bad signals in the disclosure. If the cost is not too high, there is a continuum of persuasion equilibria, and the receiver's posterior belief is supported on two points. This is akin to Bayesian persuasion. If full disclosure of signals is mandatory, the game is equivalent to costly Bayesian persuasion. Mandating full disclosure benefits the sender and hurts the receiver.

KEYWORDS: Information acquisition, Bayesian persuasion, Hard evidence

JEL CODES: C72, D82, D83

*University of Bonn (email: sjiang@uni-bonn.de). I thank V. Bhaskar, William Fuchs, John Hatfield, Sven Rady, Michael Sockin, Maxwell Stinchcombe, Yiman Sun, Caroline Thomas, Thomas Wiseman, and seminar participants at UT Austin. I gratefully acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - GZ 2047/1, Projekt-ID 390685813. Errors are my own.

1. Introduction

Evidence is often acquired over time. Consider a hedge fund manager (the sender) who wants to persuade an investor (the receiver) to invest in her hedge fund by presenting the latter with the fund's recent trading records. Trading records are accumulated over time in an incubation period using the manager's personal capital. The manager can decide *when* to seek investment and *what* to present to the investor. The investor does not know the inception date of the hedge fund, hence the manager can lie about it by truncating away older, possibly unfavorable returns and disclosing a subset of the most recent returns. The investor, consequently, will not take the disclosed returns at face value.

In this example, the sender privately learns about the profitability of the hedge fund and is not bound to share all her findings (i.e., the fund's returns) with the receiver. This is in contrast to Bayesian persuasion (Kamenica and Gentzkow, 2011) where learning is assumed to be public. However, the sender cannot falsify her findings (i.e., make false statements about past returns), nor can she cherry-pick her findings without restraint (i.e., disclose an arbitrary subset of the returns). It is only feasible to disclose a truncated version of her findings with possible omission at the beginning.

We study in this paper a persuasion problem where the sender can voluntarily and partially disclose findings. The sender sequentially acquires binary signals about an unknown binary state of the world à la Wald (1945). Learning is covert, that is, neither how much the sender learns nor what she learns is observed by the receiver, and there is a cost of acquiring each signal. When she stops learning, the sender discloses to the receiver a left truncation of her signals. The receiver takes a binary action (e.g., invest or not) to match the state of the world, but the sender wants to persuade the receiver into taking a particular action regardless of the state of the world.

If full disclosure is mandatory, the game is equivalent to costly Bayesian persuasion studied by Gentzkow and Kamenica (2014). There is a unique equilibrium. The receiver is persuaded if the difference between the number of good and bad signals in the disclosure is sufficiently high, and the sender stops acquiring additional signals if she is sufficiently optimistic or pessimistic. This paper provides a natural and tractable way to relax the assumption in Bayesian persuasion that the experiment's outcome is public.

When disclosure is voluntary but only left truncation of signals is permitted, we focus on a class of equilibria where, like in the mandatory disclosure case, the receiver is persuaded if there are sufficiently more good signals than bad signals in the disclosure. The sender stops acquiring additional signals as soon as she can persuade the receiver, or when her private belief becomes sufficiently pessimistic. The latter scenario is characterized by a single

threshold belief of the sender. Hence, whenever the sender fails to persuade the receiver, the receiver shares the same posterior belief about the state of the world with the sender. If the sender’s disclosure contains sufficiently more good signals than bad signals, the receiver is uncertain about the sender’s private belief and tries to infer what signals are omitted. However, every equilibrium is payoff equivalent to an equilibrium where the receiver cannot infer the omitted signals from the disclosure, hence his posterior belief (conditional on the event that the sender’s disclosure persuades the receiver) is independent of the disclosure. In other words, the receiver’s posterior belief is supported on two points, depending on whether he is persuaded. This is akin to the equilibrium of Bayesian persuasion.

If the cost of acquiring each signal is not too high, there is a continuum of equilibria where the receiver is persuaded with positive probability. Since the Bayesian persuasion solution is ex ante optimal for the sender, the sender’s payoff is lower in every equilibrium compared to under mandatory full disclosure. The receiver, on the other hand, is better off. His payoff is higher in all but one equilibrium than under mandatory full disclosure.

1.1. Related literature. Persuasion through public experimentation has been studied extensively. Kamenica and Gentzkow (2011) study Bayesian persuasion, where a sender commits to a costless Blackwell experiment to persuade a receiver about an unknown state of the world. Gentzkow and Kamenica (2014) study an extension where experiments have posterior separable costs (Caplin et al., 2022). Wald (1945) studies sequential sampling. A sender sequentially acquires public, noisy signals about an unknown binary state of the world before a receiver takes an action. There is a fixed cost for each signal. Morris and Strack (2019) and Jiang (2024a) show the equivalence between a sequential sampling problem and a costly Bayesian persuasion game where experiments have log-likelihood ratio costs and the state is binary. The players’ posterior belief has the same two-point distribution in the unique equilibrium of either game.

This paper studies private experimentation and assumes that the sender cannot commit to how much she learns or to fully reveal the outcome. It contributes to a growing literature relaxing the commitment assumption in Bayesian persuasion. Min (2021) and Lipnowski et al. (2022) assume that the sender privately observes the state after choosing an experiment and can substitute a cheap talk message for the outcome of the experiment with some probability. Shishkin (2019) that an experiment produces evidence only with some exogenous probability less than one, and the sender can suppress evidence. Nguyen and Tan (2021) assumes that the sender privately observes the outcome of the experiment and can change it at some cost. Guo and Shmaya (2021) assumes that the sender can ex ante assert the meaning of each outcome subject to a miscalibration cost. Lin and Liu (2024) assumes

that the sender chooses ex ante a contingency plan to replace the outcome such that the distribution of the outcome is unchanged. In this paper, we restrict the ways the sender can alter the outcome—it is only feasible to disclose a left truncation of the acquired signals. This restriction naturally arises from the sequential sampling problem, instead of from assuming a cost of lying or an exogenous probability of lying.

Our model most closely resembles Felgenhauer and Schulte (2014), but they allow the sender to disclose arbitrary subsets of her signals, hence in equilibrium, only good signals are disclosed. As we explain in Section 2.2, our assumption that only left truncations are allowed is natural and without loss of generality if signals are time-stamped. It also explains the real-life observation that bad evidence is sometimes disclosed. Henry (2009) studies a sender who commits ex ante to the number of signals and can disclose any subset of signals.

This paper is also related to hard evidence games (see, for example, Grossman (1981); Milgrom (1981); Dye (1985); Bull and Watson (2007); Hart et al. (2017)) and strategic sample selection (Di Tillio et al., 2021), where a sender privately observes some evidence, which is a random draw from a state-dependent distribution. The verifiability of evidence is modeled by assuming that the set of feasible disclosures depends on the sender’s evidence. The current paper endogenizes the distribution of evidence, and the assumption that only left truncation is permitted is natural in studying selection of time series data.

2. The Model

There is a sender and a receiver. Time $t \in \mathbb{Z}$ is discrete and unbounded in either direction. At the outset of the game, Nature chooses a binary state of the world $\omega \in \{G, B\}$ and a random date $t_0 \in \mathbb{Z}$. The sender and the receiver have a common improper prior that t_0 is uniformly distributed on \mathbb{Z} and is independent of ω . Let π_0 be the probability of the good state (i.e., $\omega = G$). The sender privately observes t_0 , and neither player observes the state of the world ω .

At every $t \geq t_0$, the sender chooses between acquiring a noisy, binary signal $s_t \in \{g, b\}$ and irreversibly stopping. If the sender chooses to acquire information, she privately observes s_t , and the game proceeds to date $t + 1$. The receiver does not observe the sender’s action or the signal. If the sender stops, she makes a disclosure m to the receiver. The receiver observes m , takes a binary action¹ $a \in \{0, 1\}$, and the game ends.

2.1. Signal technology. Conditional on the state of the world, each signal is an independent draw from a Bernoulli distribution. We assume that $\mathbb{P}(s_t = g | \omega = G) = \mathbb{P}(s_t = b | \omega = B) = p > \frac{1}{2}$. That is, each signal is aligned with the state of the world with probability p ,

¹We relax the binary action assumption in Section 4.2.

and it is misaligned with probability $q := 1 - p$.

Let $s_{t_1}, s_{t_2}, \dots, s_{t_k}$ be some k signals, and let n be the difference between the number of good signals and the number of bad signals among them. We define the *face value* of the signals, π_n , as the conditional probability of the good state

$$\pi_n := \mathbb{P}(\omega = G | s_{t_1}, s_{t_2}, \dots, s_{t_k}) = \frac{\pi_0 p^n}{\pi_0 p^n + (1 - \pi_0) q^n}.$$

It is pinned down by n . That is, adding two contradictory signals to the existing signals does not change their face value.

2.2. Disclosure technology. A disclosure $m \in \cup_{l=0}^{\infty} \{g, b\}^l =: H$ is a sequence of finitely many binary signals. We call the disclosure of length zero the *empty disclosure*. The sender can only disclose a *left truncation* of the signals she has previously acquired. That is, if the sender stops at some date $t > t_0$ after acquiring signals $(s_{t_0}, s_{t_0+1}, \dots, s_{t-1})$, a nonempty disclosure m is feasible if and only if $m = (s_{t-l}, s_{t-2}, \dots, s_{t-1})$ for some $l \leq t - t_0$. The empty disclosure is always feasible. Specifically, if the sender stops at date t_0 without acquiring any signal, it is only feasible to make the empty disclosure. We write $m \preceq h$ if m is a left truncation of h . This defines a partial order \preceq on H .

An interpretation of this disclosure technology is that each signal is time-stamped, and both the content of a signal and the date it is acquired are hard evidence. When the sender stops at date t and disclose a signal from some date $t' < t$, the receiver knows that the sender has been continually acquiring signals since t' the latest. This leads to unraveling, that is, the sender discloses a left truncation even if she is allowed to disclose arbitrary subsets of her signals.²

2.3. Payoffs. The receiver's payoff $u(a, \omega)$ depends on his action a and the state of the world ω . We assume that $u(0, G) = u(0, B) = 0$, $u(1, G) = 1$, and $u(1, B) = -\bar{\beta}/(1 - \bar{\beta})$, where $\bar{\beta} > \pi_0$ is common knowledge. Hence, it is optimal for the receiver to accept (i.e., choose $a = 1$) if and only if the probability of the good state is at least $\bar{\beta}$.

The sender wants the receiver to accept regardless of the state of the world. She receives a unit reward if the receiver accepts and pays a fixed cost $c > 0$ for each signal she acquires. Hence, the sender's payoff given the receiver's action a , the starting date t_0 and the stopping date t is $v(a, t_0, t) = a - c(t - t_0)$. Payoffs are not discounted over time.

Notice that the disclosure is not payoff relevant. The sender can disclose any left truncation of her signals at no cost. This differs from a signaling game (where signals' costs depend

²Precisely, it is without loss of generality to focus on equilibria where the sender discloses a left truncation, and if, for some dates $t' < t'' < t$, the sender discloses $s_{t'}$ but not $s_{t''}$, the receiver believes that $s_{t''} = b$.

on the sender's private information) or a cheap talk game (where feasibility of a message does not depend on the sender's private information). For the receiver, the disclosure is purely informational.

2.4. Strategies. A private history $h \in H$ of the sender is the sequence of signals since the starting date t_0 .³ A stopping strategy of the sender is a (pure) stopping time τ adapted to the natural filtration generated by private histories. In other words, τ is the number of signals the sender acquires before stopping; the calendar time at which the sender stops is $t_0 + \tau$. A pure disclosure strategy of the sender is $\mathbf{m} : H \rightarrow H$ such that $\mathbf{m}(h) \preceq h$ for all h . It selects a feasible disclosure $\mathbf{m}(h)$ in the event that the sender stops after acquiring signals h .

A pure strategy of the receiver is $\mathbf{a} : H \rightarrow \{0, 1\}$. It selects an action $\mathbf{a}(m)$ for every possible disclosure $m \in H$ of the sender. We focus on receiver strategies where his action is a function of the difference between the number of the good and bad signals in the disclosure. Formally, let $n(m)$ be the difference between the number of good and bad signals in a disclosure m . A receiver strategy is a *threshold strategy* if there exists an integer $T > 0$ such that $\mathbf{a}(m) = 1$ if and only if $n(m) \geq T$.

2.5. Beliefs. After observing the sender's disclosure, the receiver forms a belief about the state of the world. Let $\beta(m)$ be the probability he assigned to the good state after seeing disclosure m . A system of beliefs of the receiver is $\beta : H \rightarrow [0, 1]$.

2.6. Equilibrium. A weak perfect Bayesian equilibrium of the game consists of the sender's stopping strategy τ and disclosure strategy \mathbf{m} , and the receiver's strategy \mathbf{a} and system of beliefs β , such that:

(Optimal stopping) Given the sender's disclosure strategy and the receiver's strategy, the sender stops optimally. That is, given \mathbf{m} and \mathbf{a} ,

$$\mathbb{E}[\mathbf{a}(\mathbf{m}(h_\tau)) - c\tau] \geq \mathbb{E}[\mathbf{a}(\mathbf{m}(h_{\tau'})) - c\tau']$$

for all stopping times τ' , where h_τ denotes the private history when the sender stops;

(Optimal disclosure) Given \mathbf{a} , $\mathbf{a}(\mathbf{m}(h)) \geq \mathbf{a}(m')$ for all $m' \preceq h$ and $h \in H$;

(Receiver's sequential rationality) Given β , $\mathbf{a}(m) = 1$ if and only if $\beta(m) \geq \bar{\beta}$;

³We assume that the players' strategies do not depend on calendar time, and accordingly the starting date t_0 is not part of the sender's private history. For the receiver, calendar time is payoff irrelevant. For the sender, only the time elapsed since the starting date matters, and this is precisely the length of her private history. This assumption rules out the possibility that the sender signals the starting date by behaving differently in two otherwise equivalent continuation games.

(Belief consistency) The receiver updates belief $\beta(m)$ using Bayes' rule whenever possible.

Two equilibria are *equivalent* if they give each player the same equilibrium payoff. A *threshold equilibrium* (or simply, an equilibrium) is a weak PBE where the receiver uses a threshold strategy.

3. Threshold Equilibrium

3.1. Threshold stopping strategies. We start by solving the sender's best response to a threshold strategy T of the receiver. Given the receiver's strategy, the sender stops and discloses a left truncation containing at least T more good signals than bad signals whenever possible, so we focus on her optimal stopping problem. Let

$$d(h) := \max_{m \preceq h} n(m)$$

be the maximal difference between good and bad signals among all feasible disclosures given a history h . Intuitively, $d(h)$ determines what sequences of additional signals will lead to acceptance by the receiver. If the sender can already make a disclosure that contains $d < T$ more good signals than bad signals, she only needs to acquire $T - d$ more good signals than bad signals in the future to persuade the receiver. On the other hand, the sender's private belief about the state of the world $\pi_{n(h)}$ determines the distribution of future signals.

The sender's incentive to continue acquiring signals is increasing in $d(h)$ and $n(h)$. Intuitively, increasing $d(h)$ means that the sender needs fewer good signals in the future to persuade the receiver, while increasing $n(h)$ means that each future signal is more likely to be good. This implies that, *in any equilibrium*, the sender never stops with positive possibility at a history h such that $0 < d(h) < T$. Specifically, the sender never stops unsuccessfully if her last signal is good, since she would have stopped one period earlier when her private belief was lower and she needed more signals to persuade the receiver.

Hence, the sender has a best response in threshold stopping strategies. A stopping strategy is a *threshold stopping strategy* if there exist $\bar{t} > 0$ and $\underline{n} \leq 0$ such that the sender stops after acquiring signals h if and only if $d(h) \geq \bar{t}$ or $n(h) - d(h) \leq \underline{n}$.⁴ We denote by $\underline{n} = -\infty$ the strategy that the sender stops if and only if $d(h) \geq \bar{t}$.⁵

Figure 1 shows an example of a threshold stopping strategy with $\bar{t} = 3$ and $\underline{n} = -6$. The set of histories H is partitioned into *states*, each identified by a distinct pair of $n(h)$ and

⁴If the sender uses a threshold stopping strategy, histories such that $n(h) - d(h) \leq \underline{n}$ and $(n(h), d(h)) \neq (\underline{n}, 0)$ are unreachable (an example is provided in Figure 1). The sender is indifferent at those histories, and we assume that the sender stops acquiring signals immediately. The same applies to histories such that $d(h) > \bar{t}$.

⁵This is never an equilibrium strategy, but it may be a best response strategy (see Lemma 2).

$d(h)$. Notice that $d(h) \geq \max\{n(h), 0\}$, since making the empty disclosure and disclosing the entire history are always feasible. States where the sender acquires additional signals are marked by hollow points.




FIGURE 1: State transition given the sender's threshold stopping strategy $(\bar{t}, \underline{n}) = (3, -6)$

The game starts in state $(0, 0)$. If the first signal is good, the state transits to $(1, 1)$, and each time there is a good signal, the state moves from $(n(h), d(h))$ to $(n(h) + 1, d(h) + 1)$. If the first signal is bad, the state transits to $(-1, 0)$, since at history $h = (b)$, there is one bad signal and no good signal, but the maximal difference among all feasible disclosures is zero, which is obtained by making the empty disclosure. A bad signal moves the state from $(n(h), d(h))$ to $(n(h) - 1, \max\{d(h) - 1, 0\})$. Figure 1 maps all possible state transitions before the sender stops. Red (blue) arrows indicate state transitions following a good (bad) signal. Notice that the sender either stops unsuccessfully in state $(-6, 0)$ or stops successfully in a state $(n, 3)$ where $-2 \leq n \leq 3$. States marked by solid gray points are unreachable under the threshold stopping strategy.

The following lemma summarizes the results. For generic values of the cost c of acquiring signals, the sender has a unique best response, and it is a threshold stopping strategy.⁶ Hence, we focus on threshold stopping strategies in the remainder of the paper.

LEMMA 1. *Given any threshold strategy T of the receiver, for all but countably many values of c , the sender has a unique best response, and it is a threshold stopping strategy.*

Proof. See Appendix. □

3.2. The sender's best response. Restricted to threshold stopping strategies, the sender sets $\bar{t} = T$, so we simply need to find the threshold \underline{n} in the sender's best response. We focus on the sender's optimal stopping problems in states $(n, 0)$. Since the sender never stops unsuccessfully, the sender chooses in each state $(n, 0)$ between stopping immediately

⁶For nongeneric values, the sender has two best response threshold stopping strategies (see Lemma 2).

and acquiring additional signals until either state $(n - 1, 0)$ or $(n + T, T)$ is reached. If she reaches state $(n - 1, 0)$, the optimal stopping problem continues; if she reaches state $(n + T, T)$, she stops immediately, discloses an accepted truncation and gets payoff 1. Hence, letting $V(n; \underline{n})$ be the sender's continuation payoff in state $(n, 0)$ when she uses a threshold stopping strategy \underline{n} , we have

$$(1) \quad V(n; \underline{n}) = \frac{\pi_n - \pi_{n-1}}{\pi_{n+T} - \pi_{n-1}} + \frac{\pi_{n+T} - \pi_n}{\pi_{n+T} - \pi_{n-1}} V(n - 1; \underline{n}) - \kappa(n, T)$$

for all $\underline{n} < n \leq 0$, where $\frac{\pi_n - \pi_{n-1}}{\pi_{n+T} - \pi_{n-1}}$ is the probability that the sender reaches state $(n + T, T)$ first, $\frac{\pi_{n+T} - \pi_n}{\pi_{n+T} - \pi_{n-1}}$ is the probability that she reaches state $(n - 1, 0)$ first, and $\kappa(n, T)$ denotes the expected cost of acquiring signals before either state is reached. When she stops unsuccessfully, $V(\underline{n}; \underline{n}) = 0$.

If the cost c of acquiring each signal is sufficiently low, it is optimal for the sender to continue in every state $(n, 0)$. For example, suppose that $T = 1$ and $c \leq q$. That is, the sender stops successfully when she acquires a single good signal. Since the probability of getting a good signal is strictly greater than q , the sender's expected payoff from acquiring an additional signal always outweighs its cost. Therefore, she never stops until she has a good signal (i.e., $\underline{n} = -\infty$). If the cost is not too low, the sender has a unique best response $\underline{n} \leq 0$ for generic model parameters, and it is weakly increasing in c and in T . That is, the sender acquires information more aggressively if the cost per signal is lower, or if the receiver uses a lower threshold.

LEMMA 2. *Given a threshold strategy T of the receiver, the sender has a best response $\underline{n} \in \mathbb{Z}_- \cup \{-\infty\}$. Moreover, there exist $\bar{c}_0 > \bar{c}_{-1} > \bar{c}_{-2} > \dots$ such that:*

- (i) $\underline{n} = -\infty$ is the sender's unique best response if and only if $c \leq \bar{c}_{-\infty}$, where $\bar{c}_{-\infty} = \lim_{\underline{n} \rightarrow -\infty} \bar{c}_{\underline{n}} > 0$;
- (ii) $\underline{n} < 0$ is the sender's best response if and only if $c \in [\bar{c}_{\underline{n}}, \bar{c}_{\underline{n}+1}]$;
- (iii) $\underline{n} = 0$ is the sender's best response if and only if $c \geq \bar{c}_0$;
- (iv) $\bar{c}_{\underline{n}}$ is decreasing in T for all $\underline{n} \in \{-\infty\} \cup \mathbb{Z}_-$.

Proof. See Appendix. □

3.3. The receiver's belief. Given a threshold strategy T of the receiver and a threshold stopping strategy of the sender where $\underline{n} < 0$, the sender stops unsuccessfully only in state $(\underline{n}, 0)$. Hence, the receiver's posterior belief if the sender fails to persuade him, i.e., $\mathbb{P}[\omega = G | n(m) < T]$, is the same as the sender's private belief $\pi_{\underline{n}}$.

Recall that in every state $(n, 0)$ where the sender acquires information, she reaches state

$(n-1, 0)$ with probability $\frac{\pi_{n+T}-\pi_n}{\pi_{n+T}-\pi_{n-1}}$. Hence, the sender stops unsuccessfully with probability

$$\prod_{n=\underline{n}+1}^0 \frac{\pi_{n+T}-\pi_n}{\pi_{n+T}-\pi_{n-1}} = \frac{\pi_0 p^{\underline{n}} + (1-\pi_0)q^{\underline{n}}}{\lambda^{\underline{n}}},$$

where

$$\lambda := \frac{pq(p^T - q^T)}{p^{T+1} - q^{T+1}}.$$

By the martingale property of beliefs, the receiver's *average* posterior belief conditional on seeing an accepted disclosure m , i.e., $\mathbb{P}[\omega = G | n(m) \geq T]$, is

$$\tilde{\beta}(\underline{n}, T) := \frac{\pi_0(\lambda^{\underline{n}} - p^{\underline{n}})}{\pi_0(\lambda^{\underline{n}} - p^{\underline{n}}) + (1-\pi_0)(\lambda^{\underline{n}} - q^{\underline{n}})},$$

which is increasing in \underline{n} and T . The following lemma shows that we can focus on equilibria where the receiver's posterior belief after seeing any accepted disclosure is precisely $\tilde{\beta}(\underline{n}, T)$. That is, $\beta(m) = \tilde{\beta}(\underline{n}, T)$ for all disclosures m such that $n(m) = T$.

LEMMA 3. *Every equilibrium is equivalent to an equilibrium where the receiver's posterior belief after seeing any accepted disclosure is $\tilde{\beta}(\underline{n}, T)$.*

Proof. Given an equilibrium, we replace the sender's disclosure strategy by the following: she discloses the shortest optimal left truncations. This merely changes how the sender breaks ties when disclosing, hence it yields an equivalent equilibrium.⁷ Moreover, given the tie-breaking rule, disclosed signals and omitted signals are independent. Intuitively, consider two disclosures m_1 and m_2 that are accepted on the equilibrium path. Since $n(m_1) = n(m_2) = T$, they differ by the same number of good and bad signals. Suppose that m_1 contains k more good (bad) signals than m_2 . Regardless of the omitted signals, m_1 is disclosed $(pq)^k$ times more likely than m_2 . Therefore, the receiver cannot infer the omitted signals from the disclosure. \square

That the receiver cannot infer the omitted signals from the disclosure is not an artifact of the particular tie-breaking rule.⁸ It can be justified by introducing a small uncertainty to the receiver's payoff à la Harsanyi (1973).⁹ If the sender is uncertain about the receiver's

⁷In the original equilibrium, the receiver's average posterior belief when persuaded is $\tilde{\beta}(\underline{n}, T)$. We are to show that given the new disclosure strategy, the receiver's posterior belief after seeing any accepted disclosure is $\tilde{\beta}(\underline{n}, T)$, hence the receiver's sequential rationality is intact.

⁸It is true as long as the sender does not signal her omitted evidence through tie-breaking of her disclosure strategy. That is, for all $h, h' \in H$ such that $\arg\max_{m \preceq_h} \mathbf{a}(m) = \arg\max_{m' \preceq_{h'}} \mathbf{a}(m')$, $\mathbf{m}(h) = \mathbf{m}(h')$. The tie-breaking rule in favor of the shortest truncation is just one example of many.

⁹Since the sender's disclosure is payoff irrelevant, our model is a game with nongeneric payoffs, and

threshold belief for acceptance, the sender discloses the left-truncation that maximizes the receiver's posterior belief. Hence, the receiver's posterior belief conditional on being persuaded must be the same in any purifiable equilibrium (i.e., robust to small payoff uncertainties).

To summarize, despite the different settings of our model and Bayesian persuasion, the receiver's posterior belief is supported on two points. On the equilibrium path, the receiver's posterior belief is $\pi_{\underline{n}}$ if he sees a disclosure m such that $n(m) < T$ and $\tilde{\beta}(\underline{n}, T)$ if he sees a disclosure m such that $n(m) \geq T$.

3.4. Equilibria. The receiver's threshold T and the sender's threshold \underline{n} determines the players' payoffs. Since the sender's best response is essentially unique, we focus on the receiver's equilibrium threshold strategy T .

By Lemma 2, the sender has a best response $\underline{n} < 0$ if and only if T is bounded by some T^* . If $T > T^*$, acquiring no signal (i.e., $\underline{n} = 0$) is a best response of the sender, and this is always an equilibrium. Since the receiver can form any belief upon seeing a nonempty disclosure, he can form beliefs that justify the threshold strategy.

If $T \leq T^*$, the sender has a best response $\underline{n} < 0$. A *persuasion equilibrium* is an equilibrium where the sender is persuaded with positive probability, and it exists if and only if the receiver's posterior belief $\tilde{\beta}(\underline{n}, T) \geq \bar{\beta}$. The following proposition shows a tight upper bound on c such that a persuasion equilibrium exists.

PROPOSITION 4. *Let \bar{T} be the smallest integer such that $\pi_{\bar{T}} \geq \bar{\beta}$, and $\bar{c}_0(\bar{T})$ the threshold defined in Lemma 2 given $T = \bar{T}$. A persuasion equilibrium exists only if $c \leq \bar{c}_0(\bar{T})$, and a persuasion equilibrium exists if $c = \bar{c}_0(\bar{T})$.*

Proof. Observe that $\tilde{\beta}(-1, T) = \pi_T$ for all $T > 0$. Intuitively, if the sender stops unsuccessfully whenever her belief is below the prior, full disclosure is optimal when she stops successfully, so the receiver take the disclosure at face value. If $c > \bar{c}_0(\bar{T})$, the sender's unique best response to any $T \geq \bar{T}$ is $\underline{n} = 0$, and the receiver's posterior belief $\tilde{\beta}(\underline{n}, T) \leq \pi_T < \bar{\beta}$ for all $T < \bar{T}$ and $\underline{n} \leq 0$. Therefore, no persuasion equilibrium exists. If $c = \bar{c}_0(\bar{T})$, the threshold stopping strategy $\underline{n} = -1$ is a best response to the receiver's threshold strategy \bar{T} . And since $\tilde{\beta}(-1, \bar{T}) = \pi_{\bar{T}} \geq \bar{\beta}$, \bar{T} is an equilibrium threshold strategy. \square

The condition $c \leq \bar{c}_0(\bar{T})$ is not sufficient for equilibrium existence due to the discreteness of the players' strategies.¹⁰ Suppose that a persuasion equilibrium exists, and let T_* be

purification can be used as a refinement. Jiang (2024b) applies purification as a refinement in disclosure games with exogenous evidence. Relatedly, Bhaskar and Thomas (2019) applies it as a refinement in dynamic games, and Diehl and Kuzmics (2021) in cheap talk games.

¹⁰For example, suppose that $\bar{c}_0(\bar{T} + 1) < c < \bar{c}_{-1}(\bar{T})$ and $\pi_{\bar{T}} = \bar{\beta}$. The sender's does not acquire any signal if $T > \bar{T}$. For all to $T \leq \bar{T}$, her best response $\underline{n} < -1$, hence $\tilde{\beta}(\underline{n}, T) < \pi_T \leq \bar{\beta}$. This leads to nonexistence

the receiver's smallest equilibrium threshold. Recall that the sender's best response \underline{n} is a weakly increasing step function of T , and $\tilde{\beta}(\underline{n}, T)$ is increasing in \underline{n} and T . Hence, every integer $T \in [T_*, T^*]$ is an equilibrium threshold strategy, and the set of persuasion equilibria is ordered by the receiver's posterior belief $\tilde{\beta}(\underline{n}, T)$ when he is persuaded.

4. Discussion and Extensions

4.1. Mandatory full disclosure. We study a benchmark where only full disclosure is possible. This is equivalent to a symmetric information game where signals are publicly observed. Hence, the players have common belief $\pi_{n(h)}$ about the state at every history h .

The following lemma shows that a persuasion equilibrium exists under mandatory full disclosure if and only if the cost c of acquiring signals is bounded by the same upper bound in Proposition 4. That is, if signals are too costly such that persuasion is not possible when left truncation is permitted, mandating full disclosure cannot help.

LEMMA 5. *A persuasion equilibrium exists under mandatory full disclosure if and only if $c \leq \bar{c}_0(\bar{T})$.*

Proof. Every equilibrium is equivalent to one where the sender acquires signals if and only if $\underline{n} < n(h) < \bar{T}$ for some threshold $\underline{n} \in \mathbb{Z}_- \cup \{-\infty\}$. We show in Appendix A.4 that the sender's payoff is a single-peaked function of \underline{n} . Therefore, a persuasion equilibrium exists if and only if the sender's payoff from choosing $\underline{n} = -1$ is nonnegative. Notice that this strategy induces the same outcome as the strategy profile $T = \bar{T}$ and $\underline{n} = -1$ in the model where left truncation is permitted. Hence, the sender's payoff is the same. By Lemma 2, it is nonnegative if and only if $c \leq \bar{c}_0(\bar{T})$. \square

The mandatory full disclosure benchmark is equivalent to costly Bayesian persuasion with the log-likelihood ratio cost of experiments (see Morris and Strack (2019) and Jiang (2024a)). Since we model learning in discrete time, we have the additional constraint that the posterior belief is supported on a subset of $\{\pi_n\}_{n \in \mathbb{Z}}$. For the next result, we assume that this constraint is satisfied, so mandatory full disclosure yields the Bayesian persuasion solution. We compare the players' equilibrium payoffs when left truncation is permitted to their payoffs in the Bayesian persuasion solution. We show that allowing left truncation benefits the receiver but hurts the sender.

PROPOSITION 6. *Assuming that $\pi_{\bar{T}} = \bar{\beta}$ and a persuasion equilibrium exists when left truncation is permitted:*

of persuasion equilibrium.

- (i) *The receiver's payoff is weakly higher in all equilibria and strictly higher in all but one equilibrium than in the Bayesian persuasion solution;*
- (ii) *The sender's payoff is lower in all persuasion equilibria than in the Bayesian persuasion solution if $c \neq c_B$ for some $c_B \in [\bar{c}_{-1}(\bar{T}), \bar{c}_0(\bar{T})]$; if $c = c_B$, the Bayesian persuasion solution is equivalent to the unique persuasion equilibrium.*

Proof. (i) The receiver's payoff in the Bayesian persuasion solution is 0. That in any equilibrium is nonnegative. Since equilibria are ordered by the receiver's posterior belief $\tilde{\beta}(\underline{n}, T)$, except for possibly one equilibrium, $\tilde{\beta}(\underline{n}, T) > \bar{\beta}$, and the receiver's payoff is positive.

(ii) Let $\underline{\beta}$ and $\bar{\beta}$ be the posterior beliefs induced by the Bayesian persuasion solution. $\underline{\beta}$ is decreasing in c , so there exists a unique c_B such that $\underline{\beta} = \pi_{-1}$. When $c = c_B$, the same posterior distribution is induced by the strategy profile $T = \bar{T}$ and $\underline{n} = -1$. Since the Bayesian persuasion solution is ex ante optimal for the sender and $\tilde{\beta}(-1, \bar{T}) = \bar{\beta}$, this is the unique equilibrium strategy profile. By Lemma 2, $c_B \in [\bar{c}_{-1}(\bar{T}), \bar{c}_0(\bar{T})]$.

Let T and \underline{n} be a persuasion equilibrium strategy profile, and assume that $c \neq c_B$. The Bayesian persuasion solution gives the sender a higher payoff than the experiment that induces posterior beliefs $\pi_{\underline{n}}$ and $\tilde{\beta}(\underline{n}, T)$. The latter gives the sender a weakly higher payoff than her equilibrium payoff, since the sender's posterior distribution in the equilibrium is a mean-preserving spread of the two-point distribution $\{\pi_{\underline{n}}, \tilde{\beta}(\underline{n}, T)\}$, so the cost of acquiring signals in the equilibrium is at least the log-likelihood ratio cost of the experiment. \square

4.2. Finite receiver actions. Our analysis can be extended to finite receiver actions. Suppose that the receiver chooses one of $K + 1$ actions $a_0 < a_1 < \dots < a_K$. We assume that the receiver's strategy is according to some thresholds $0 < T_1 \leq T_2 \leq \dots \leq T_K$, that is, after seeing disclosure m , he takes action a_0 if $n(m) < T_1$, takes action a_k if $T_k \leq n(m) < T_{k+1}$ for all $1 \leq k \leq K - 1$, and takes action a_K if $n(m) \geq T_K$.

Lemma 2 shows that, if $K = 1$, the sender has a best response threshold \underline{n}_1 to the receiver's threshold strategy T_1 . Now consider the case of $K = 2$, and the receiver uses thresholds T_1, T_2 . The sender's best response features a threshold \underline{n}_2 such that in state $(\underline{n}_2, 0)$, the sender forgoes the possibility of reaching the higher threshold T_2 ; if $\underline{n}_2 > \underline{n}_1$, she will continue acquiring signals until she reaches the lower threshold T_1 or state $(\underline{n}_1, 0)$. The threshold \underline{n}_2 can be solved from a dynamic programming problem similar to (1), but the sender's continuation payoff in state $(\underline{n}_2, 0)$ is that from continuing according to the threshold stopping strategy (T_1, \underline{n}_1) . It is easy to see that \underline{n}_2 is a weakly decreasing step function of a_2 . If action a_2 is sufficiently lucrative compared to a_1 for the sender, $\underline{n}_2 \leq \underline{n}_1$, and the sender never stops at the lower threshold T_1 . This is similar to the result in Bayesian persuasion: the sender never induces an action if her payoff given the action lies below the concavification

of her indirect utility function. By induction on K , the sender has an essentially unique best response featuring thresholds $\underline{n}_1 \leq \underline{n}_2 \leq \dots \underline{n}_K \leq 0$, such that she continues in state $(n, n + d)$ if and only if $\underline{n}_k \leq n$ and $0 \leq d < T_k$ for some k .¹¹

If $\underline{n}_k < \underline{n}_{k+1}$, the receiver's posterior belief $\tilde{\beta}_k$ after seeing any disclosure with T_k more good signals than bad signals can be similarly derived as in the binary action case. The receiver's strategy $\{T_k\}_{k=1}^K$ and the sender's best response $\{\underline{n}_k\}_{k=1}^K$ are supported by an equilibrium if and only if action a_k is receiver optimal given belief $\tilde{\beta}_k$ for all k such that $\underline{n}_k < \underline{n}_{k+1}$.

Notice that in equilibrium, all K actions may be taken with positive probability. This is different from the Bayesian persuasion solution, where at most two actions are induced. This is because in Bayesian persuasion, the sender can induce any posterior distribution that is a mean-preserving spread of the prior, but in our model, the receiver's posterior belief is endogenously determined by the players' thresholds. This restricts the set of posterior distributions the sender can induce.

5. Conclusion

We study persuasion via sequentially acquired evidence. A sender privately and sequentially acquire noisy signals about a binary state of the world and can disclose a left truncation to persuade a receiver. This relaxes the commitment assumption in Bayesian persuasion. Indeed, the Bayesian persuasion solution is the equilibrium of a game where full disclosure of signals is mandatory.

We show that in equilibrium, the receiver's posterior belief is supported on two points when his action is binary. This is akin to the Bayesian persuasion solution. When the cost is not too high, there exists a continuum of persuasion equilibria. Compared to the Bayesian persuasion solution, the sender is worse off in all equilibria, while the receiver is weakly better off.

¹¹If the sender never stops at threshold T_k , it is without loss to set $\underline{n}_k = \underline{n}_{k+1}$.

APPENDIX A. OMITTED PROOFS

We first prove Lemma 2 and some auxiliary results on the sender's continuation value. We then prove Lemma 1. The last section contains results relating to the mandatory full disclosure benchmark, which are used to prove Lemma 5.

A.1. Proof of Lemma 2

Proof. By Proposition 10 of Jiang (2024a),

$$\kappa(n, t) = H(\pi_n) - \frac{\pi_n - \pi_{n-1}}{\pi_{n+t} - \pi_{n-1}} H(\pi_{n+t}) - \frac{\pi_{n+t} - \pi_n}{\pi_{n+t} - \pi_{n-1}} H(\pi_{n-1}),$$

where $H(x) := -\frac{c}{(p-q)\ln \frac{p}{q}}(2x-1)\ln\left(\frac{x}{1-x}\right)$. Notice that $\ln\left(\frac{\pi_n}{1-\pi_n}\right) - \ln\left(\frac{\pi_0}{1-\pi_0}\right) = n \ln \frac{p}{q}$. Therefore,

$$\begin{aligned} \kappa(n, t) &= \frac{c}{p-q} \left[n(2\pi_n - 1) - \frac{\pi_n - \pi_{n-1}}{\pi_{n+T} - \pi_{n-1}} (n+T)(2\pi_{n+T} - 1) - \frac{\pi_{n+T} - \pi_n}{\pi_{n+T} - \pi_{n-1}} (n-1)(2\pi_{n-1} - 1) \right] \\ &= -\frac{c}{p-q} \left[T(2\pi_n - 1) - \frac{\pi_{n+T} - \pi_n}{\pi_{n+T} - \pi_{n-1}} (T+1)(2\pi_{n-1} - 1) \right], \end{aligned}$$

where

$$\frac{\pi_{n+T} - \pi_n}{\pi_{n+T} - \pi_{n-1}} = \lambda \frac{\pi_0 p^{n-1} + (1-\pi_0)q^{n-1}}{\pi_0 p^n + (1-\pi_0)q^n}$$

is the probability of reaching state $(n-1, 0)$ before $(n+T, T)$. Hence, letting

$$(A.1) \quad \tilde{V}(n; \underline{n}) = [\pi_0 p^n + (1-\pi_0)q^n] \left[V(n; \underline{n}) - 1 + \frac{c}{p-q} T(2\pi_n - 1) \right],$$

the difference equation (1) can be rewritten as

$$\tilde{V}(n; \underline{n}) = \lambda \tilde{V}(n-1; \underline{n}) + \frac{c\lambda}{p-q} [\pi_0 p^{n-1} - (1-\pi_0)q^{n-1}].$$

It has a unique solution

$$(A.2) \quad \tilde{V}(n; \underline{n}) = \lambda^n \tilde{C}(\underline{n}) + \frac{c\lambda}{p-q} \left(\frac{\pi_0 p^n}{p-\lambda} - \frac{(1-\pi_0)q^n}{q-\lambda} \right),$$

where

$$(A.3) \quad \tilde{C}(\underline{n}) = \left[\frac{c}{p-q} \left(T - \frac{\lambda}{p-\lambda} \right) - 1 \right] \frac{\pi_0 p^n}{\lambda^n} + \left[\frac{c}{p-q} \left(\frac{\lambda}{q-\lambda} - T \right) - 1 \right] \frac{(1-\pi_0)q^n}{\lambda^n}$$

is solved from the boundary condition $V(\underline{n}; \underline{n}) = 0$. Since $V(\underline{n}; \underline{n})$ is an affine transformation of $\tilde{C}(\underline{n})$ for all $\underline{n} \leq n \leq 0$, \underline{n} is a best response if and only if it maximizes $\tilde{C}(\underline{n})$.

Notice that

$$\frac{\lambda}{q-\lambda} - T = \sum_{t=1}^T \left[\left(\frac{p}{q} \right)^t - 1 \right] > \sum_{t=1}^T \left[1 - \left(\frac{q}{p} \right)^t \right] = T - \frac{\lambda}{p-\lambda} > 0$$

for all $T \geq 1$. Therefore, if

$$c \leq \bar{c}_{-\infty} := \frac{p-q}{\frac{\lambda}{q-\lambda} - T},$$

$\tilde{C}(\underline{n})$ is decreasing in \underline{n} , and the sender's unique best response is $\underline{n} = -\infty$. Notice that $\bar{c}_{-\infty} > 0$ and is decreasing in T . If $c \geq \frac{p-q}{T - \frac{\lambda}{p-\lambda}}$, $\tilde{C}(\underline{n})$ is increasing in \underline{n} , and the sender's unique best response is $\underline{n} = 0$.

Now suppose that $c \in \left(\bar{c}_{-\infty}, \frac{p-q}{T - \frac{\lambda}{p-\lambda}} \right)$. We show that \tilde{C} is single-peaked. Taking derivative of (A.3) with respect to \underline{n} , we have

$$(A.4) \quad \tilde{C}'(\underline{n}) = \left[\frac{c}{p-q} \left(T - \frac{\lambda}{p-\lambda} \right) - 1 \right] \pi_0 \left(\frac{p}{\lambda} \right)^{\underline{n}} \ln \frac{p}{q} + \ln \frac{q}{\lambda} \tilde{C}(\underline{n}).$$

It has at most one zero on $(-\infty, 0]$, since

$$\tilde{C}''(\underline{n}) = \left[\frac{c}{p-q} \left(T - \frac{\lambda}{p-\lambda} \right) - 1 \right] \pi_0 \left(\frac{p}{\lambda} \right)^{\underline{n}} \ln \frac{p}{\lambda} \ln \frac{p}{q} + \ln \frac{q}{\lambda} \tilde{C}'(\underline{n}) < 0$$

if $\tilde{C}'(\underline{n}) = 0$. Moreover, (A.4) is positive when \underline{n} is sufficiently large in absolute value. Therefore, \tilde{C} is maximized at at most two (adjacent) integers in $\underline{n} \in \mathbb{Z}_-$ if (A.4) is negative at $\underline{n} = 0$; otherwise, it is maximized at $\underline{n} = 0$.

By (A.3), for all $\underline{n} < 0$, $\tilde{C}(\underline{n}) > \tilde{C}(\underline{n} + 1)$ if and only if

$$\left[\frac{c}{p-q} \left(T - \frac{\lambda}{p-\lambda} \right) - 1 \right] \pi_{\underline{n}} \frac{p-\lambda}{\lambda} + \left[\frac{c}{p-q} \left(\frac{\lambda}{q-\lambda} - T \right) - 1 \right] (1 - \pi_{\underline{n}}) \frac{q-\lambda}{\lambda} < 0,$$

that is,

$$(A.5) \quad \frac{c}{p-q} \left[1 - 2\pi_{\underline{n}} + T \left(\frac{\pi_{\underline{n}} - q}{\lambda} + 1 - 2\pi_{\underline{n}} \right) \right] < \frac{\pi_{\underline{n}} p + (1 - \pi_{\underline{n}}) q}{\lambda} - 1.$$

Taking derivative of the bracket on the left-hand side of (A.5) with respect to T yields

$$(A.6) \quad \frac{\pi_{\underline{n}} - q}{pq} \underbrace{\frac{p^{2T+1} + q^{2T+1} - p^T q^T \left[1 + T(p - q) \ln \frac{p}{q} \right]}{(p^T - q^T)^2}}_{=: \chi} + 1 - 2\pi_{\underline{n}}.$$

Notice that

$$\begin{aligned} \chi - q &= \frac{p^T q^T (p - q)}{(p^T - q^T)^2} \left[\left(\frac{p}{q} \right)^T - 1 - T \ln \frac{p}{q} \right] > 0, \\ \chi - p &= -\frac{p^T q^T (p - q)}{(p^T - q^T)^2} \left[\left(\frac{q}{p} \right)^T - 1 - T \ln \frac{q}{p} \right] < 0, \end{aligned}$$

and (A.6) is linear in $\pi_{\underline{n}}$. Hence, (A.6) is bounded from below by

$$\min \left\{ 1 - \frac{\chi}{p}, \frac{\chi}{q} - 1 \right\} > 0.$$

That is, the bracket on the left-hand side of (A.5) is increasing in T . Since it is linear in $\pi_{\underline{n}}$, it is bounded from below by

$$\min \left\{ 1 + T \left(-\frac{q}{\lambda} + 1 \right), -1 + T \left(\frac{p}{\lambda} - 1 \right) \right\} \geq \min \left\{ \frac{p - q}{q}, \frac{p - q}{p} \right\} > 0.$$

Since $\lambda < q$ is increasing in T , the right-hand side of (A.5) is positive and decreasing in T . Hence, (A.5) defines an upper bound $\bar{c}_{\underline{n}+1} > 0$ of c . $\bar{c}_{\underline{n}}$ is decreasing in T , and $\tilde{C}(\underline{n}) > \tilde{C}(\underline{n}+1)$ if and only if $c < \bar{c}_{\underline{n}+1}$. Since \tilde{C} is single-peaked, $\bar{c}_{\underline{n}}$ is increasing in \underline{n} . As $\underline{n} \rightarrow -\infty$, (A.5) becomes $c < \bar{c}_{-\infty}$. Hence, $\lim_{\underline{n} \rightarrow -\infty} \bar{c}_{\underline{n}} = \bar{c}_{-\infty}$. \square

A.2. Results Relating to the Sender's Continuation Value

We first show that, given a threshold strategy of the receiver and a best response threshold stopping strategy of the sender, the sender's continuation value in state $(n, 0)$ is increasing in n .

LEMMA A.1. *Fix a threshold strategy T of the receiver, and let $\underline{n} < 0$ be a best response threshold stopping strategy of the sender. The sender's continuation value $V(n; \underline{n})$ is increasing in n for all $n > \underline{n}$.*

Proof. By (A.1), (A.2) and (A.3),

$$\begin{aligned}
V(n; \underline{n}) &= \lambda^n \tilde{C}(\underline{n}) + \frac{c\lambda}{p-q} \left(\frac{\pi_0 p^n}{p-\lambda} - \frac{(1-\pi_0)q^n}{q-\lambda} \right) \\
&= \left[\frac{c}{p-q} \left(T - \frac{\lambda}{p-\lambda} \right) - 1 \right] \pi_n \left[\left(\frac{p}{\lambda} \right)^{\underline{n}-n} - 1 \right] \\
&\quad + \left[\frac{c}{p-q} \left(\frac{\lambda}{q-\lambda} - T \right) - 1 \right] (1-\pi_n) \left[\left(\frac{q}{\lambda} \right)^{\underline{n}-n} - 1 \right].
\end{aligned}
\tag{A.7}$$

If $\underline{n} = -\infty$, that is, $c \leq \bar{c}_{-\infty}$,

$$V(n; -\infty) = - \left[\frac{c}{p-q} \left(T - \frac{\lambda}{p-\lambda} \right) - 1 \right] \pi_n - \left[\frac{c}{p-q} \left(\frac{\lambda}{q-\lambda} - T \right) - 1 \right] (1 - \pi_n).$$

Recall that $\frac{\lambda}{q-\lambda} - T > T - \frac{\lambda}{p-\lambda}$. Hence, $V(n; -\infty)$ is increasing in n .

If $-\infty < \underline{n} < 0$, $c \in (\bar{c}_{-\infty}, \bar{c}_0)$. That is,

$$\frac{c}{p-q} \left(T - \frac{\lambda}{p-\lambda} \right) - 1 < 0 < \frac{c}{p-q} \left(\frac{\lambda}{q-\lambda} - T \right) - 1.$$

Hence, by (A.7), for all $n > \underline{n}$, $V(n+1; \underline{n}+1) > V(n; \underline{n})$. But \underline{n} is a best response of the sender, so $V(n; \underline{n}) \geq V(n+1; \underline{n}+1)$. Therefore, $V(n+1; \underline{n}) > V(n; \underline{n})$. That is, the sender's continuation value $V(n; \underline{n})$ is increasing in n for all $n > \underline{n}$. \square

Fixing any $n > \underline{n}$, we now show that the sender's continuation value in state $(n+d, d)$, which we denote by $\hat{V}(n+d, d)$, is increasing in d .

LEMMA A.2. *The sender's continuation value $\hat{V}(n+d, d)$ is increasing in d for all $n > \underline{n}$ and $0 \leq d \leq T$.*

Proof. Notice that for all $n > \underline{n}$ and $0 < d < T$,

$$(A.8) \quad \hat{V}(n+d, d) = \gamma_{n+d} \hat{V}(n+d+1, d+1) + (1-\gamma_{n+d}) \hat{V}(n+d-1, d-1) - c,$$

where $\pi_{n+d} = \pi_{n+d}p + (1-\pi_{n+d})q$ is the probability that the next signal is good in state $(n+d, d)$. (A.8) can be rewritten as

$$(A.9) \quad \gamma_{n+d} [\hat{V}(n+d+1, d+1) - \hat{V}(n+d, d)] = (1-\gamma_{n+d}) [\hat{V}(n+d, d) - \hat{V}(n+d-1, d-1)] + c,$$

In state $(n, 0)$, we have

$$\hat{V}(n, 0) = \gamma_n \hat{V}(n+1, 1) + (1 - \gamma_{n+d}) \hat{V}(n-1, 0) - c.$$

By Lemma A.1, $\hat{V}(n, 0) > \hat{V}(n-1, 0)$. Therefore, $\hat{V}(n+1, 1) > \hat{V}(n, 0)$. By (A.9), we conclude that $\hat{V}(n+d+1, d+1) > \hat{V}(n+d, d)$ for all $n > \underline{n}$ and $0 \leq d < T$. \square

As a corollary, if \underline{n} is the sender's unique best response, i.e., $c \in (\bar{c}_{\underline{n}}, \bar{c}_{\underline{n}+1})$, the sender's continuation payoff is positive in all states (n, d) such that $n > \underline{n}$ and $0 \leq d \leq T$.

COROLLARY A.3. *If $c \in (\bar{c}_{\underline{n}}, \bar{c}_{\underline{n}+1})$, the sender's continuation value $\hat{V}(n+d, d) > 0$ for all $n > \underline{n}$ and $0 \leq d \leq T$.*

A.3. Proof of Lemma 1

Proof. Given a threshold strategy T of the receiver, the sender discloses an accepted truncation whenever possible, so she gets payoff $\mathbb{1}(d(h) \geq T)$ from stopping at history h . Her optimal stopping problem is to choose τ to maximize $\mathbb{E}[\mathbb{1}(d(h_\tau) \geq T) - c\tau]$.

We first show that the sender has a best response in threshold stopping strategies. Let τ be an optimal stopping strategy, and at every history h that is reached with positive probability under τ , denote by $\hat{V}(h)$ the sender's continuation payoff. We show that $\hat{V}(h)$ is a function of $n(h)$ and $d(h)$, and it is weakly increasing in $d(h)$. Let h and h' be two histories such that $n(h) = n(h')$ and $d(h) \geq d(h')$. The distributions of future signals conditional on h and h' are the same, and for all sequences of future signals \tilde{h} , $d((h, \tilde{h})) \geq d((h', \tilde{h}))$. Consider the following strategy in the continuation game starting from h : the sender stops at (h, \tilde{h}) if and only if she stops at (h', \tilde{h}) under τ . It gives the sender at least payoff $\hat{V}(h')$ in the continuation game. Therefore, $\hat{V}(h) \geq \hat{V}(h')$. This also implies that, if the sender continues at h' under τ , it is optimal to continue at h . By abuse of notation, we will write $\hat{V}(n(h), d(h))$ for $\hat{V}(h)$.

Let $\underline{n} \in \mathbb{Z}_- \cup \{-\infty\}$ be the infimum of integers n such that state $(n, 0)$ is reached under τ .¹² We show that $\hat{V}(0, 0) = V(0; \underline{n})$. That is, the threshold stopping strategy \underline{n} gives the sender the same payoff as τ . Hence, it is optimal.

If $\underline{n} = 0$, this is trivial. If $\underline{n} < 0$, let $\hat{n} > \underline{n}$ be the largest integer such that state $(\hat{n} + T, T)$ is reached under τ . That is, it is optimal to continue in all states $(\hat{n} + d, d)$ where $0 \leq d < T$. By monotonicity, it is optimal to continue in all states $(\hat{n} - 1 + d, d)$ where $1 \leq d < T$; it is also optimal to continue in state $(\hat{n} - 1, 0)$ if $\hat{n} - 1 > \underline{n}$. Hence, by induction on n , it is

¹²Since τ is not necessarily stationary, we say state (n, d) is reached under τ if there exists a history h that is reached with positive probability under τ , and $n(h) = n$ and $d(h) = d$.

optimal for the sender to continue in all states (n, d) such that $\underline{n} < n \leq \hat{n}$ and $0 \leq d < T$. That is, \underline{n} is optimal among all threshold strategies, and $\hat{V}(\hat{n}, 0) = V(\hat{n}; \underline{n})$. We are left to show that $\hat{n} = 0$. Suppose by way of contradiction that $\hat{n} < 0$. Then $\hat{V}(0, 0) < \hat{V}(\hat{n}, 0) - c$. This is because the sender needs to reach state $(\hat{n}, 0)$ before she can persuade the receiver, and this requires acquiring at least one signal. But by A.1, $V(n; \underline{n})$ is increasing in n , so $\hat{V}(0; \underline{n}) > V(\hat{n}; \underline{n}) = \hat{V}(\hat{n}, 0)$. That is, the threshold strategy is a profitable deviation, a contradiction.

We now show uniqueness. By Lemma 2 and Corollary A.3, for all but countably many values of c , $\hat{V}(n + d, d) > 0$ for all $n > \underline{n}$ and $0 \leq d \leq T$. Hence, under all optimal stopping strategies, the sender strictly prefers continuing when she continues, and she stops if $n(h) = \underline{n}$ and $d(h) = 0$. Hence, the threshold strategy \underline{n} is the unique optimal stopping strategy. \square

A.4. Results Relating to Mandatory Full Disclosure

LEMMA A.4. *Let $v(\underline{n})$ be the sender's payoff from acquiring signals if and only if $\underline{n} < n(h) < \bar{T}$ in the game with mandatory full disclosure. $v(\underline{n})$ is a single-peaked function.*

Proof. By Jiang (2024a), for all $\underline{n} < 0$,

$$(A.10) \quad v(\underline{n}) = \frac{\pi_0 - \pi_{\underline{n}}}{\pi_{\bar{T}} - \pi_{\underline{n}}} - \left[H(\pi_0) - \frac{\pi_0 - \pi_{\underline{n}}}{\pi_{\bar{T}} - \pi_{\underline{n}}} H(\pi_{\bar{T}}) - \frac{\pi_{\bar{T}} - \pi_0}{\pi_{\bar{T}} - \pi_{\underline{n}}} H(\pi_{\underline{n}}) \right],$$

Since $\pi_{\underline{n}}$ is increasing in \underline{n} , it is sufficient to show that the right-hand side of (A.10) is a single-peaked function of $\pi_{\underline{n}}$. With abuse of notation, we write $v(\pi_{\underline{n}})$ for $v(\underline{n})$. Notice that

$$v'(\pi_{\underline{n}}) = -\frac{\pi_{\bar{T}} - \pi_0}{(\pi_{\bar{T}} - \pi_{\underline{n}})^2} [1 + H(\pi_{\bar{T}}) - H(\pi_{\underline{n}})] + \frac{\pi_{\bar{T}} - \pi_0}{\pi_{\bar{T}} - \pi_{\underline{n}}} H'(\pi_{\underline{n}}),$$

and

$$v''(\pi_{\underline{n}}) = \frac{2}{\pi_{\bar{T}} - \pi_{\underline{n}}} v'(\pi_{\underline{n}}) - \frac{\pi_{\bar{T}} - \pi_0}{\pi_{\bar{T}} + \pi_{\underline{n}}} H''(\pi_{\underline{n}}).$$

Since H is concave, $v''(\pi_{\underline{n}}) < 0$ if $v'(\pi_{\underline{n}}) = 0$. That is, (A.10) is a single-peaked function of $\pi_{\underline{n}}$. \square

REFERENCES

Bhaskar, V. and Thomas, C. (2019). Community enforcement of trust with bounded memory. *The Review of Economic Studies*, 86(3):1010–1032.

- Bull, J. and Watson, J. (2007). Hard evidence and mechanism design. *Games and Economic Behavior*, 58(1):75–93.
- Caplin, A., Dean, M., and Leahy, J. (2022). Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *Journal of Political Economy*, 130(6):1676–1715.
- Di Tillio, A., Ottaviani, M., and Sørensen, P. N. (2021). Strategic sample selection. *Econometrica*, 89(2):911–953.
- Diehl, C. and Kuzmics, C. (2021). The (non-)robustness of influential cheap talk equilibria when the sender’s preferences are state independent. *International Journal of Game Theory*, 50(4):911–925.
- Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research*, 23(1):123–145.
- Felgenhauer, M. and Schulte, E. (2014). Strategic private experimentation. *American Economic Journal: Microeconomics*, 6(4):74–105.
- Gentzkow, M. and Kamenica, E. (2014). Costly persuasion. *American Economic Review*, 104(5):457–462.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483.
- Guo, Y. and Shmaya, E. (2021). Costly miscalibration. *Theoretical Economics*, 16(2):477–506.
- Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1):1–23.
- Hart, S., Kremer, I., and Perry, M. (2017). Evidence games: Truth and commitment. *American Economic Review*, 107(3):690–713.
- Henry, E. (2009). Strategic disclosure of research results: The cost of proving your honesty. *The Economic Journal*, 119(539):1036–1064.
- Jiang, S. (2024a). Costly persuasion by a partially informed sender. *Working Paper*.
- Jiang, S. (2024b). Equilibrium refinement in finite action evidence games. *International Journal of Game Theory*, 53(1):43–70.
- Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.
- Lin, X. and Liu, C. (2024). Credible persuasion. *Journal of Political Economy*, 132(7):2228–2273.
- Lipnowski, E., Ravid, D., and Shishkin, D. (2022). Persuasion via weak institutions. *Journal of Political Economy*, 130(10):2705–2730.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2):380–391.
- Min, D. (2021). Bayesian persuasion under partial commitment. *Economic Theory*, 72(3):743–764.
- Morris, S. and Strack, P. (2019). The Wald problem and the relation of sequential sampling and ex-ante information costs. *Working Paper*.
- Nguyen, A. and Tan, T. Y. (2021). Bayesian persuasion with costly messages. *Journal of Economic*

Theory, 193:105212.

Shishkin, D. (2019). Evidence acquisition and voluntary disclosure. *Working Paper*.

Wald, A. (1945). Sequential method of sampling for deciding between two courses of action. *Journal of the American Statistical Association*, 40(231):277–306.