
ECONtribute
Discussion Paper No. 435

**Robust Rules for Voluntary Disclosure: When
Commitment Matters**

Sarah Auster

Helene Mass

October 2026

www.econtribute.de



**UNIVERSITÄT
ZU KÖLN**

Robust Rules for Voluntary Disclosure: When Commitment Matters*

Sarah Auster Helene Mass[†]

December 22, 2025

Abstract

We study the value of commitment in verifiable disclosure games when the receiver faces non-quantifiable uncertainty about the evidence environment. The sender has stochastic access to hard evidence and may disclose or suppress it; the receiver chooses a decision rule to maximize her payoff guarantee. When ambiguity concerns the payoff state, the receiver has strict incentives to soften her standards ex post after non-disclosure, making commitment strictly valuable. When ambiguity concerns the sender's expertise, incentives reverse: the receiver prefers a stricter reaction than is ex-ante optimal, so prespecified, hard-wired rules for non-disclosure improve payoffs for both sides.

1 Introduction

In many economic and organizational settings, informed parties possess verifiable information but retain discretion over whether to reveal it. Managers may choose not to release internal forecasts; auditors and rating agencies may refrain from emphasizing unfavorable evidence; regulators, courts, or boards often learn about misconduct or underperformance only when someone comes forward with hard proof. In all these cases, evidence is (at least partly) verifiable but disclosure is voluntary: it is difficult to lie but easy to remain silent.

A central policy question in such environments is how decision makers should react to non-disclosure. To what extent should investors, supervisors, or judges treat silence as a negative signal? When should the consequences of non-disclosure be left to case-by-case

*Sarah Auster gratefully acknowledges funding from the European Research Council (ERC)—grant No 101165999, the German Research Foundation (DFG) through Germany's Excellence Strategy—EXC 2126/1—390838866 and CRC TR 224 (Project B02).

[†]Auster: Department of Economics, University of Bonn, (email: sauster@uni-bonn.de);
Mass: Department of Economics, University of Vienna, (email: helene.mass@univie.ac.at)

discretion, and when should they be governed by fixed rules that are announced and committed to in advance? In practice, we frequently see the latter. For example, financial regulators specify how supervisory scrutiny escalates following a bank’s silence. Shareholders condition their support on the provision of corporate updates. Clinical guidelines prescribe default courses of action in response to non-reporting by laboratories. These institutions effectively encode simple rules of thumb—“no news is bad news”—into hard constraints on the receiver’s behavior.

From a theoretical perspective, the widespread reliance on prespecified rules for interpreting non-disclosure is difficult to reconcile with existing models of voluntary disclosure. A large literature on evidence games, starting from Grossman (1981); Milgrom (1981); Dye (1985); Jung and Kwon (1988), has clarified when and why informed agents voluntarily reveal even unfavorable information. Later work by Glazer and Rubinstein (2004, 2006), Sher (2011), and Hart et al. (2017) has established a remarkably strong benchmark: in a broad class of sender–receiver games with verifiable disclosure, giving the receiver commitment power over her message-contingent actions does *not* improve her payoff. In these models, the best equilibrium without commitment implements the same outcome as the optimal mechanism in which the receiver can commit in advance to a disclosure-contingent action rule. Taken at face value, these “no-value-of-commitment” results suggest that real-world institutions need not codify how non-disclosure is treated, highlighting a gap between theory and practice. We introduce ambiguity about the evidence process as one explanatory factor bridging this gap.

To study how ambiguity affects the value of commitment, we start from a simple Dye-style disclosure game with a one-dimensional payoff-relevant state $\theta \in [0, 1]$. A sender has probabilistic access to hard evidence proving the state: when evidence is available, he can either reveal θ or send an empty message. The receiver commits beforehand to a disclosure-contingent action rule, mapping each possible message into an action $a \in \mathbb{R}$. The sender then observes his evidence and, if evidence is available, chooses whether to disclose. In contrast to the classical model, the receiver does not have sufficient past data to confidently quantify the stochastic process governing the sender’s evidence. Formally, we capture this by letting her hold a *set of priors* over the sender’s evidence types and assuming that she seeks robustness, choosing the disclosure-contingent rule that maximizes her payoff guarantee.

To evaluate the receiver’s benefit from commitment, we assume that, conditional on the observed message, she updates each prior in the initial set by Bayes’ rule and then applies the maxmin criterion to the resulting posterior set (*Full Bayesian Updating*). It is well known that maxmin solutions are often dynamically inconsistent, which can make commitment valuable even in single-agent decision problems. Our goal here is

not to confirm that such dynamic inconsistencies can arise in the present setting, but to understand, in a natural and widely used disclosure environment, *when* commitment helps, *how* it helps (through which incentives and distortions), and *how much* it helps (in terms of payoff gains). To this end, we study two distinct sources of ambiguity and characterize the structure and magnitude of the commitment benefit in each case.

First, we consider ambiguity about the *payoff-relevant state*, assuming the receiver knows how likely the sender is to hold evidence but not the distribution of θ . Under the maxmin-optimal rule against all distributions of θ , the receiver’s action following non-disclosure hedges between two polar scenarios: one in which θ is low and the sender always withholds information, and one in which θ is high and non-disclosure occurs because the sender has nothing to reveal. Anticipating that silence makes the low state relatively more likely, the optimal rule prescribes a tough response to non-disclosure. Under some mild conditions, however, once non-disclosure occurs, the receiver has strict incentives to deviate by *raising* her action: ex post, her best hedge equalizes payoffs across the extreme states $\theta = 0$ and $\theta = 1$, effectively ignoring disclosure incentives. From the sender’s perspective, this ex-post temptation toward leniency is good news and weakens his incentives to disclose in the first place. Commitment lets the receiver credibly maintain a stricter non-disclosure response, thereby inducing more disclosure.

We next characterize how large the commitment benefit is and how it varies with the sender’s information. We show that the receiver’s benefit from commitment is strictly increasing in the sender’s likelihood of obtaining evidence. The higher this likelihood, the stricter the ex-ante optimal rule and the stronger the incentives to deviate from this rule when the sender provides no evidence. We also demonstrate that the qualitative findings do not hinge on the receiver optimizing against the worst case over all distributions of θ . Bridging the standard setting with ours, we consider the case where the receiver’s prior set consists of “small” perturbations around a *reference belief* $\hat{F}$ —the receiver’s best guess—capturing the case where she has a good understanding of the stochastic process governing θ but wants to hedge against the scenario where she is misspecified with some probability $\varepsilon > 0$. We show that for every $\varepsilon > 0$, the receiver finds it optimal ex post to raise her action following non-disclosure, so commitment remains strictly valuable.

Finally, we study the case where the receiver is confident about the distribution of the payoff state but perceives ambiguity about the *sender’s expertise*, i.e., the likelihood with which he obtains provable evidence. We show that ex ante, the worst-case scenario always corresponds to the lowest such likelihood. After observing non-disclosure, however, the receiver becomes concerned about scenarios in which silence is more likely to reflect strategic withholding of bad news rather than lack of evidence. We derive a simple necessary and sufficient condition under which this concern makes commitment strictly

valuable. When it holds, the receiver has strict incentives to revise her action following non-disclosure *downwards*. In this case, commitment benefits not only the receiver but also the sender, providing a strong rationale for the prevalence of prespecified disclosure standards and hard-wired responses to non-disclosure observed in practice.

Further related literature. Our paper is related to a growing literature that combines strategic communication with ambiguity aversion or robustness. On the “soft”-information side, several papers study cheap talk with ambiguity-sensitive players: e.g., Kellner and Le Quement (2018) allow the sender to use ambiguous randomization and demonstrate that equilibria featuring ambiguous communication can Pareto-dominate standard cheap-talk equilibria,¹ Li (2022) provides a related analysis for multiplier preferences, while Colo (2024) allows the sender to communicate distributions to a receiver facing ambiguity about the underlying state. A second strand considers Bayesian persuasion with ambiguity-averse agents: Beauchêne et al. (2019) and Hedlund et al. (2020) analyze persuasion when players have multiple-prior or α -maxmin preferences; Cheng (2025) and Pahlke (2022) focus explicitly on dynamic consistency in ambiguous persuasion problems. Our focus differs in that we keep the information structure simple and verifiable, as in classical disclosure games, and ask how ambiguity about the evidence environment affects the value of commitment to disclosure-contingent actions. Closer to our setting, Silva (2020) challenges the headline result that commitment is never valuable for the receiver in an expected-utility framework with stochastic verification that nests deterministic verifiability as a limit case. We complement his work by highlighting an alternative channel—ambiguity about the evidence process—through which commitment becomes valuable.

2 Environment

We consider a disclosure game between a sender (referred to as “he”) and a receiver (referred to as “she”). The sender is informed about the state of the world $\theta \in \Theta = [0, 1]$. After observing the evidence disclosed by the sender, the receiver chooses an action $a \in A = \mathbb{R}$. The sender aims to maximize the receiver’s action independently of the state. His utility is given by $u(a) = a$, while the receiver’s utility is described by the function $v(a, \theta)$.

Assumption 1. *a) The function v is twice continuously differentiable and jointly concave in (a, θ) .*

¹See also related work on ambiguous communication devices, e.g. Bose and Renou (2014).

b) For all $\theta, \theta' \in \Theta$, the optimal action $a^*(\theta) := \arg \max_a v(a, \theta)$ is strictly increasing in θ .

c) $v(a^*(0), 0) > v(a^*(0), 1)$ and $v(a^*(1), 0) < v(a^*(1), 1)$.

Parts a) and b) require concavity and single crossing. Part c) imposes that, for any given state, the highest achievable payoff is better than another state's worst payoff; this assumption ensures that ambiguity plays a non-trivial role in the receiver's optimization problem.

Evidence structure. The evidence structure follows Dye (1985): the sender can either perfectly verify the state or remain silent. Formally, types are described by the state and the sender's evidence endowment. The set of types is $T = \Theta \times \{0, 1\}$, where we interpret $t = (\theta, 1)$ as a sender who can prove that the state is θ , and $t = (\theta, 0)$ as one who cannot. The set of possible messages is $M \equiv \Theta \cup \{m_\emptyset\}$, where $m_\emptyset$ denotes the “no evidence” message. Letting $M_t \subseteq M$ be the set of feasible messages for type $t \in T$, we have, for all $\theta \in \Theta$,

$$M_{(\theta,0)} = \{m_\emptyset\} \quad \text{and} \quad M_{(\theta,1)} = \{\theta, m_\emptyset\}.$$

Thus, only those types endowed with verifiable evidence can disclose the true state, while all types can always remain silent.

Ambiguity. While, in the standard setting, the receiver has a unique prior over the set of types T , we are interested in situations where the receiver is uncertain about the evidence-generating process. To model this uncertainty, we describe the receiver's knowledge via a compact set of multiple priors $\mathcal{P} \subseteq \Delta(T)$ with the following product structure: each element of $\mathcal{P}$ is a pair (F, q) , where $F \in \Delta(\Theta)$ is a distribution over the payoff-relevant state θ and $q \in [0, 1]$ is the probability that the sender has verifiable evidence.

A central question in this paper is whether commitment is valuable to the receiver and, if so, what incentives she faces to deviate from an ex-ante optimal decision rule. To assess the value of commitment under ambiguity, we must specify the receiver's dynamic preferences. We assume that, conditional on observing any message, the receiver updates each prior in the initial set $\mathcal{P}$ via Bayes' rule and then evaluates actions according to the maxmin criterion over the resulting posterior set (*Full Bayesian Updating*).

Timing. The timing of the game is as follows. First, the receiver publicly commits to a schedule of message-contingent actions, that is, a mapping from messages to actions. Second, the sender's type $t = (\theta, e)$ is realized, and he sends a message from his feasible

message set M_t ; in particular, if $e = 1$, he chooses whether to disclose the state θ or to send the no-evidence message $m_\emptyset$. Finally, the receiver observes the message and chooses the corresponding action according to her schedule.

Optimization problem. We can represent the receiver's problem as choosing both a reporting strategy for the sender and a mapping from messages to actions. Formally, the receiver chooses a reporting rule $m : T \rightarrow M$ that respects feasibility,

$$m(\theta, 0) = m_\emptyset \quad \text{and} \quad m(\theta, 1) \in \{\theta, m_\emptyset\} \quad \text{for all } \theta \in \Theta,$$

and an action rule $a : M \rightarrow \mathbb{R}$ to solve

$$\max_{(m,a)} \min_{\mu \in \mathcal{P}} \sum_{e \in \{0,1\}} \int_{\theta \in \Theta} v(a(m(\theta, e)), \theta) d\mu(\theta, e) \quad (1)$$

subject to the sender's incentive compatibility constraint

$$a(m(\theta, 1)) \geq a(m_\emptyset) \quad \text{for all } \theta \in \Theta. \quad (2)$$

The sender's payoff is increasing in the receiver's action, so (2) ensures that any type who is able to disclose θ (i.e., with $e = 1$) is willing to follow the intended reporting rule rather than remaining silent.

Before turning to the main analysis, it is useful to establish the following preliminary lemma: it shows that the receiver's optimal decision rule follows a simple cutoff structure, exactly like in the standard Dye setting without ambiguity. To state the result, let $\theta^*(x)$ be the state at which x maximizes $v(\cdot, \theta)$ so that $a^*(\theta^*(x)) = x$.²

Lemma 1. *For any $\mathcal{P} \subseteq \Delta(T)$, a maximizer (m, a) of problem (1) satisfies*

$$a(m(t)) = \begin{cases} a^*(\theta) & \text{if } t = (\theta, 1) \text{ and } \theta > \theta^*(x) \\ x & \text{otherwise} \end{cases}$$

for some $x \in A$.

Given this result, the receiver's optimization problem can be reduced to the choice of a scalar $x \in A$, specifying the action taken upon receiving the no-evidence message. Her expected payoff as a function of x and a distribution-probability pair (F, q) is then given

²All proofs are relegated to the Appendix.

by

$$V(x; F, q) := q \left(\int_0^{\theta^*(x)} v(x, \theta) dF(\theta) + \int_{\theta^*(x)}^1 v(a^*(\theta), \theta) dF(\theta) \right) + (1 - q) \int_0^1 v(x, \theta) dF(\theta). \quad (3)$$

The receiver's optimization problem amounts to

$$\max_{x \in A} \min_{(F, q) \in \mathcal{P}} V(x; F, q). \quad (4)$$

3 Ambiguity about the Payoff-Relevant State

In this section, we consider the case where the receiver knows the probability $q \in (0, 1)$ with which the sender has verifiable evidence but is uncertain about how the underlying payoff state θ is distributed. We start with the case where the receiver optimizes her payoff guarantee over all possible distributions of θ . Fixing $q \in (0, 1)$,

$$\mathcal{P} = \Delta(\Theta) \times \{q\}. \quad (5)$$

Proposition 1. *Let Assumption 1 be satisfied and let $\mathcal{P}$ be specified by (5). The unique solution x^* to problem (4) is:*

1. *If*

$$[v(a^*(0), 0) - v(a^*(0), 1)] - q[v(a^*(1), 1) - v(a^*(0), 1)] \leq 0, \quad (6)$$

then $x^ = a^*(0)$.*

2. *Otherwise, $x^* \in (a^*(0), a^*(1))$ is uniquely characterized by*

$$v(x^*, 0) = q v(a^*(1), 1) + (1 - q) v(x^*, 1). \quad (7)$$

As usual, we can interpret (4) as the value of a zero-sum game between the receiver and an adversarial Nature, where, given (5), Nature chooses a distribution $F \in \Delta(\Theta)$ to minimize the receiver's payoff. Solving problem (4) amounts to finding a saddle point of this game.

In the proof of Proposition 1, we first observe that, for any given x and any distribution F , Nature prefers a mean-preserving spread of F . This is a consequence of Assumption 1(a) and implies that Nature optimally distributes all probability weight on the extreme

states 0 and 1, thereby making it most difficult for the receiver to tailor her action to the actual state. The receiver's payoff in the scenario $\theta = 0$ is determined by the sender's non-disclosure and equals $v(x, 0)$, while her payoff in the scenario $\theta = 1$ is an average of the payoffs when the sender has evidence and when he does not.

There are two candidate saddle points. If (6) is satisfied, the receiver's worst-case scenario is $\theta = 0$ no matter which value of x she commits to. The payoff guarantee is thus maximized at $x^* = a^*(0)$. If, instead, condition (6) is violated, the worst-case scenario depends on the receiver's choice of x . Problem (4) is then solved by the value of x that equalizes the receiver's expected payoff across the two extreme scenarios $\theta = 0$ and $\theta = 1$, as captured by (7). Condition (6) reveals that this latter case arises whenever q is sufficiently small. Intuitively, if q is small (the sender is rarely informed), the receiver's worst-case concern is relatively balanced across states 0 and 1, and it is optimal to choose an interior x^* that trades off losses in both states. If q is large and $v(a^*(1), 1) > v(a^*(0), 0)$, the receiver's potential loss in state 1 is limited, as the sender discloses with high probability; in this case, the worst-case payoff is determined by state 0 independently of x , and the receiver optimally chooses the corner solution $x^* = a^*(0)$.

The value of commitment. Suppose the receiver cannot commit to an action conditional on the sender's non-disclosure. Since, for each θ , there is a positive probability $1 - q$ that the sender has no evidence to show, observing the empty message does not allow her to rule out any state in Θ . The receiver's updated belief set following the empty message is thus equal to her initial belief set $\mathcal{P}$ and, in particular, includes all scenarios $F \in \Delta(\Theta)$. The receiver's choice problem following $m_\emptyset$ is then

$$\max_{x \in A} \min_{F \in \Delta(\Theta)} \int_0^1 v(x, \theta) dF(\theta). \quad (8)$$

Given Assumption 1, this problem has a unique solution $\tilde{x}$, characterized by

$$v(\tilde{x}, 0) = v(\tilde{x}, 1).$$

As before, the receiver's worst-case distribution places all mass on whichever of the two extreme states 0 or 1 yields the lower payoff; at the optimum $\tilde{x}$, these two payoffs are equalized.

Comparing x^* and $\tilde{x}$ reveals a dynamic inconsistency. Ex ante, when designing the commitment rule, the receiver optimally chooses x^* taking into account how the sender will disclose evidence. Ex post, after observing the empty message, she strictly prefers to switch to $\tilde{x} > x^*$, the robust action that equalizes her payoff across states 0 and 1. The

caution that leads the receiver to commit to a tough non-disclosure rule ex ante is thus undermined once she actually faces non-disclosure and starts worrying that the sender might simply have had nothing to report. From the sender's perspective, this ex-post deviation makes the receiver more lenient after receiving no evidence, which weakens his incentives to disclose when he has evidence to show. Commitment therefore calls for a tight rule that encourages more disclosure through a harsher response to non-disclosure.

Proposition 2. *Let Assumption 1 be satisfied and let $\mathcal{P}$ be specified by (5). For every $q \in (0, 1)$, the gain from commitment, defined as*

$$\Delta(q) := \min_{F \in \Delta(\Theta)} V(x^*; F, q) - \min_{F \in \Delta(\Theta)} V(\tilde{x}; F, q),$$

is strictly positive and non-decreasing in q . Moreover, $\tilde{x} > x^$.*

If the commitment problem has a corner solution, $x^* = a^*(0)$, the gain from commitment is strictly positive but constant in q : under the receiver's ex-ante worst-case distribution, which concentrates on $\theta = 0$, disclosure is irrelevant and neither x^* nor $\tilde{x}$ depends on q . By contrast, if the commitment solution is interior, the receiver benefits from anticipating that the sender will disclose in state 1 whenever evidence is available. The larger this probability, the smaller the receiver's concern for scenario $\theta = 1$ and the smaller x^* . Consequently, when the solution is interior, the gain from commitment is strictly increasing in q .

For an illustration of the magnitude of the commitment gain, consider the case of quadratic loss, $v(a, \theta) = -(a - \theta)^2$, where $\tilde{x} = \frac{1}{2}$ and

$$x^* = \frac{\sqrt{1-q}}{1 + \sqrt{1-q}}.$$

Then the commitment gain ranges from 0 as $q \rightarrow 0$ to $1/4$ as $q \rightarrow 1$. In the latter case, the loss from not being able to commit is as large as the loss from not having access to the sender.

3.1 Moderate ambiguity

A setting in which the receiver treats any distribution over Θ as possible represents an extreme case of ambiguity. We now extend the main insights from the previous analysis to the case where the receiver has a reference belief $\hat{F}$ but recognizes that this belief may

be misspecified.³ Her decision criterion is

$$\max_{x \in A} \left[(1 - \varepsilon)V(x; \hat{F}, q) + \varepsilon \min_{F \in \Delta(\Theta)} V(x; F, q) \right] \quad (9)$$

for some exogenously given $\varepsilon \in (0, 1]$. This is equivalent to the receiver following a maxmin decision criterion where the set of priors is

$$\mathcal{P} = \left\{ (1 - \varepsilon)\hat{F} + \varepsilon F \mid F \in \Delta(\Theta) \right\} \times \{q\}. \quad (10)$$

Let x_ε^* be the optimal policy when the ambiguity level is ε and let $\hat{x} := \arg \max_x V(x; \hat{F}, q)$ be the optimal policy without ambiguity. To avoid case distinctions, we restrict attention to situations where the problem with extreme ambiguity has an interior solution (i.e., (6) is violated) and abstract from the non-generic case $x^* = \hat{x}$.

Proposition 3. *Let Assumption 1 be satisfied and let $\mathcal{P}$ be specified by (10). Assume that (6) is violated. Then there exists some $\bar{\varepsilon} \in (0, 1)$ such that the unique solution x_ε^* to problem (9) satisfies:*

1. For all $\varepsilon \geq \bar{\varepsilon}$, $x_\varepsilon^* = x^*$.
2. For all $\varepsilon < \bar{\varepsilon}$,
 - (a) if $\hat{x} < x^*$, then $x_\varepsilon^* \in (\hat{x}, x^*)$;
 - (b) if $\hat{x} > x^*$, then $x_\varepsilon^* \in (x^*, \hat{x})$.

If the receiver's confidence in the reference belief is sufficiently low (that is, if ε is sufficiently high), then the solution under moderate ambiguity coincides with the Knightian solution. To see why, note that Nature's choice does not depend on the magnitude of ε . Hence, whenever x_ε^* differs from the Knightian solution x^* , Nature optimally assigns the entire mass ε either to $\theta = 0$ (when $x_\varepsilon^* > x^*$) or to $\theta = 1$ (when $x_\varepsilon^* < x^*$). For ε sufficiently large, this would induce cycling of best responses. Therefore, at the saddle point, Nature's indifference condition must still be satisfied. By contrast, when ε is sufficiently small, the receiver is willing to move closer to the optimal action prescribed by the reference belief $\hat{F}$, even though, with probability ε , this action will be punished.

The value of commitment. Given the commitment policy x_ε^* , the sender discloses whenever $\theta > \theta^*(x_\varepsilon^*)$. Conditional on observing the empty message, the receiver's payoff

³The proofs related to this extension are relegated to Online Appendix B

maximizing action solves

$$\max_{x \in A} \min_{F \in \Delta(\Theta)} \left\{ \varepsilon \int_0^1 v(x, \theta) dF(\theta) + (1 - \varepsilon) \mathbb{E}_{\hat{F}}[v(x, \theta) \mid m = \emptyset] \right\}, \quad (11)$$

where

$$\mathbb{E}_{\hat{F}}[v(x, \theta) \mid m = \emptyset] = \frac{q \int_0^{\theta^*(x)} v(x, \theta) d\hat{F}(\theta) + (1 - q) \int_0^1 v(x, \theta) d\hat{F}(\theta)}{q\hat{F}(\theta^*(x)) + (1 - q)}.$$

The following result shows that commitment is valuable for every level of ambiguity, i.e., even when ε is arbitrarily small. In other words, the classical result that in Dye-type disclosure games the receiver cannot improve her payoff by committing ex ante hinges critically on her having full confidence in her probabilistic assessment of the underlying state (the special case $\varepsilon = 0$).

Proposition 4. *Let Assumption 1 be satisfied and let $\mathcal{P}$ be specified by (10). Assume that (6) is violated. For every $q \in (0, 1)$ and every $\varepsilon \in (0, 1)$, the solution to (11) is strictly greater than x_ε^* .*

4 Ambiguity about the Sender's Expertise

Suppose now the receiver knows the distribution of payoff states F but faces ambiguity about the sender's expertise, captured by the probability q with which the sender obtains evidence. Let F be differentiable on support Θ and let $[\underline{q}, \bar{q}]$ denote the set of all q the receiver considers possible. Under this restriction, the receiver solves optimization problem (4) for the set

$$\mathcal{P} = \{F\} \times [\underline{q}, \bar{q}]. \quad (12)$$

Proposition 5. *Let Assumption 1 be satisfied and let $\mathcal{P}$ be specified by (12). The unique solution x^* to problem (4) is characterized by*

$$\underline{q} \int_0^{\theta^*(x^*)} v_a(x^*, \theta) dF(\theta) + (1 - \underline{q}) \int_0^1 v_a(x^*, \theta) dF(\theta) = 0.$$

Proposition 5 shows that the receiver's maxmin problem is solved by the action that is optimal when the sender's likelihood of obtaining evidence is known to be $\underline{q}$. The proof

is straightforward. For every x ,

$$\int_0^{\theta^*(x)} v(x, \theta) dF(\theta) + \int_{\theta^*(x)}^1 v(a^*(\theta), \theta) dF(\theta) > \int_{\theta^*(x)}^1 v(x, \theta) dF(\theta),$$

so the coefficient on q is strictly larger than the coefficient on $(1-q)$. Hence, the worst-case scenario is always $q = \underline{q}$, independently of the receiver's choice of x , and the maxmin optimal strategy is obtained by solving the problem in which the sender is least likely to hold any evidence. Economically, the ambiguity-averse receiver behaves as if she were facing the weakest plausible expert: she safeguards against the possibility that evidence is rarely available with a relatively generous response (high x) to non-disclosure.

The value of commitment. The commitment solution asks the receiver to act as if the sender's expertise were given by $\underline{q}$. We now inquire whether this continues to be optimal once the receiver observes the empty message and updates her beliefs about the state and the sender's information accordingly.

Without commitment, upon observing $m = \emptyset$, the receiver takes as given that the sender discloses whenever he can and the state lies above the disclosure cutoff implied by x^* . She then solves

$$\max_{x \in \mathbb{R}} \min_{q \in [\underline{q}, \bar{q}]} \frac{1}{1 - (1 - F(\theta^*(x^*)))q} \left[\int_0^{\theta^*(x^*)} v(x, \theta) dF(\theta) + \int_{\theta^*(x^*)}^1 v(x, \theta)(1 - q) dF(\theta) \right],$$

and we denote the solution by $\tilde{x}$.

The next result characterizes when the optimal commitment plan is dynamically inconsistent and, consequently, when the gain from commitment is strictly positive.

Proposition 6. *Let Assumption 1 be satisfied and let $\mathcal{P}$ be specified by (12). The gain from commitment, defined by*

$$\min_{q \in [\underline{q}, \bar{q}]} V(x^*; F, q) - \min_{q \in [\underline{q}, \bar{q}]} V(\tilde{x}; F, q),$$

is strictly positive if and only if

$$\mathbb{E}_F[v(x^*, \theta)] < \mathbb{E}_F[v(x^*, \theta) \mid \theta \geq \theta^*(x^*)]. \quad (13)$$

In that case $\tilde{x} < x^$.*

Whenever (13) holds, not only the receiver but also the sender benefits from the receiver having commitment power: since $\tilde{x} < x^*$, the sender obtains a (weakly) higher

action under commitment, while the receiver achieves a strictly higher worst-case payoff. This provides a natural rationale for prespecified disclosure rules in practice—such as fixed reporting standards or hard-wired communication protocols—that are agreed upon *ex ante* and then rigidly followed. It also contrasts with the case of ambiguity about the state distribution, where the receiver’s optimal commitment rule called for a lower action following non-disclosure compared to the conditionally optimal response.

The receiver’s incentives to deviate from the optimal commitment rule in the current setting can be understood as follows. Although the receiver’s *ex-ante* payoff is always higher when the sender obtains evidence with higher probability, it is not immediately clear how her expected payoff varies with q conditional on the event that the sender discloses no evidence. Upon observing $m = \emptyset$, a higher q implies a higher probability that non-disclosure arises because the state lies in the non-disclosure region $[0, \theta^*(x^*)]$, rather than because the sender has no evidence. If the receiver’s expected utility from x^* in states $\theta \in [0, \theta^*(x^*)]$ is sufficiently low, her conditional payoff may actually be *lower* at higher values of q . Upon receiving the empty message, the receiver’s worst-case scenario may thus shift from $\underline{q}$ —where non-disclosure is mostly driven by a lack of evidence—to $\bar{q}$ —where non-disclosure is more likely to reflect strategic withholding of unfavorable evidence. Condition (13) captures precisely this situation: the receiver’s conditional payoff at x^* is minimized at $q = \bar{q}$, making her strictly prefer to deviate downward from the *ex-ante* optimal rule whenever she cannot commit.

To gain further intuition, consider the quadratic loss case $v(a, \theta) = -(a - \theta)^2$ and $\underline{q} = 0$. The *ex-ante* optimal value of x is then $x^* = E_F[\theta] =: \mu$, and condition (13) amounts to

$$\mathbb{E}_F[(\theta - \mu)^2] > \mathbb{E}_F[(\theta - \mu)^2 \mid \theta \geq \mu] \iff \text{Cov}_F((\theta - \mu)^2, \mathbb{1}_{\{\theta \geq \mu\}}) < 0.$$

Starting from the commitment solution, the receiver has strict incentives to lower her response to non-disclosure whenever the unconditional variance is strictly greater than the conditional variance for states above the mean. Intuitively, relative to a symmetric benchmark such as the uniform distribution, F must be sufficiently skewed to the left for commitment to be valuable.⁴

The benefit from commitment depends not only on the underlying distribution of the payoff state but also on the set $[\underline{q}, \bar{q}]$. As $\underline{q}$ increases, the commitment solution x^* , characterized in Proposition 5, shifts downward, since a higher baseline expertise (in the worst case) leads the receiver to put less weight on high states when choosing x^* . This makes it harder for (13) to be satisfied: as x^* decreases, the conditional

⁴Under any distribution symmetric around the mean, it holds that $\mathbb{E}_F[(\theta - \mu)^2] = \mathbb{E}_F[(\theta - \mu)^2 \mid \theta \geq \mu]$.

expected utility in the high states falls relative to the unconditional expected utility. Hence, even when F is sufficiently left-skewed to generate a commitment gain for small $\underline{q}$, the gain from commitment gradually vanishes as $\underline{q}$ becomes large. In contrast to the previous case, where a higher q implies a larger commitment gain, greater expertise in the current setting—captured by an upward shift of $[\underline{q}, \bar{q}]$ —lowers the receiver’s benefit from commitment.

5 Conclusion

This paper revisits a canonical verifiable disclosure game through the lens of ambiguity about the evidence environment. In the standard formulation of the problem, commitment to disclosure-contingent actions has no value for the receiver. We show that once the receiver is uncertain about how evidence is generated, commitment often becomes valuable, bringing the theory in line with many real-world institutions that hard-wire responses to non-disclosure.

When ambiguity concerns the underlying fundamentals, the receiver’s ex-ante rule disciplines the sender by prescribing a stringent response to non-disclosure. Commitment prevents her from softening this standard ex post, even though doing so would be conditionally optimal once her concern shifts toward the scenario in which the sender truly has nothing to show. By contrast, when ambiguity concerns the sender’s expertise, giving the sender the benefit of the doubt hedges against the possibility that evidence is genuinely hard to obtain. Ex post, however, the receiver is tempted to react too harshly to non-disclosure, so commitment serves to restrain excessive ex-post punishment.

A Proofs

A.1 Preliminaries

The following observation will be used repeatedly in the proofs.

Lemma 2. *Under Assumption 1, for every $F \in \Delta(\Theta)$ and every $q \in [0, 1]$, the function*

$$x \mapsto V(x; F, q)$$

defined in (3) is concave on A .

Proof. Fix F and q . We can rewrite V as

$$\begin{aligned} V(x; F, q) &= \int_0^1 \left[(1 - q) v(x, \theta) + q \mathbf{1}_{\{a^*(\theta) \leq x\}} v(x, \theta) + q \mathbf{1}_{\{a^*(\theta) > x\}} v(a^*(\theta), \theta) \right] dF(\theta) \\ &= \int_0^1 \varphi(x, \theta) dF(\theta), \end{aligned}$$

where, for each θ ,

$$\varphi(x, \theta) := \begin{cases} (1 - q) v(x, \theta) + q v(a^*(\theta), \theta), & \text{if } x < a^*(\theta), \\ v(x, \theta), & \text{if } x \geq a^*(\theta). \end{cases}$$

On the region $x < a^*(\theta)$, $\varphi(\cdot, \theta)$ is an affine combination of $v(\cdot, \theta)$ and a constant, hence concave in x . On the region $x \geq a^*(\theta)$, $\varphi(\cdot, \theta) = v(\cdot, \theta)$ is again concave in x . We next check the behavior at the kink $x = a^*(\theta)$. The left and right derivatives at $x = a^*(\theta)$ are

$$\partial_x \varphi(x, \theta) \Big|_{x \uparrow a^*(\theta)} = (1 - q) v_a(a^*(\theta), \theta) \quad \text{and} \quad \partial_x \varphi(x, \theta) \Big|_{x \downarrow a^*(\theta)} = v_a(a^*(\theta), \theta),$$

which both equal zero because $a^*(\theta)$ maximizes $v(\cdot, \theta)$. This, together with piecewise concavity, implies $\varphi(\cdot, \theta)$ is concave on A for each θ .

Finally, $V(\cdot; F, q)$ is an integral of concave functions $\varphi(\cdot, \theta)$ with respect to F , so concavity is preserved. $\square$

A.2 Proof of Lemma 1

Proof. Fix an arbitrary feasible rule (m, a) and let $x := a(m_\emptyset)$ be the action taken when the no-evidence message $m_\emptyset$ is sent. Since every type can send $m_\emptyset$, incentive compatibility (2) implies $a(m') \geq x$, for all $m' \in m(T)$.

Consider an alternative rule, $(\tilde{m}, \tilde{a})$ such that $\tilde{a}(\tilde{m}(t)) = a^*(\theta)$ if $t = (\theta, 1)$ with $\theta > \theta^*(x)$ and $\tilde{a}(\tilde{m}(t)) = x$ otherwise. We will show that this rule constitutes a (weak) improvement relative to the arbitrary feasible rule. Since $\tilde{a}(\tilde{m}(t)) \geq x = \tilde{a}(m_\emptyset)$ for all t , this rule is also incentive compatible. Comparing the receiver's payoff type by type, note that for all $t = (\theta, 1)$ with $\theta > \theta^*(x)$,

$$\tilde{a}(\tilde{m}((\theta, 1))) = a^*(\theta) \geq a(m((\theta, 1))),$$

for all $t = (\theta, 1)$ with $\theta \leq \theta^*(x)$,

$$a^*(\theta) \leq \tilde{a}(\tilde{m}((\theta, 1))) = x \leq a(m((\theta, 1))),$$

and for all $t = (\theta, 0)$, $\tilde{a}(\tilde{m}(t)) = a(m(t))$. Assumption 1 (a) then implies that moving from (a, m) to $(\tilde{a}, \tilde{m})$ constitutes a pointwise improvement for the receiver. $\square$

A.3 Proof of Proposition 1

Proof. To show that x^* solves problem (4), it suffices to show that x^* , together with some $F^* \in \Delta(\Theta)$, forms a saddle point of the associated zero-sum game. The candidate saddle-point distribution has support $\{0, 1\}$ with probability mass

$$F^*(0) = \frac{(1 - q)v_a(x^*, 1)}{(1 - q)v_a(x^*, 1) - v_a(x^*, 0)}$$

assigned to 0 and $1 - F^*(0)$ assigned to 1. Note that $F^*(0) = 1$ if $x^* = a^*(0)$ (case 1) since $v_a(a^*(0), 0) = 0$, and $F^*(0) \in (0, 1)$ if $x^* \in (a^*(0), a^*(1))$ (case 2) since $v_a(x^*, 1) > 0$ and $v_a(x^*, 0) < 0$.

To show that (x^*, F^*) forms a saddle point, we need to verify that for all $x \in A$ and all $F \in \Delta(\Theta)$,

$$V(x; F^*, q) \leq V(x^*; F^*, q) \leq V(x^*; F, q). \quad (14)$$

1. For the first inequality, fix Nature's choice to F^* and consider the receiver's optimization problem

$$\max_{x \in A} F^*(0)v(x, 0) + (1 - F^*(0))(qv(a^*(1), 1) + (1 - q)v(x, 1)).$$

The first-order condition for this problem is

$$F^*(0)v_a(x, 0) + (1 - F^*(0))(1 - q)v_a(x, 1) = 0.$$

Given our specification of F^* , this condition is satisfied at $x = x^*$. As we have shown in Section A.1, the first-order condition is not only necessary but also sufficient for x^* to maximize $V(x; F^*, q)$.

2. For the second inequality, we want to verify that, given x^* , Nature cannot lower the receiver's payoff by choosing a distribution other than F^* . We first show that Nature cannot benefit from distributions that assign positive mass to realizations of θ other than 0 or 1.

To see this, define the receiver's payoff at x^* as a function of θ :

$$\begin{cases} v(x^*, \theta) & \text{if } \theta < \theta^*(x^*), \\ qv(a^*(\theta), \theta) + (1 - q)v(x^*, \theta) & \text{if } \theta \geq \theta^*(x^*). \end{cases}$$

This piecewise function is concave in θ : by Assumption 1(a), for each fixed x , the function $\theta \mapsto v(x, \theta)$ is concave. Moreover, since $v(a, \theta)$ is jointly concave in (a, θ) and A is convex, the value function

$$\theta \mapsto v(a^*(\theta), \theta) = \max_{a \in A} v(a, \theta)$$

is obtained by partially maximizing a concave function over a and is therefore concave as well (see, for instance, Hiriart-Urruty and Lemaréchal (2001), Chapter B.2). Therefore, for any F that is a mean-preserving spread of G , we have $V(x^*; F, q) \leq V(x^*; G, q)$. Since any distribution G with mass on $(0, 1)$ is a mean-preserving contraction of some distribution F with support $\{0, 1\}$, we may thus restrict attention to distributions F with support $\{0, 1\}$.

For such an F , write the receiver's payoff as

$$V(x^*; F, q) = F(0)v(x^*, 0) + (1 - F(0))(qv(a^*(1), 1) + (1 - q)v(x^*, 1)),$$

with

$$\frac{\partial V}{\partial F(0)} = v(x^*, 0) - qv(a^*(1), 1) - (1 - q)v(x^*, 1).$$

Define

$$g(x) := v(x, 0) - qv(a^*(1), 1) - (1 - q)v(x, 1).$$

Then

$$g'(x) = v_a(x, 0) - (1 - q)v_a(x, 1) < 0 \quad \text{for all } x \in [a^*(0), a^*(1)], \quad q \in (0, 1),$$

since $a^*(0)$ maximizes $v(\cdot, 0)$ and $a^*(1)$ maximizes $v(\cdot, 1)$, so $v_a(x, 0) \leq 0$ and $v_a(x, 1) \geq 0$, with at least one of these inequalities holding strictly. Moreover,

$$g(a^*(1)) = v(a^*(1), 0) - v(a^*(1), 1) < 0$$

by Assumption 1(c).

We distinguish two cases: If $g(a^*(0)) \leq 0$, then $x^* = a^*(0)$ (case 1). For $x = a^*(0)$, the receiver's value $V(a^*(0); F, q)$ is (weakly) decreasing in $F(0)$, so $F(0) = 1$ satisfies Nature's optimality condition, and $F^* = \delta_0$ is a worst-case distribution. If $g(a^*(0)) > 0$, then by continuity and strict monotonicity of g , there exists a unique interior $x^* \in (a^*(0), a^*(1))$ such that $g(x^*) = 0$, which is precisely condition (7) (case 2). For this x^* , the receiver's value $V(x^*; F, q)$ is constant in $F(0)$, so Nature is indifferent among all distributions with support $\{0, 1\}$.

Taken together, x^* is the unique solution to problem (4). □

A.4 Proof of Proposition 2

Proof. We can follow the same steps as in the proof of Proposition 1 to show that $(\tilde{x}, \tilde{F})$ forms a saddle point for problem (8), where $\tilde{x}$ solves $v(\tilde{x}, 0) = v(\tilde{x}, 1)$ and $\tilde{F}$ is a distribution on $\{0, 1\}$ with

$$\tilde{F}(0) = \frac{v_a(\tilde{x}, 1)}{v_a(\tilde{x}, 1) - v_a(\tilde{x}, 0)}.$$

By Assumption 1(c), there always exists an interior solution $\tilde{x} \in (a^*(0), a^*(1))$.

To show $\tilde{x} > x^*$, recall that $\tilde{x}$ is such that $v(\tilde{x}, 0) = v(\tilde{x}, 1)$, whereas x^* is either equal to $a^*(0) < \tilde{x}$ or is an interior solution satisfying (7). Consider $g(\tilde{x})$, where g was defined in the proof of Proposition 1. Then

$$g(\tilde{x}) = \underbrace{v(\tilde{x}, 0) - v(\tilde{x}, 1)}_{=0} - q(v(a^*(1), 1) - v(\tilde{x}, 1)) < 0.$$

As shown in the proof of Proposition 1, g is strictly decreasing on $(a^*(0), a^*(1))$ with $g(x^*) = 0$, so $\tilde{x} > x^*$.

This strict inequality implies $\Delta(q) > 0$ for all $q \in (0, 1)$; it remains to show that $\Delta(q)$ is non-decreasing in q .

1. If $x^* = a^*(0)$ (corner solution), then neither x^* nor $\tilde{x}$ depends on q , so the gain from commitment is constant in q .

2. If x^* is an interior solution, it is defined as the action that makes Nature indifferent between a degenerate distribution centered at $\theta = 0$ and a degenerate distribution centered at $\theta = 1$, so

$$\min_{F \in \Delta(\Theta)} V(x^*; F, q) = v(x^*, 0).$$

Since $\tilde{x} > x^*$ and $g(x)$ is strictly decreasing in x , the distribution assigning all mass to 0 becomes the unique worst-case distribution for $\tilde{x}$, implying

$$\min_{F \in \Delta(\Theta)} V(\tilde{x}; F, q) = v(\tilde{x}, 0).$$

Moreover, $\tilde{x}$ does not depend on q , so

$$\frac{d\Delta(q)}{dq} = \frac{d}{dq} [v(x^*, 0) - v(\tilde{x}, 0)] = v_a(x^*, 0) \frac{dx^*}{dq}.$$

To sign $\frac{dx^*}{dq}$, note that x^* satisfies

$$g(x^*) = v(x^*, 0) - qv(a^*(1), 1) - (1 - q)v(x^*, 1) = 0.$$

Differentiating implicitly with respect to q ,

$$\frac{dx^*}{dq} = \frac{\overbrace{v(a^*(1), 1) - v(x^*, 1)}^{>0}}{\underbrace{v_a(x^*, 0) - (1 - q)v_a(x^*, 1)}_{=g'(x^*) < 0}} < 0.$$

Together with $v_a(x^*, 0) < 0$, this yields

$$\frac{d\Delta(q)}{dq} = v_a(x^*, 0) \frac{dx^*}{dq} > 0.$$

□

A.5 Proof of Proposition 6

Proof. It is straightforward that at $q = \underline{q}$, the receiver's payoff conditional on non-disclosure

$$W(x; F, q) = \frac{1}{1 - (1 - F(\theta^*(x^*)))q} \left(\int_0^1 v(x, \theta) dF(\theta) - q \int_{\theta^*(x^*)}^1 v(x, \theta) dF(\theta) \right)$$

is maximized by x^* . This reflects the fact that, without ambiguity about q , commitment has no value.

For $(x^*, \underline{q})$ to form a saddle point at the interim stage, it must hold that at $x = x^*$, $W(x; F, q)$ is minimized at $q = \underline{q}$. Letting $S^* := 1 - F(\theta^*(x^*))$ be the probability that θ lies in the disclosure region and taking the derivative with respect to q yields

$$\frac{\partial W(x^*; F, q)}{\partial q} = \frac{S^*}{(1 - S^*q)^2} \left(\int_0^1 v(x^*, \theta) dF(\theta) - \frac{1}{S^*} \int_{\theta^*(x^*)}^1 v(x^*, \theta) dF(\theta) \right). \quad (15)$$

$W(x^*; F, q)$ is minimized at $q = \underline{q}$ if and only if this expression is non-positive, i.e., if and only if (13) holds.

It remains to show that if (13) fails, the receiver's interim problem has a saddle point $(\tilde{x}, \tilde{q})$ such that $\tilde{x} < x^*$. The saddle-point condition is

$$W(x; F, \tilde{q}) \geq W(\tilde{x}; F, \tilde{q}) \geq W(\tilde{x}; F, q), \quad \forall x \in \mathbb{R}, \forall q \in [\underline{q}, \bar{q}]. \quad (16)$$

Considering

$$W_x(x; F, q) = \frac{1}{1 - S^*q} \left[\int_0^1 v_a(x, \theta) dF(\theta) - q \int_{\theta^*(x^*)}^1 v_a(x, \theta) dF(\theta) \right],$$

let $\bar{x}$ be such that $W_x(\bar{x}; F, \bar{q}) = 0$. Noting that the term in square brackets is strictly decreasing in q and that $W_x(x^*; F, \underline{q}) = 0$, we obtain $\bar{x} < x^*$.

We now distinguish two cases. Recall from (15) that Nature's incentives to increase or decrease q given the receiver's choice x are determined by the sign of

$$\Phi(x) := \int_0^1 v(x, \theta) dF(\theta) - \frac{1}{S^*} \int_{\theta^*(x^*)}^1 v(x, \theta) dF(\theta). \quad (17)$$

If $\Phi(\bar{x}) \leq 0$ so that $W(x; F, q)$ is weakly decreasing in q , we set $(\tilde{x}, \tilde{q}) = (\bar{x}, \bar{q})$. Since $\bar{x}$ solves $W_x(x; F, \bar{q}) = 0$ and $W(x; F, \bar{q})$ is concave in x , the first inequality of (16) holds. Since $\Phi(\bar{x}) \leq 0$, $W(\bar{x}; F, q)$ is minimized at $q = \bar{q}$, so the second inequality of (16) holds as well. The profile $(\bar{x}, \bar{q})$ thus forms a saddle point.

Assume next that $\Phi(\bar{x}) > 0$ and recall that $\Phi(x^*) < 0$ by the assumption that (13) fails. By continuity, there exists $\tilde{x} \in (\bar{x}, x^*)$ such that $\Phi(\tilde{x}) = 0$. Given $\tilde{x}$, Nature is indifferent between all $q \in [\underline{q}, \bar{q}]$, so the second inequality of (16) holds. Noting that $W_x(\tilde{x}; F, \underline{q}) > 0$ and $W_x(\tilde{x}; F, \bar{q}) < 0$, we can choose $\tilde{q} \in (\underline{q}, \bar{q})$ such that $W_x(\tilde{x}; F, \tilde{q}) = 0$, thereby satisfying the first inequality of (16).

In both cases we obtain a saddle point $(\tilde{x}, \tilde{q})$ with $\tilde{x} < x^*$ whenever (13) fails. $\square$

References

- Beauchêne, D., J. Li, and L. Li (2019). Ambiguous persuasion. *Journal of Economic Theory* 181, 80–106.
- Bose, S. and L. Renou (2014). Mechanism design with ambiguous communication devices. *Econometrica* 82(5), 1853–1872.
- Cheng, X. (2025). Ambiguous persuasion: An ex-ante formulation. *Games and Economic Behavior*. Forthcoming. Earlier version: arXiv:2010.05376.
- Colo, P. (2024). Communicating about confidence: Cheap talk with an ambiguity-averse receiver. *American Economic Journal: Microeconomics* 16(3), 43–75.
- Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research* 23(1), 123–145.
- Glazer, J. and A. Rubinstein (2004). On the optimal rules of persuasion. *Econometrica* 72(6), 1715–1736.
- Glazer, J. and A. Rubinstein (2006). A study in the pragmatics of persuasion: A game theoretical approach. *Theoretical Economics* 1(4), 395–410.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics* 24(3), 461–483.
- Hart, S., I. Kremer, and M. Perry (2017). Evidence games: Truth and commitment. *American Economic Review* 107(3), 690–713.
- Hedlund, J., T. F. Kauffeldt, and M. Lammert (2020). Persuasion under ambiguity. *Theory and Decision* 89(3–4), 455–482.
- Hiriart-Urruty, J.-B. and C. Lemaréchal (2001). *Fundamentals of Convex Analysis*. Berlin: Springer.
- Jung, W.-O. and Y. K. Kwon (1988). Disclosure when the market is unsure of information endowment of managers. *Journal of Accounting research*, 146–153.
- Kellner, C. and M. Le Quement (2018). Endogenous ambiguity in cheap talk. *Journal of Economic Theory* 173, 1–17.
- Li, L. (2022). Ambiguous cheap talk. Working paper.

- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics* 12(2), 380–391.
- Pahlke, M. (2022). Dynamic consistency and ambiguous communication.
- Sher, I. (2011). Credibility and determinism in a game of persuasion. *Games and Economic Behavior* 71(2), 409–419.
- Silva, F. (2020). The importance of commitment power in games with imperfect evidence. *American Economic Journal: Microeconomics* 12(4), 99–113.