

Bayesian Elicitation

Mark Whitmeyer*

February 6, 2024

Abstract

We study how a decision-maker can acquire more information from an agent by reducing her own ability to observe what the agent transmits. In a large class of binary-action games, opacity design is just as good as full commitment to actions and also guarantees that *ex ante* information acquisition always benefits the receiver, even though without opacity design this learning might actually lower the receiver's expected payoff.

Keywords: Signaling, Opacity, Mediation, Information Design, Bayesian Persuasion

JEL Classifications: C72; D82; D83

*Arizona State University

Email: mark.whitmeyer@gmail.com. Earlier versions of this paper were titled “On Optimal Transparency in Signaling Games,” and “Opacity Design in Signaling Games.” This paper also incorporates the previously separate working paper “In Simple Communication Games, When Does *Ex-ante* Fact-finding Benefit the Receiver?” I thank V. Bhaskar, William Fuchs, Rosemary Hopcroft, Vasudha Jain, Meg Meyer, Vasiliki Skreta, Max Stinchcombe, Yiman Sun, Tymon Tatur, Alex Teytelboym, Thomas Wiseman, Joseph Whitmeyer, Kun Zhang, and various seminar and conference audiences for their feedback. An earlier version of this work was the third chapter of my dissertation at UT Austin. I worked on this while at the Hausdorff (T_2) Center for Mathematics in Bonn, supported under the DFG Project 390685813

1 Introduction

There is a decision-maker who must choose an action in an uncertain world. She does not have direct access to information about the state of the world, but there is a second agent who does. The second agent (or sender) observes the state before sending a message¹, which the decision-maker (or receiver) observes before taking an action. This scenario is called a signaling game (or a communication game), and the class of such games includes cheap-talk games, like that studied in Crawford and Sobel (1982); and costly signaling games, like that studied in Spence (1973).

In many situations, the receiver's only concern is information transmission: she values only the information content of the message, and so her welfare increases as the sender becomes more informative. However, there are a number of potential frictions that could impede this communication. First, the sender's and the receiver's preferences over the action taken may be imperfectly aligned. Second, the messages may be costly to the sender, with these costs affecting the messages that are chosen in equilibrium. As a result, less than full information may be transmitted at equilibrium, and the frictions may even be so severe that no information is transmitted.

Now, suppose that the receiver is not forced to observe the sender's message directly, but may commit *ex ante* to observe a noisy signal of the message instead. Can such opacity design help the receiver? How should the receiver design an optimal information structure?

If this were a decision problem in which the information were exogenous, then the answer to the first question would be no. The receiver would always (at least weakly) prefer to observe the message itself rather than some noisy signal. Here, however, the message is not exogenous but is instead an equilibrium choice of the sender. The sender is aware of the information structure, which therefore affects the message that he sends at equilibrium. Consequently, a noisier signal of the sender's message may beget a more informative message, to the receiver's gain.

¹We term the sender's action a message in order to distinguish it from the receiver's action. In some settings, like for instance cheap talk games, the moniker is fitting. In others, like for instance the Spence (1973) model of signaling through education attainment, labeling his action a message is less appropriate.




Figure 1: The Parameterized “Beer-Quiche” Game

The classic “Beer-Quiche” example of [Cho and Kreps \(1987\)](#), reproduced in Figure 1, provides a useful illustration of this idea. The sender is privately informed about whether he is wimpy or strong (θ_W or θ_S , respectively). His message is his choice of fare, either beer (B) or quiche (Q). All else equal, wimps prefer quiche and strong guys beer, but there are strategic concerns as well: the receiver observes the sender’s order before choosing whether to fight him. The receiver prefers to fight the wimp and not fight the strong guy, but the sender does not wish to be fought, regardless of the state. Moreover, in each state, the sender would rather consume his least favorite fare and not be fought than consume his favorite fare and be fought. Both sender and receiver share the common prior $\mu_0 := \Pr(\theta_S) = 2/3$: the sender is less likely to be the wimp.

If the receiver perfectly observes the sender’s message, there are no equilibria in which any information is transmitted. The logic behind this is simple: there can be no equilibrium in which both B and Q are sent in such a way that the receiver strictly prefers to take a different action after each message. In such a circumstance, the wimp always prefers to deviate to the message that is followed by not fight. The receiver’s payoff is $2/3$.

Suppose the scenario is amended to include a neutral third party, a waiter. The receiver no longer observes the sender’s message directly; instead, the waiter witnesses the message before communicating to her. Because the waiter is neutral, he can be represented as a signal, conditional distributions $\pi(\cdot|B)$ and $\pi(\cdot|Q)$ on some set of signal realizations.² Such obfuscation can strictly increase the receiver’s payoff. The optimal signal

²Throughout this paper, the signal, π , is conditioned on the message, m , and not the state, θ .

is

$$\begin{aligned}\pi(f|Q) &= \frac{1}{2}, & \pi(f|B) &= 0 \\ \pi(nf|Q) &= \frac{1}{2}, & \pi(nf|B) &= 1\end{aligned}$$

After nf the receiver does not fight and after f the receiver fights. The signal begets a separating equilibrium—one in which the wimp chooses Q and the strong guy chooses B . This gain in informativeness outweighs the garbling by the signal, and the receiver obtains a strictly higher equilibrium payoff than without opacity ($5/6$ versus $2/3$).

Throughout this paper, we explore the problem of solving for the receiver-optimal³ information structure in signaling games. We restrict attention to the environment in which the sender and receiver share a common prior and focus on signaling games in which the receiver has no intrinsic preferences over the message chosen by the sender.

There is a useful upper bound for the solution to the opacity design problem; the *commitment* solution for the receiver, which corresponds to the scenario in which the receiver can commit to a distribution over actions conditioned on the sender’s choice of message. Consequently, if the commitment strategy that maximizes the receiver’s payoff also satisfies the receiver’s obedience constraints, then it must correspond to the optimal information structure. The commitment problem is much simpler, and we establish that it reduces to a linear program.

In the first main result of this paper, Proposition 3.3, we discover that in any signaling game with two actions, “opacity equals commitment”. Namely, for any number of states and messages, provided the receiver has a binary decision, the commitment solution satisfies the receiver’s obedience constraints.

We also investigate the benefits and drawbacks to the receiver of *ex ante* information acquisition, and how the consequences thereof are affected by the receiver’s ability to design the opacity in the ensuing game. While it is clear that in any setting, information acquisition may benefit the receiver, our goal is to discover in which settings *ex ante* learning is always beneficial. Surprisingly, even in games with just two actions, *ex ante*

³We focus throughout on the receiver-optimal Perfect Bayesian Equilibrium.

information acquisition may hurt the receiver if she may not design the opacity in the ensuing game. However, the second main result of this paper, Proposition 5.4, establishes that opacity design flips this result. In two-action games, the value of information is always (weakly) positive, provided the receiver may design the opacity in the game that follows.

1.1 Discussion of Assumptions and Relevance

Let us briefly discuss the relevance of this paper’s theoretical environment.

1.1.1 Real-World Instances of Opacity Design

There are coarse ratings or grades, in which case the signal is the policy that maps the sender’s behavior to ratings or grades; recommender or aggregation systems, in which case the signal is the aggregator or recommender who distills product information like warranty details, return policies, advertising and even price details,⁴ into a purchase recommendation for a consumer; and interactions through third parties (mediators), in which case it is the mediator itself that serves as the signal. The “Beer-Quiche” game discussed above can be viewed as an allegory for an incumbent firm signaling to a potential entrant through its pricing decisions. Accordingly, we argue a role for a third party consultant who advises the aspiring entrant. Transparency is often lauded as a desirable feature in politics. However, political decision makers have private information, so their actions can carry information. Our framework shows how a neutral free press that in some circumstances obfuscates and hides an incumbent’s actions may benefit a populace.

1.1.2 Who Are These Mediators?

One profession that arguably fills this role is headhunting or staffing firms. There are many non-productive activities that prospective job candidates take part in that are nevertheless useful to a hiring firm due to the information these behaviors carry about the

⁴These are all classic signaling environments; see [Milgrom and Roberts \(1986\)](#); [Nelson \(1974\)](#) and [Gal-Or \(1989\)](#).

hidden productivity of the candidates. In [Spence \(1973\)](#) it is education that is such a (intrinsically) wasteful activity. This paper, therefore, illustrates the usefulness of hiring through third party agencies, who serve as intermediaries and thereby improve welfare even if they do not possess any special expertise in identifying talent.

1.1.3 Focusing on the Receiver's Design Problem.

In a number of examples, it seems realistic to have the receiver design the mechanism. For instance, consider a university that is faced with the problem of admitting students. It is the school that designs the admissions form, which is effectively the opacity-design problem under study in this work: the university can allow the student to enter his standardized test score in full detail, only a coarse categorization of the score, or no information about the score. Naturally, the university does this to optimize its (the receiver's) welfare. Similarly, in the staffing-firm and consulting examples mentioned above, it is the receiver—hiring firm and potential entrant, respectively—that typically pays the staffing firm (or consultant) for its service. Given this, the intermediary (in order to maximize its own rents) designs the mechanism to be receiver-optimal.⁵

1.1.4 Relevance of Two Actions.

Our main results are based on a binary-action specification, which may be limiting in some environments, as receivers may have more than two available actions. However, many important scenarios *do* only have two options. For instance, a university admits or rejects an applicant, a populace keeps or ousts an incumbent, and a consumer purchases or refrains from buying a product. Additionally, a hiring firm either hires or rejects an applicant, and a firm considering market entry chooses whether to enter or stay out of the market.

⁵Moreover, it is in the intermediary's interest to be *unbiased* at an interim point, as otherwise its services are worthless. This can be sustained, for instance, through reputation.

1.2 Related Work

The idea that introducing noise to communication can enhance welfare is not new to this paper. In particular, a number of papers study this in various cheap talk settings. Forges (1990) explores mediation in a (cheap-talk) job-market example; and Myerson (1991) describes a game in which communication is improved if occasionally messages are not transmitted.⁶ Blume et al. (2007) shows how noise can improve welfare in a uniform-quadratic setting, and other papers follow—Goltsman et al. (2009); Ivanov (2010); and Ganguly and Ray (2011)—that look at the benefits of mediation in the uniform-quadratic environment. More recently, Salamanca (2021) explores a general version of optimal mediation using duality theory from linear programming.

Goltsman et al. (2009) and Blume et al. (2007) focus on receiver-optimal equilibria, as do we in this paper.⁷ Salamanca (2021), in contrast, looks at the design problem from the sender’s perspective. In that light, his paper is more true to the standard Bayesian persuasion paradigm, in which a sender produces information to influence a receiver. Here, we tackle an extraction (or elicitation) problem instead: how can a receiver induce a sender to reveal information to her?

Because the messages are cheap talk, these papers view the mediation problem as a centralized problem and use the *Communication Equilibrium* concept, as formulated in Myerson (1986) and Forges (1986). Here, provided the game is cheap talk and the set of messages is sufficiently large, the design of the optimal signal, π , is equivalent to the centralized mediator-driven problem analyzed in these cheap talk papers (we can think of the separating equilibrium as the reporting of the state). However, if the message set is not big enough, then this equivalence is lost. Moreover, if the messages are costly (i.e., the game is not cheap talk), then the problem we explore is different from the centralized problems of the literature. With costly messages, it is not even without loss of generality to restrict the sender to pure strategies.

It is more difficult to obtain general results when messages are not cheap. In his

⁶Amusingly, the medium for communication is a (somewhat wayward) carrier pigeon.

⁷Ivanov (2010) and Ganguly and Ray (2011) investigate slightly different questions within the uniform quadratic framework, rather than tackle a design problem.

exploration of signaling games, [Rick \(2013\)](#)'s main result is that welfare is improved by a signal only if either i. the signal removes some profitable deviations to unused messages and/or ii. the signal garbles the distribution over posteriors induced by the sender's strategy non-trivially. [Ball \(2019\)](#) investigates the design of a scoring rule in order to elicit information from a sender who can distort multiple features (what the receiver can observe) about himself. Analogously, he finds that coarser information can benefit the receiver due to the endogeneity of the information, which is produced by the sender. In a similar setting, [Frankel and Kartik \(2022\)](#) look at the commitment problem for a principal acquiring manipulable information from an agent.

In contrast to these papers, this paper highlights the power of opacity design in signaling games. First, it is just as good for the receiver as full commitment; and second, it guarantees that the value of *ex ante* information is always positive. This first equivalence is especially useful: it transforms a large and difficult problem into a much easier linear program. It is not uncommon for applied papers to include a commitment benchmark, and here we find that this benchmark can be obtained through information alone.

The equivalence of opacity design and commitment also sheds light on how increased noise can improve the receiver's welfare. Because the obedience constraints are always satisfied by the commitment solution, the receiver's interim incentives (at the optimum) are inconsequential. Accordingly, maximizing the receiver's payoff is constrained only by the sender's incentives.

This paper is also related to the recent line of information design papers in which the information generation process (or even the state) is endogenous. Those papers include [Boleslavsky and Kim \(2017\)](#), who study a persuasion problem in which the endogenous state depends on the hidden effort of a third party; [Georgiadis and Szentos \(2020\)](#), who ask how to optimally monitor an agent; and [Asriyan et al. \(2021\)](#), who investigate how a planner should reveal information about past trading volume. The closest of these to this work is [Boleslavsky and Kim \(2017\)](#), whose problem reduces formally to a constrained persuasion problem, the constraints to which correspond to the incentive compatibility conditions of the effort-providing agent. The opacity design problem is a constrained persuasion problem, and we use the Lagrangian approach introduced by [Boleslavsky and](#)

Kim (2017) (see also Doval and Skreta (2023)) to solve the game of § 4.2.

Finally, this paper is also related to the rational inattention literature, which explores the effects of costly information processing and/or acquisition on an agent’s (or agents’) behavior. Mackowiak et al. (2022) provides an excellent survey of this area. At a fundamental level, our receiver can be viewed as solving a rational inattention problem, where the menu of feasible information acquisition strategies is determined by the various incentive constraints (obedience and sender incentive compatibility).⁸ In contrast to the literature; however, in which more information is typically assumed to be more expensive, this effective cost of “acquiring” information can take unusual forms. For instance, a game that admits a separating equilibrium but not equilibria with certain varieties of partial state revelation would be such that full information (learning the state) is costless for the receiver, but some garblings of the prior are infinitely costly.

2 The Model

The setup is a version of the standard signaling game. There are two players: a sender, S ; and a receiver, R . The sender is privately informed about the state $\theta \in \Theta$ and chooses a message, m , from the set of messages M . The receiver observes m , but not θ , updates her belief about the state using Bayes’ law, then chooses a mixture over actions a from the set of actions A . We assume that sets A , M , and Θ are compact.

S and R share a common prior over the state of the world, $\mu \in \Delta(\Theta)$, where $\mu(\theta) := \Pr(\theta)$. Each player, S and R , has state-dependent preferences over the message sent, and the action taken, which are represented by the continuous utility functions $u_S: A \times M \times \Theta \rightarrow \mathbb{R}$ and $u_R: A \times \Theta \rightarrow \mathbb{R}$. A behavioral strategy for S , $\sigma(\cdot|\theta)$, is a family of probability distributions over M . Similarly, a behavioral strategy for R , $\rho(\cdot|m)$, is a family of probability distributions over A . The special case in which messages are costless and the message space is large is the classical cheap-talk setting. We focus on receiver-optimal Perfect Bayesian Equilibria (PBE), defined in the standard manner.⁹

⁸I thank a referee for pointing out this perspective.

⁹Strictly speaking, we focus on receiver-optimal Perfect Bayesian Equilibria of the extended game in

2.1 Opacity Design

Suppose that R can design the game's opacity by committing *ex ante* to observe the realization of a signal π , rather than m . A signal, π , is a mapping $\pi: M \rightarrow \Delta(X)$, where $\pi(x|m) := \Pr(X = x|M = m)$ and X is a (large) finite set of signal realizations. The timing is as follows: first, R commits (publicly) to π ; second, S mixes over actions according to $\sigma(\cdot|\theta)$; and third, R observes the signal realization x , before choosing a behavioral strategy $\rho: X \rightarrow \Delta(A)$.

The receiver solves

$$\sup_{\sigma, \pi, \rho} \mathbb{E}_{\mu, \sigma, \pi, \rho} [u_R(a, \theta)] ,$$

such that

$$\rho(a|x) \in \arg \max_{\rho(a|x)} \mathbb{E}_{\mu, \sigma, \pi, \rho} [u_R(a, \theta)|x] , \quad (O_x)$$

for all $x \in X$;¹⁰ and

$$\sigma(m|\theta) \in \arg \max_{\sigma(m|\theta)} \mathbb{E}_{\sigma, \pi, \rho} [u_S(a, m, \theta)|\theta] , \quad (IC_\theta)$$

for all $\theta \in \Theta$.

The second group of conditions (Conditions IC_θ) collects the optimality conditions for the sender typically found in signaling games (with noise): given the (expected) best responses by the receiver to a sender's message, in each state the sender's message must be optimal. The first set of conditions (Conditions O_x) collects the obedience constraints for the receiver, which mandate that her choices following each signal realization x are sequentially rational. The term "obedience" hints at the first result, in which we note that a revelation principle applies. This allows us to greatly simplify the receiver's problem by restricting attention to direct signals that recommend actions.

Definition 2.1. Call a signal π **Direct** if $X = A$. That is, a direct signal recommends actions.

which the receiver commits *ex ante* to a garbling of the sender's message, π .

¹⁰Note that following any signal realization, the receiver will have beliefs concerning the sender's message and the state. Concordant with our focus on PBE, these beliefs must be consistent with Bayes' law whenever possible.

Then, the following remark establishes that it is without loss of generality to restrict attention to direct signals, which recommend actions to the receiver.

Remark 2.2. For any equilibrium triple, (σ, π, ρ) , there is another equilibrium triple, (σ', π', ρ') , that yields the same (terminal) joint distribution over messages and actions, and hence the same payoffs; where π' is a direct signal, $\sigma = \sigma'$, and $\rho'(a|a') = 1$ for $a' = a$ and $\rho'(a|a') = 0$ for $a' \neq a$.

In the remainder of the paper, we restrict attention to direct signals, π , and thus the receiver's problem can be reduced to

$$\sup_{\sigma, \pi} \mathbb{E}_{\mu, \sigma, \pi} [u_R(a, \theta)] ,$$

such that

$$\mathbb{E}_{\mu, \sigma, \pi} [u_R(a, \theta)|a] \geq \mathbb{E}_{\mu, \pi, \sigma} [u_R(a', \theta)|a] , \quad (O_a)$$

for all $a, a' \in A$; and

$$\mathbb{E}_{\pi} [u_S(a, m, \theta)|\theta] \geq \mathbb{E}_{\pi} [u_S(a, m', \theta)|\theta] , \quad (IC_{\theta})$$

for all $m \in \text{supp } \sigma(\cdot|\theta)$, for all $m' \in M$, for all $\theta \in \Theta$. We term this program the **Opacity Design Problem**, and label a pair (σ, π) that solves this program the **Opacity Design Solution**. Define $V = V(\mu)$ to be the value of this constrained optimization problem as a function of the prior. We call this the receiver's **Opacity Design Payoff**.

It is important to note that if one were to restrict the *sender* to a pure strategy in each state, one would sacrifice generality. That is, the opacity design solution may require the sender to mix in some states (which is easy to show through a counterexample). However, as noted earlier, if the game is cheap talk (and there are at least as many messages as states) then the sender may be assumed to "report" the state in her message and hence play a pure strategy in the opacity design solution (this is just the Revelation Principle).

3 Opacity Design = Commitment

If the receiver has the stronger form of commitment to actions—i.e., instead of choosing a signal π , she chooses a conditional distribution over actions as a function of the sender’s message, (with abuse of notation) π .¹¹ This program, which we term the **Commitment Problem** is

$$\sup_{\sigma, \pi} \mathbb{E}_{\mu, \sigma, \pi} [u_R(a, \theta)] ,$$

such that

$$\mathbb{E}_{\pi} [u_S(a, m, \theta) | \theta] \geq \mathbb{E}_{\pi} [u_S(a, m', \theta) | \theta] , \quad (IC_{\theta})$$

for all $m \in \text{supp } \sigma(\cdot | \theta)$, for all $m' \in M$, for all $\theta \in \Theta$. We label a pair (σ, π) that solves this problem the **Commitment Solution**. The receiver’s value is her **Commitment Payoff**.

Since the receiver’s commitment problem is merely the opacity design problem absent the obedience constraints, the following observation is evident:

Remark 3.1. Any (terminal) joint distribution over messages and actions that can be obtained through opacity design can be obtained through commitment to actions.

As the next result illustrates, the commitment problem is much easier to solve than the opacity design problem, since we may restrict the sender to pure strategies in each state. Indeed, the result is trivial: the receiver’s objective is linear in the probability of any message in the support of a sender’s mixed strategy; and in any state in which the sender is mixing, he must be indifferent over all pure strategies in the support of his mixed strategy. Formally,

Lemma 3.2. *There is a commitment solution in which the sender chooses a pure strategy in each state.*

One consequence of this result is that there is a commitment solution for the receiver in which at most $|\Theta|$ messages are used. Furthermore, for a fixed vector of pure strategies, the receiver’s commitment problem is a linear program. Thus, the receiver’s commit-

¹¹We are merely replacing “recommendations” with “orders.”

ment problem is merely a finite collection of linear programming problems. This also immediately implies existence of an optimum in the commitment problem.

Next, we discover the first main result:

Proposition 3.3 (Opacity Design Equals Commitment). *If $|A| = 2$, the opacity design solution coincides with the commitment solution.*

Proof. Recall that in the commitment problem it is without loss of generality to restrict the sender to pure strategies in each state. Thus, the receiver's commitment payoff simplifies to

$$\sum_{i=1}^n \mu(\theta_i) [\pi_i v_i + (1 - \pi_i) w_i] ,$$

where $v_i := u_R(a_1, \theta_i)$, $w_i := u_R(a_2, \theta_i)$ and $\pi_i := \pi(a_1 | m_i)$. Without loss of generality, the receiver's optimal action under the prior is a_1 . Accordingly,

$$\sum_{i=1}^n \mu(\theta_i) [\pi_i v_i + (1 - \pi_i) w_i] \geq \sum_{i=1}^n \mu(\theta_i) [\pi_i v_i + (1 - \pi_i) v_i] \geq \sum_{i=1}^n \mu(\theta_i) [\pi_i w_i + (1 - \pi_i) w_i] ,$$

where the first inequality uses the fact that the receiver's commitment payoff must be weakly higher than her payoff from taking the action that is optimal under the prior and the second inequality follows by the optimality of a_1 under the prior.

This chain of inequalities implies that the two obedience constraints are satisfied:

$$\sum_{i=1}^n \mu(\theta_i) \pi_i v_i \geq \sum_{i=1}^n \mu(\theta_i) \pi_i w_i , \tag{O1}$$

$$\sum_{i=1}^n \mu(\theta_i) (1 - \pi_i) w_i \geq \sum_{i=1}^n \mu(\theta_i) (1 - \pi_i) v_i . \tag{O2}$$

■

3.1 Multiple Senders

In some sense, the insight provided by Proposition 3.3 is quite general: in fact, although this paper focuses two-player signaling games, opacity equals commitment for costly sig-

naling games with an arbitrary number of senders provided the receiver has only two actions. To put differently, from the receiver’s (principal’s) perspective, *information design is just as good as mechanism design (without transfers)*. The exact same logic from the proof of Proposition 3.3 goes through (though it is more cumbersome notationally since senders may be mixing): the receiver cannot strictly benefit from deviating at the interim stage (following a recommendation) since that would contradict *ex ante* optimality.

A special case of such a multi-sender environment is the classic mechanism design without transfers setting. There, the binary-action case is of particular relevance since it pertains to a principal’s (in our parlance, receiver’s) decision to approve a policy or candidate or convict a defendant—see for instance Wolinsky (2002), Feddersen and Gradwohl (2015), Battaglini (2017), Gradwohl and Feddersen (2018), Ali and Bohren (2019), or Katwinkiel and Winter (2021), all of whom center their analysis around a principal’s ability to commit to a binary decision. Our results imply that the strong form of commitment (to actions) assumed in these papers can be weakened to one of commitment to information, without affecting the receiver-optimal outcome.

Formally, suppose that rather than there only being one sender, there are $k \geq 1$ ($k \in \mathbb{N}$) senders S_i ($i = 1, \dots, k$), each of which simultaneously sends a message from her (possibly idiosyncratic) set of messages M_i . Each sender is informed of her dimension of the multi-dimensional state $\theta_i \in \Theta_i$ and chooses a behavioral strategy $\sigma_i(\cdot|\theta_i)$, which is a probability distribution over messages. For this subsection alone, we redefine $M := \times_i M_i$, $\Theta := \times_i \Theta_i$, and $\sigma := (\sigma_1, \dots, \sigma_k)$. Sender i ’s utility function is $u_{S_i}: A \times M \times \Theta \rightarrow \mathbb{R}$ and the receiver’s utility is $u_R: A \times \Theta \rightarrow \mathbb{R}$. All are continuous. μ denotes the common prior and μ_{-i} the joint distribution over states other than θ_i , given state θ_i . As is standard, let $\sigma_{-i}(m_{-i})$ denote the vector of mixed (pure) strategies chosen by senders other than i .

As in the single-sender setting, the receiver commits *ex ante* to a signal, mapping $\pi: M \rightarrow \Delta(A)$, where we have used an analog of Remark 2.2 to allow us to restrict attention to direct signals. The receiver’s opacity design problem is now

$$\sup_{\sigma, \pi} \mathbb{E}_{\mu, \sigma, \pi} [u_R(a, \theta)] ,$$

such that

$$\mathbb{E}_{\mu,\sigma,\pi} [u_R(a, \theta) | a] \geq \mathbb{E}_{\mu,\pi,\sigma} [u_R(a', \theta) | a] , \quad (O_a)$$

for all $a, a' \in A$; and

$$\mathbb{E}_{\pi,\mu_{-i},\sigma_{-i}} [u_{S_i}(a, m_i, m_{-i}, \theta_i) | \theta_i] \geq \mathbb{E}_{\pi,\mu_{-i},\sigma_{-i}} [u_{S_i}(a, m'_i, m_{-i}, \theta_i) | \theta_i] , \quad (IC_{\theta_i})$$

for all $m_i \in \text{supp } \sigma_i(\cdot | \theta_i)$, for all $m'_i \in M_i$, for all $\theta_i \in \Theta_i$, for all $i = 1, \dots, k$. As in the single-sender case, the commitment problem is just this program without the collection of obedience constraints (O_a).

An exact analog of Proposition 3.3 holds:

Proposition 3.4 (Opacity Design Equals Commitment). *With $k \geq 1$ senders, if $|A| = 2$, the opacity design solution coincides with the commitment solution.*

We omit the proof as it is identical to the proof of Proposition 3.4 modulo some minor notational changes.

4 A “Match-the-Binary-State” Game

Let us briefly explore a simple binary-state game in order to i. investigate how the sender’s and receiver’s conflict of interest affects the optimal signal, ii. study the limits of “Opacity = Commitment,” and iii. illustrate how to solve the opacity design via Lagrangian methods.

Throughout this section, we impose that there are just two states, $\Theta = \{0, 1\}$; and the set of actions for the receiver is a subset of $[0, 1]$. The receiver’s payoff, given action a and realized state θ , is the well-known quadratic loss function:

$$U_R(a, \theta) = -(a - \theta)^2 .$$

The sender has just two messages $M = \{0, 1\}$. In addition, he has bias $b \in [\frac{1}{2}, 1)$. We also

stipulate that the sender suffers a “lying cost” of $c \in [0, 2b - 1)$. His payoff is, therefore,

$$U_S(a, \theta, m) = \begin{cases} -(a - (\theta + b))^2 & \text{if } m = \theta \\ -(a - (\theta + b))^2 - c & \text{if } m \neq \theta \end{cases} .$$

We begin by assuming the receiver has just two actions, in which environment Proposition 3.3 holds. We then add one additional action and show that it fails to hold in general before, finally, solving a continuum of actions version of the game.

4.1 Two Versus Three Actions

First, suppose $A = \{0, 1\}$. By Proposition 3.3, we need only solve the commitment problem for the receiver: she commits to two conditional distributions over actions $p := \mathbb{P}(1|1)$ and $q := \mathbb{P}(1|0)$. By the revelation principle (for scenarios with reporting costs, see [Kephart and Conitzer \(2016\)](#)) it is optimal for the receiver to have the sender choose message 1 when the state is 1 and message 0 otherwise. Thus, the receiver solves the program (where we have omitted constants)

$$\max_{p, q} \{-\mu(1 - p) - (1 - \mu)q\} ,$$

subject to

$$-pb^2 - (1 - p)(1 + b)^2 \geq -qb^2 - (1 - q)(1 + b)^2 - c , \quad (IC_1)$$

and

$$-q(1 - b)^2 - (1 - q)b^2 \geq -p(1 - b)^2 - (1 - p)b^2 - c . \quad (IC_0)$$

By our parametric assumptions Constraint IC_1 does not bind, and so the optimum is given by a binding Constraint IC_0 and either $q = 0$ or $p = 1$, the latter if and only if $\mu \geq 1/2$. Thus,

Remark 4.1. Unless $\mu = 1/2$, the optimal signal is unique. It is given as follows:

- (i) If $\mu > 1/2$, $p^* = 1$ and $q^* = 1 - c/(2b - 1)$.
- (ii) If $\mu < 1/2$, $p^* = c/(2b - 1)$ and $q^* = 0$.

(iii) If $\mu = 1/2$, p^* ranges from $c/(2b - 1)$ to 1 and $q^* = p^* - c/(2b - 1)$.

Now suppose $A = \{0, \hat{a}, 1\}$, where $\hat{a} \in (2b - 1, 1)$. In the receiver's commitment problem, she commits to four conditional distributions over actions $p := \mathbb{P}(1|1)$, $\hat{p} := \mathbb{P}(\hat{a}|1)$, $q := \mathbb{P}(1|0)$ and $\hat{q} := \mathbb{P}(\hat{a}|0)$. Leaving the details to Appendix A.2, we have:

Remark 4.2. Except for a knife-edge case, the commitment solution is unique. If the prior is sufficiently high, the commitment solution is $q^* = \hat{p}^* = 0$, $p^* = 1$, and $\hat{q}^* = \frac{2b-c-1}{\hat{a}(2b-\hat{a})}$. Otherwise, the commitment solution is $q^* = \hat{p}^* = \hat{q}^* = 0$, and $p^* = \frac{c}{2b-1}$.

Notably, the commitment solution does not satisfy the obedience constraints when the prior is sufficiently high: when the receiver is recommended \hat{a} by the mechanism, she knows the sender is type 0 and is unwilling to follow the recommendation. Furthermore, the explicit expression for a “sufficiently high” prior is

$$\frac{\mu}{1-\mu} \geq \frac{(2b-1)\hat{a}}{2b-\hat{a}}.$$

The left-hand side of this is strictly increasing in b , so we see that commonality of interest (a smaller b) need not make the commitment solution easier to obtain via opacity design. If the prior is not “sufficiently high” (the stated inequality does not hold), the commitment solution satisfies the obedience constraints, so the receiver can obtain the commitment solution with the weaker form of commitment corresponding to opacity.

This exercise illustrates that the finding of Proposition 3.3 need not persist if the receiver has three or more actions.¹² What is it that makes the two-action case special? The crucial aspect of the receiver's binary decision is that, following a recommendation by the signal, there is *only one way in which she can be disobedient*. Consequently, disobeying a recommendation would yield her the same payoff from an *ex ante* perspective as if there were no information transmitted whatsoever, since she would be taking just one action. However, she can always trivially obtain that payoff by choosing a completely uninformative signal. Thus, disobedience (of the signal corresponding to the commitment solution) must not be profitable.

¹²The setting inhabited by Ball (2019) is another in which commitment is better for the receiver than opacity.

In the commitment solution, when μ is high, it is important for the receiver to take the correct action when the state is 1, as it occurs so frequently. In order to maintain incentive compatibility for the sender in state 0, this requires taking a non-zero action with positive probability. Moreover, it is always best (with commitment) to take action \hat{a} . This cannot be done in the opacity design problem; however, as it exposes the low state.

4.2 A Continuum of Actions

Now let $A = [0, 1]$.

4.2.1 Commitment

Suppose first the receiver can commit to a distribution over actions as a function of the sender's message. Formally, the receiver commits to two conditional distributions over actions $P(a) := \mathbb{P}(A \leq a|1)$ and $Q(a) := \mathbb{P}(A \leq a|0)$. Appealing to the revelation principle, the receiver solves

$$\max_{P, Q} \left\{ -\mu \int_0^1 (a-1)^2 dP(a) - (1-\mu) \int_0^1 a^2 dQ(a) \right\},$$

subject to

$$-\int_0^1 (a-1-b)^2 dP(a) \geq -\int_0^1 (a-1-b)^2 dQ(a) - c, \quad (IC_1)$$

and

$$-\int_0^1 (a-b)^2 dQ(a) \geq -\int_0^1 (a-b)^2 dP(a) - c. \quad (IC_0)$$

Absent the IC conditions, the receiver would choose a degenerate P and Q on 1 and 0, respectively. However, this is not incentive compatible: the sender prefers to lie in state 0. Despite this, this problem is not too difficult. We introduce the following function of the bias b :

$$\gamma(b) := 2b - 1 - b^2 \mu(2 - \mu),$$

and state the following proposition, leaving its derivation to the appendix.

Proposition 4.3. *If the prior is sufficiently high, $c \geq \gamma(b)$, in the optimal commitment solution,*

both P and Q are degenerate distributions. Otherwise, Q is degenerate but P is binary, with support on $\{0, 1\}$.

This proposition is similar in spirit to Remark 4.2. There is a tension the receiver must resolve in choosing P and Q . When μ is high, it is relatively more damaging to randomize (which hurts the receiver due to the strict concavity of her objective) in state 1. Accordingly, the receiver prefers a degenerate distribution, even though it makes deviating attractive for the sender in state 0.

In this game, the receiver does strictly better by having a strong form of commitment to actions instead of only information design:

Corollary 4.4. *In this game opacity \neq commitment.*

4.2.2 Opacity Design

Because of the quadratic loss specification, the obedience constraint can be folded into the optimization problem in a concise way. In particular, given any posterior belief $x := \mathbb{P}(1)$, the receiver's optimal action is merely x , yielding a payoff (as a function of the belief) $-x(1-x)$. It is again without loss of generality to restrict the sender to truthful "reports" of the state, and the receiver's choice variable is any distribution F in the set of Bayes-plausible distributions $\mathcal{F}(\mu)$. Naturally, the sender's reports must be incentive compatible, so the receiver solves

$$\max_{F \in \mathcal{F}(\mu)} \left\{ - \int_0^1 x(1-x) dF(x) \right\},$$

subject to

$$- \int_0^1 (x-b-1)^2 \frac{x}{\mu} dF(x) dx \geq - \int_0^1 (x-b-1)^2 \frac{1-x}{1-\mu} dF(x) dx - c, \quad (IC_1)$$

and

$$- \int_0^1 (x-b)^2 \frac{1-x}{1-\mu} dF(x) dx \geq - \int_0^1 (x-b)^2 \frac{x}{\mu} dF(x) dx - c. \quad (IC_0)$$

It is clear that the first constraint (as in the commitment problem) must be slack. Moreover, following [Doval and Skreta \(2023\)](#) we can write the following Lagrangian corresponding to the receiver's *minimization problem*:

$$\mathcal{L}(x) = x(1-x) + \lambda \left(-c + \frac{(\mu-x)(x-b)^2}{(1-\mu)\mu} \right).$$

Leaving the details (and a figure) to the appendix, we unearth the following proposition:

Proposition 4.5. *There is a uniquely optimal signal. It has two realizations and yields posteriors $\{\underline{x}, 1\}$, where $\underline{x} \in (0, \mu)$.*

At the optimum, the sender in state 1 is occasionally exposed, or else is pooled with the sender in state 0 by the signal. The natural relationship between the conflict of interest and the informativeness of the signal established in [Remark 4.1](#) persists:

Corollary 4.6. *As incentives become more aligned, or the cost of misrepresentation increases (b decreases or c increases, respectively), the optimal signal becomes more informative.*

In this game, any lack of informativeness is due to the problem of providing incentives to the sender. As this concern lessens, the optimal signal adjusts to improve the receiver's decision-making.

A comparison between the opacity design solution in this game to the commitment solution is illuminating. In particular, observe that with opacity design the optimal signal is always stochastic: the sender's message in state 1 is occasionally mixed with the sender's message in state 0. In contrast, in the commitment solution, it is only when incentives are not aligned and the cost of misrepresentation is low that the receiver's actions are random. With opacity design, the recommendations by the signal must be random, in order for the receiver to be willing to obey its directives. There is no such issue with commitment to actions, randomization is only needed as a direct incentive device that mechanically makes a deviation by the sender in state 0 less enticing.

5 The Rewards to *ex ante* Learning

We established above that in two-action games opacity design is as good as commitment. Its usefulness is not limited merely to the game itself. As we discover in this section, it also guarantees that the value of *ex ante* information is always (weakly) positive.

Let us revisit the “Beer-Quiche” example. Suppose that *ex ante*, prior to the signaling game, the receiver has a chance to acquire public information. She may for instance scrutinize the sender’s clothing, or the company that he keeps, both of which are correlated with the state. Must such *ex ante* information acquisition benefit the receiver?

We focus on two cases: that in which the receiver may design the opacity in the ensuing game, and that in which she may not. Recall that the opacity design payoff is $V = V(\mu)$. We denote the receiver’s payoff without opacity design as $V^T = V^T(\mu)$.¹³ We wish to discover whether these functions are convex. Indeed, the convexity of these functions is equivalent to the value of information being (weakly) positive—any secant line lies above the function so any splitting of a belief must yield an (expected) payoff that is weakly higher than the payoff at the belief.

In the example, it is straightforward to verify that V and V^T are

$$V(\mu) = \begin{cases} 1 - \frac{\mu}{2}, & \mu \leq \frac{1}{2} \\ \frac{1+\mu}{2}, & \mu \geq \frac{1}{2} \end{cases}, \quad \text{and} \quad V^T(\mu) = \begin{cases} 1 - \mu, & \mu \leq \frac{1}{2} \\ \mu, & \mu \geq \frac{1}{2} \end{cases},$$

which are convex. Hence, regardless of whether the receiver can design the opacity, *ex ante* learning is always (at least weakly) beneficial. As we will see shortly, in two-action games *ex ante* information acquisition is always good, provided the receiver can design the opacity in the ensuing game. If she cannot, then information up front may actually hurt the receiver. That is, there is information that she would refuse, even if it were free.

Formally, we model *ex ante* information acquisition or fact-finding as follows. Fix a signaling game, and suppose that prior to participating in the game, the receiver may acquire public information. Initially, the receiver and the sender share some prior, μ_0 ,

¹³Formally, the **Payoff Without Opacity Design** is the receiver’s maximal payoff in the receiver-optimal PBE of the signaling game without opacity design when the prior is μ .

and there is some finite (or at least compact) set of signal realizations Y and a signal, $\zeta: \Theta \rightarrow \Delta(Y)$, whose realization is public. Call ζ the **Initial Experiment**. This begets a distribution over posteriors, where the posterior following signal realization y is μ_y . The sender and receiver then take part in the signaling game, where the common prior for the game is now μ_y .

Obviously, regardless of the game, there are some experiments that benefit the receiver (or at least do not hurt her): for instance a fully informative initial experiment always benefits the receiver; but surprisingly, there are also experiments that hurt the receiver.

Example 5.1. There are four states, $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$, and a belief is a quadruple $(\mu_1, \mu_2, \mu_3, \mu_4)$, where $\mu_i := \Pr(\Theta = \theta_i)$ for all $i = 1, 2, 3, 4$ and $\mu_1 + \mu_2 + \mu_3 + \mu_4 = 1$.

The belief can be fully described with just three variables; hence, depicting the receiver's payoff as a function of the belief requires four dimensions. The current (two-dimensional) medium of this paper renders this impossible, so instead we restrict attention to a family of experiments that involve learning on just one dimension. That is, we fix $\mu_1 = 1/3$ and $\mu_3 = 1/8$, and consider only the receiver's payoff as a function of her (prior) belief about states θ_2 and θ_4 . Learning is on just one dimension, and so (abusing notation) we rewrite the receiver's belief μ_2 as μ and μ_4 as $13/24 - \mu$, where $\mu \in [0, 13/24]$.

In states θ_1 and θ_2 , action a_2 is the correct action for the receiver; and in states θ_3 and θ_4 , action a_1 is correct:

Action	θ_1	θ_2	θ_3	θ_4
a_1	0	0	1	2
a_2	1	1	0	0

Likewise, the sender's state-dependent payoffs from message, action pairs are

state	θ_1			θ_2			θ_3			θ_4		
	m_1	m_2	m_3	m_1	m_2	m_3	m_1	m_2	m_3	m_1	m_2	m_3
a_1	1	0	0	0	1	0	0	1	3	-1	2	-2
a_2	1	0	0	0	1	0	2	4	0	5/4	0	-1

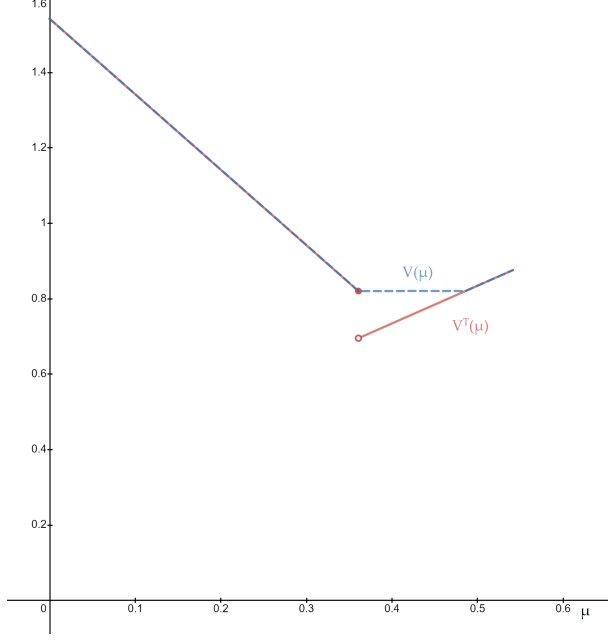


Figure 2: Payoffs With and Without Opacity Design

In Figure 2 we depict the opacity design payoff (V) and the payoff without opacity design (V^T). Explicitly, those functions are

$$V(\mu) = \begin{cases} \frac{37}{24} - 2\mu, & \mu \leq \frac{13}{36} \\ \frac{59}{72}, & \frac{13}{36} \leq \mu \leq \frac{35}{72} \\ \frac{1}{3} + \mu, & \frac{35}{72} \leq \mu \leq \frac{13}{24} \end{cases}, \quad \text{and} \quad V^T(\mu) = \begin{cases} \frac{37}{24} - 2\mu, & \mu \leq \frac{13}{36} \\ \frac{1}{3} + \mu, & \frac{13}{36} < \mu \leq \frac{13}{24} \end{cases}.$$

Without opacity design, *ex ante* information acquisition can hurt the receiver by eliminating particularly lucrative equilibria. In this example, if $\mu \leq 13/36$, there is an equilibrium with partial separation, which provides some useful information to the receiver. On the other hand, should μ strictly exceed $13/36$, the only equilibria beget the pooling payoff—the receiver becomes too sure that the state is not θ_3 or θ_4 . Thus, no useful information is transmitted, to the receiver's detriment.

As Figure 2 illustrates, the issue engendered by *ex ante* learning—that the resulting belief may beget a strictly worse equilibrium in the signaling game—is ameliorated if the receiver can design the game's opacity. The optimal π allows for a moderately informative equilibrium on the interval $[13/36, 35/72]$, to the receiver's benefit.

This counterexample suffices to prove the following proposition:

Proposition 5.2. *In signaling games with four or more states, three or more messages, and two actions, the receiver’s payoff without opacity design is not generally convex in the prior.*

Although the example illustrates that learning may be injurious when the receiver is unable to design the opacity in the game that follows, it also suggests that things may be different if the receiver can design the opacity in the ensuing game. In the example, opacity design renders the value of *ex ante* information positive. As the next pair of results illustrate, this result is general.

With commitment power, information acquisition is always beneficial:

Lemma 5.3. *The commitment payoff is convex in the prior, μ .*

Recall the earlier result—that opacity equals commitment. Intuitively, the receiver’s commitment problem is simply a decision problem, once the sender’s best response (to the commitment) is folded into the objective. Thus, Proposition 3.3 and Lemma 5.3 combine to yield

Proposition 5.4. *In two-action signaling games, the opacity design payoff is convex in the prior, μ .*

Not only is opacity design as good as commitment in binary-action games, it also guarantees that *ex ante* learning can only benefit the receiver, thus revealing a second avenue through which opacity design is useful. Furthermore, just as Proposition 3.3 can be extended to allow for more than one sender, so can Proposition 5.4. *Viz.*, in binary-action multi-sender signaling games, the value of (public) information is always positive for the receiver, provided the receiver can commit to an information structure. Here is an easy argument: at initial belief μ_0 , let $\pi^*: M \rightarrow \Delta\{a_1, a_2\}$ be an optimal commitment strategy for the receiver and let V^* be the receiver’s payoff from such a strategy. Following any realization of the initial experiment it is obvious that receiver can choose the same π^* as it did at the initial belief. Thus, the receiver can do no worse than V^* in expectation (with respect to ζ), hence we have convexity. That opacity equals commitment completes the reasoning.

Even without opacity design, in two-action cheap-talk games *ex ante* information cannot hurt the receiver:

Proposition 5.5. *In two-action cheap-talk games, the value of information is always positive for the receiver*

This proposition is closely related to Proposition 3.3. It is easy to see that at any prior, there must be a receiver-optimal equilibrium in which at most two messages are sent—akin to the action recommendations of an optimal signal. Then, just as the opacity equals commitment result follows from the fact that disobedience cannot be interim optimal—since the receiver could always have just committed to take the same action in the first place—the value of *ex ante* information must be optimal in binary-action cheap-talk games even without opacity design. At any new prior for the game, either the same strategy vector for the sender is part of the receiver optimal equilibrium—the receiver is still willing to follow the “recommendations”—or it isn’t. But if it isn’t, the receiver strictly prefers to choose the same action after each message instead of following the “recommendations.” An equilibrium that yields this payoff always exists—for instance, the babbling equilibrium—and so the value of initial learning must be positive.

Beyond cheap talk, it is easy to identify other (strong) sufficient conditions that guarantee a positive value of information in signaling games without opacity design. Indeed, any conditions that engender a fully separating equilibrium ensure a positive value of information, since in such equilibria the receiver’s posterior beliefs do not depend on her prior. In the online appendix, we specify further nontrivial sufficient conditions.

5.1 Related Analyses of the Value of Information

The closest work to this portion of the paper is [Chen \(2012\)](#), who shows that in a two state quadratic cheap talk setting, the value of information may be negative for the receiver.¹⁴ Also similar is [de Barreda \(2010\)](#), who shows that the value of *private* information for the receiver in the uniform-quadratic setting of [Crawford and Sobel \(1982\)](#) may be negative.

¹⁴Note that her result does not contradict Proposition 5.5 below—in her setting the sender, himself, is only partially informed about the state and thus learns from the public signal as well.

Weksler and Zik (2021) introduce testing into a signaling environment, wherein a receiver chooses a test (whose realization is private) which the sender observes before sending a costly signal. Notably, the receiver may prefer less informative tests. The rationale is similar to the results on the potential negative value of information of this paper—the additional information provided by the test may lead to a worse equilibrium. In contrast, Lichtig and Weksler (2021) show that increased information for the *sender* benefits the receiver in a general class of voluntary disclosure games.

There are also a number of other papers, including Kamien et al. (1990); Gossner and Mertens (2001); Bassan et al. (2003); Lehrer and Rosenberg (2006); Pęski (2008); Meyer et al. (2010); and Ui and Yoshizawa (2015), that look at the value of information in games. However, none of these explore the value of information in the same sense as this paper. Here, we are interested in something much closer to a decision problem for the receiver, with the *caveat* that the information generation process is endogenous and generated by equilibrium play of the sender. Outside of that there are no strategic concerns, and the sender is perfectly informed, so there is no learning on his part. Consequently, this work's (and Chen's) analysis is more similar in spirit to e.g. Blackwell (1951, 1953); and Ramsey (1990), who explore the value of information in decision problems.

6 Discussion

Opacity design, in two-action signaling games, enables the receiver to achieve the same payoff as in the case wherein she had the power of committing *ex ante* to actions conditioned on the signal choice of the sender. This is both unexpected, that for a broad class of games information design is as powerful as commitment; and useful, since the receiver's commitment problem is merely a linear program.

Also surprising is that even in games with just two actions, *ex ante* learning may hurt the receiver, provided she cannot design the opacity in the game. The resulting drop in (useful) information provided in equilibrium at the new priors may be so great so as to overpower the initial gain in information. On the other hand, if she may design the opacity, the value of information is always positive—the receiver always benefits from

information, no matter the form it takes.

One final comment: there is recent experimental evidence that senders respond to different levels of opacity, which advocates for the usefulness of this paper’s results. Specifically, [Blume et al. \(2023\)](#) find that senders do respond as predicted to the introduction of noise in a communication setting. Relatedly, the randomized response technique of [Warner \(1965\)](#), in which a survey respondent’s answers are garbled (ostensibly to preserve anonymity), is in similar spirit to the information design in signaling games proposed by this work. The success of that technique—refer, e.g., to [Blume et al. \(2019\)](#), who show experimentally that randomized response can induce senders to be significantly more truthful—provides suggestive evidence that the information design we studied may be useful in practice. We leave this for future work.

A Omitted Proofs

A.1 Remark 2.2 Proof

Proof. Consider any equilibrium triple (π, σ, ρ) . Now introduce for each signal realization x_e and action a_l in the support of $\rho(\cdot|x_e)$ two new mappings:

1. $\hat{\pi}: M \rightarrow \Delta(\hat{A})$, where $\hat{\pi}(a_l^e|m_j) := \pi(x_e|m_j)\rho(a_l|x_e)$; and
2. $\hat{\rho}$, where $\hat{\rho}(a_l|a_l^e) = 1$ and $\hat{\rho}(a_l|a_{l'}^e) = 0$ for all $l' \neq l$.

That is, a_l^e is the instruction to play a_l that induces the same belief as x_e . Clearly set \hat{A} may be larger than A —it may have multiple “duplicate” recommendations.

By construction, for each a_l^e , action a_l is a best response. Moreover, both the obedience and IC constraints are satisfied, and the expected payoff for the receiver is the same. Finally, introduce garbling $g: \hat{A} \rightarrow \Delta(A)$, where $g(a_l|a_{l'}^e) = 1$ for $l = l'$ and $g(a_l|a_{l'}^e) = 0$ for $l \neq l'$.

Define $\pi' := g \circ \hat{\pi}$. It is easy to see that the IC constraints remain satisfied (since each message will lead to the same distribution of actions chosen by the receiver). Moreover, the obedience constraints must be satisfied as well since π' is less (Blackwell) informative than $\hat{\pi}$. Hence (π', σ', ρ') is also an equilibrium, where $\rho'(a_l|a_{l'}) = 1$ for $l = l'$, $\rho'(a_l|a_{l'}) = 0$

for $l \neq l'$, and $\sigma' = \sigma$; and the receiver's expected payoff remains the same. ■

A.2 Remark 4.2 Proof

Proof. The receiver solves

$$\max_{p, \hat{p}, q, \hat{q}} \left\{ -\mu \left(\hat{p} (1 - \hat{a})^2 + 1 - p - \hat{p} \right) - (1 - \mu) (q + \hat{q} \hat{a}^2) \right\},$$

subject to

$$-pb^2 - \hat{p}(\hat{a} - 1 - b)^2 - (1 - p - \hat{p})(1 + b)^2 \geq -qb^2 - \hat{q}(\hat{a} - 1 - b)^2 - (1 - q - \hat{q})(1 + b)^2 - c, \quad (IC_1)$$

and

$$-q(1 - b)^2 - \hat{q}(\hat{a} - b)^2 - (1 - q - \hat{q})b^2 \geq -p(1 - b)^2 - \hat{p}(\hat{a} - b)^2 - (1 - p - \hat{p})b^2 - c. \quad (IC_0)$$

It is straightforward to verify that IC_1 does not bind. On the other hand, IC_0 must hold with equality at the optimum. Furthermore, we must have $p + \hat{p} \leq 1$ (in addition to $p, \hat{p}, q, \hat{q} \in [0, 1]$ and $q + \hat{q} \leq 1$). Let us first guess that $p + \hat{p} = 1$. We can substitute in for \hat{p} into IC_0 , and rearrange to get p as a function of q and \hat{q} :

$$p = \frac{(2b - 1)q + 2\hat{a}b(\hat{q} - 1) - \hat{a}^2(\hat{q} - 1) + c}{(\hat{a} - 1)(-2b + \hat{a} + 1)}.$$

As $0 \leq p \leq 1$, this becomes

$$\frac{2bq - q + c - 2b + 1}{\hat{a}(\hat{a} - 2b)} \leq \hat{q} \leq \frac{2bq - q + c - 2\hat{a}b + \hat{a}^2}{\hat{a}(\hat{a} - 2b)}.$$

We also substitute for p and \hat{p} into the objective to obtain a two-variable linear program, which we solve for q and \hat{q} . All in all, we obtain $q = \hat{p} = 0$, $p = 1$, and

$$\hat{q} = \frac{2b - c - 1}{\hat{a}(2b - \hat{a})},$$

which produces a payoff for the receiver of

$$-\frac{(1-\mu)\hat{a}(2b-1-c)}{2b-\hat{a}}. \quad (A_1)$$

It is also easy to check that at the optimum we cannot have $q + \hat{q} = 1$. This leaves only—as there exists a solution to a linear program at a vertex of the constraint polytope—three of p, \hat{p}, q , and \hat{q} equaling 0; and the fourth pinned down by those and the binding IC_0 as a solution. Checking them one by one, we see that only $q = \hat{q} = \hat{p} = 0$ and

$$p = \frac{c}{2b-1},$$

is ever optimal. It produces a payoff of

$$-\frac{(2b-1-c)\mu}{2b-1}.$$

Comparing it to Expression A_1 yields the first part of the result. Finally, if the solution is not unique, then there must be multiple optimal solutions at vertices of the constraint polytope. However, this is only true if the parameters satisfy

$$\frac{\mu}{1-\mu} = \frac{(2b-1)\hat{a}}{2b-\hat{a}},$$

which is a knife-edge case. ■

A.3 Proposition 4.3 Proof

Proof. It is clear that Constraint IC_1 must be slack. Therefore, by the concavity of $-a^2$ and $-(a-b)^2$, it is optimal for Q to be degenerate on some action $y \in [0, 1]$. Now let us define v as the expectation of P . Given this, the (necessarily) binding Constraint IC_0 can be rewritten as

$$\int_0^1 a^2 dP(a) = 2bv - 2by - c + y^2.$$

Substituting this into the objective and reducing (and eliminating constants), we obtain the following maximization problem

$$\max_{y,v \in [0,1]} \left\{ 2\mu v(1-b) + 2b\mu y - y^2 \right\},$$

subject to

$$v^2 \leq 2bv - 2by - c + y^2 \leq v, \quad (A2)$$

where this chain of inequalities follows from the bounds on the possible variance of P . Ignoring this constraint, we obtain that $v = 1$ and $y = b\mu$. However, this solution always violates one of the inequalities except in the knife-edge case where $c = 2b - 1 - b^2\mu(2 - \mu)$. The result follows immediately. ■

A.4 Corollary 4.4 Proof

Proof. If both P and Q are deterministic, then opacity = commitment if and only if $v = y$. Moreover, we can substitute this into Inequality A2, which yields $v^2 - c \geq v^2$, a contradiction. If P is not deterministic, the result is obvious. ■

A.5 Proposition 4.5 Proof

Proof. A direct computation reveals that

$$\mathcal{L}(x) = x(1-x) + \lambda \left(-c + \frac{(\mu-x)(x-b)^2}{(1-\mu)\mu} \right),$$

is concave on $[0, 1]$ if λ is sufficiently low or else is convex then concave. Any line tangent to $\mathcal{L}(x)$ at a point $a \in [0, 1]$ has equation

$$t(x) := \left[\lambda \left(\frac{2(\mu-a)(a-b)}{(1-\mu)\mu} - \frac{(a-b)^2}{(1-\mu)\mu} \right) - 2a + 1 \right] x - \frac{2\lambda a^3 - ((\mu+2b)\lambda + (\mu-1)\mu)a^2 + (c(\mu-1) + b^2)\mu\lambda}{(\mu-1)\mu}.$$

The equation $t(1) - \mathcal{L}(1) = 0$ rearranges to

$$a = a(\lambda) = \frac{(\mu + 2b - 1)\lambda + (\mu - 1)\mu}{2\lambda}.$$

This function is strictly increasing in λ and ranges from $-\infty$ to $\frac{\mu+2b-1}{2}$. Now, let us look for the line that intersects the two points $(\mu, 0)$ and $\left(1, -c + \frac{(\mu-1)(1-b)^2}{(1-\mu)\mu}\right)$. It is

$$-\frac{(c\mu + b^2 - 2b + 1)(\mu - x)}{(\mu - 1)\mu},$$

and it intersects the curve

$$\tau(x) := -c + \frac{(\mu - x)(x - b)^2}{(1 - \mu)\mu},$$

at the point

$$\underline{x} := \frac{-\sqrt{4b^2 + (-4\mu - 4)b + \mu^2 + (4c + 2)\mu + 1} + 2b + \mu - 1}{2}.$$

\underline{x} is strictly decreasing in c and equals 0 when $c = 2b - 1$. Accordingly $\underline{x} > 0$. Moreover, it is clear that $\underline{x} < \frac{\mu+2b-1}{2}$. Thus, by the intermediate value theorem, there exists a $\lambda > 0$ such that $a = \underline{x}$.

We have established the result: the support of the optimal distribution is $\{\underline{x}, 1\}$, which is given by the convexification of the objective for the receiver. In this objective, the Lagrange multiplier, λ , is an endogenous object pinned down by the binding constraint for the sender in state 0 (necessarily, $\mathbb{E}_F[\tau(x)] = 0$).¹⁵ Figure 3 illustrates the Lagrangian (solid black), its convexification (thick orange), the curve $\tau(x)$ (solid blue), and the splitting corresponding to the optimal F (dotted red). ■

A.6 Proposition 5.2 Proof and Payoff Function Derivation

We derive the receiver's payoff as a function of the belief μ through a pair of claims. First,

Claim A.1. *Without opacity design, for any prior $\mu > 13/36$, there exists no equilibrium in which a message is played that induces a belief such that action a_1 is strictly optimal.*

¹⁵This is why we solved for the line that goes through the point $(\mu, 0)$.




Figure 3: The Opacity Design Problem and its Solution

Proof. We can exhaustively proceed through each message:

1. Suppose a_1 is strictly optimal following m_1 . Then, θ_4 must have support of his mixed strategy on m_1 . That gives him a payoff of -1 , so he can deviate profitably to m_2 .
2. Suppose a_1 is strictly optimal following m_2 . Both θ_3 and θ_4 must have support of their mixed strategies on m_2 . Moreover, so much of their support must be on m_2 that the receiver must choose a_2 following m_1 (which will always be chosen by θ_1). Hence, θ_3 can deviate profitably to m_1 .
3. Suppose a_1 is strictly optimal following m_3 . Consequently, θ_3 must have some support of his mixed strategy on m_3 (since θ_4 will never choose m_3). Moreover, θ_3 cannot be choosing a pure strategy since otherwise he would have a profitable deviation to m_2 . Hence, θ_3 must be mixing over m_3 and m_2 (since 3 is strictly larger than either of θ_3 's payoffs for m_1). The receiver must also mix following m_2 , so as to leave θ_3 willing to mix. In particular, the receiver must choose a_2 with probability $2/3$ following m_2 and a_1 with probability $1/3$. But note that θ_4 must also have support on m_2 , which message would thus yield it a payoff of $2/3$, which is less than $5/4$, the payoff he would get from deviating profitably to m_1 . ■

As a result $V^T = \frac{1}{3} + \mu$ for all $\mu > 13/36$. Second,

Claim A.2. For any $\mu \leq 13/36$, the payoff without opacity design is $37/24 - 2\mu$.

Proof. First, an equilibrium that begets such a payoff exists. In state θ_1 the sender chooses m_1 , in states θ_2 and θ_4 he chooses m_2 , and in state θ_3 he chooses m_3 . Upon observing m_1 , the receiver chooses a_2 , and upon observing either m_2 or m_3 the receiver chooses a_1 .

It is immediately evident that neither θ_1 nor θ_2 have profitable deviations, since they are choosing strictly dominant strategies. On path, θ_4 obtains 2, whereas his payoff from deviating is less than 2. Finally, θ_3 obtains 3 following m_3 and less than 3 following any other message. The receiver's equilibrium payoff is

$$\frac{1}{3} + \frac{1}{8} + 2\left(\frac{13}{24} - \mu\right) = \frac{37}{24} - 2\mu.$$

This payoff is the opacity design payoff and so must be maximal. ■

By Proposition 3.3, to obtain V we need only solve the commitment problem. Thus, its derivation reduces to a simple (though tedious) linear program, which is omitted.

A.7 Lemma 5.3 Proof

Proof. Recall that from Lemma 3.2 it is without loss of generality in the receiver's commitment problem to restrict the sender to pure strategies. Hence, the receiver solves

$$\max_{\pi, \mathbf{s}} \left\{ \sum_{i=1}^n \mu(\theta_i) \sum_{l=1}^k \pi(a_l | m_i) u_R(m_i, a_l, \theta_i) \right\},$$

such that

$$\sum_{l=1}^k \pi(a_l | m_i) u_S(m_i, a_l, \theta_i) \geq \sum_{l=1}^k \pi(a_l | m'_i) u_S(m'_i, a_l, \theta_i),$$

for all θ_i, m'_i , where \mathbf{s} is a vector of pure strategies chosen by the sender and message m_i is the message chosen in state θ_i (of course it is possible that $m_i = m_k$ for $k \neq i$ if the sender chooses the same message in states θ_i and θ_k). For a fixed vector of pure strategies this is a linear program.

Naturally, the payoffs of the game may be such that for some vector of pure strategies, there exists no signal that is incentive compatible. Define set F as the set of all pairs of signals and strategy vectors that are incentive compatible, with element $f := (\pi, \mathbf{s})$. For any $f \in F$, we have

$$V_f(\mu) = \sum_{i=1}^n \mu(\theta_i) \sum_{l=1}^k \pi(a_l | m_i) u_R(m_i, a_l, \theta_i),$$

which is linear, and hence convex, in μ .

Furthermore, by definition, $V(\mu) = \max_f V_f(\mu)$. Since V_f is convex, $\text{epi } V_f$ is also convex. Then, $\text{epi } V = \text{epi } \max_f V_f = \cap_{f \in F} \text{epi } V_f$. Since the intersection of convex sets is also convex, $\text{epi } V$ is also convex, hence V is convex. ■

A.8 Proposition 5.5 Proof

Proof. First, we observe that it is WLOG to impose that at most two messages are sent in the receiver-optimal equilibrium. Take any receiver-optimal equilibrium in which more than two messages are sent and both actions are taken with probability 1 following at least one message. Let messages m_1 and m_2 be two such on-path messages, each of which is followed by a different action (a_1 and a_2 , respectively). Suppose only those two messages are used by the sender (m_1 if the sender is in a state where previously he had sent a message that was followed by a_1 and m_2 otherwise), followed by the two actions, respectively. This leaves the sender's incentives and (useful) information content of the messages unchanged. Consequently, this is also an equilibrium and yields the same payoffs to both players.

Given this, the initial experiment either results in a posterior at which the receiver is willing to take action a_1 after m_1 and a_2 after m_2 , or not. In the former case, the result is immediate: the incentives for the sender are unchanged so it is as if (from the receiver's perspective) the information was exogenous. In the latter case, the proof is not much more difficult. Suppose at the new belief the receiver is unwilling to take action a_1 following m_1 . This means that the receiver prefers to always take a_2 (after each message) than the status quo. However, there is always a babbling equilibrium in cheap-talk games, so at this new belief, there is an equilibrium in which the receiver takes a single action, which must yield the receiver a weakly higher payoff than under the prior. ■

References

- S Nageeb Ali and J Aislinn Bohren. Should straw polls be banned? *Games and Economic Behavior*, 118:284–294, 2019.
- Vladimir Asriyan, William Fuchs, and Brett Green. Aggregation and design of information in asset markets with adverse selection. *Journal of Economic Theory*, 2021.

- Ian Ball. Scoring strategic agents. *Forthcoming, American Economic Journal: Microeconomics*, sep 2019.
- Bruno Bassan, Olivier Gossner, Marco Scarsini, and Shmuel Zamir. Positive value of information in games. *International Journal of Game Theory*, 32(1):17–31, Dec 2003.
- Marco Battaglini. Public protests and policy making. *The Quarterly Journal of Economics*, 132(1):485–549, 2017.
- David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, Berkeley, Calif., 1951. University of California Press.
- David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.
- Andreas Blume, Oliver J. Board, and Kohei Kawamura. Noisy talk. *Theoretical Economics*, 2(4):395–440, 2007.
- Andreas Blume, Ernest K. Lai, and Wooyoung Lim. Eliciting private information with noise: The case of randomized response. *Games and Economic Behavior*, 113:356 – 380, 2019.
- Andreas Blume, Ernest K. Lai, and Wooyoung Lim. Mediated talk: An experiment. *Journal of Economic Theory*, 2023.
- Raphael Boleslavsky and Kyungmin Kim. Bayesian persuasion and moral hazard. *Mimeo*, February 2017.
- Ying Chen. Value of public information in sender–receiver games. *Economics Letters*, 114(3):343 – 345, 2012.
- In-Koo Cho and David M. Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–221, 1987.
- Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.

- Inés Moreno de Barreda. Cheap talk with two-sided private information. *Mimeo*, nov 2010.
- Laura Doval and Vasiliki Skreta. Constrained information design. *Mathematics of Operations Research*, 2023.
- Timothy Feddersen and Ronen Gradwohl. Persuasion, transparency, and commitment. *Mimeo*, 2015.
- Francoise Forges. An approach to communication equilibria. *Econometrica*, 54(6):1375–1385, 1986.
- Francoise Forges. Equilibria with communication in a job market example. *The Quarterly Journal of Economics*, 105(2):375–398, 1990.
- Alex Frankel and Navin Kartik. Improving information from manipulable data. *Journal of the European Economic Association*, 20(1):79–115, 2022.
- Esther Gal-Or. Warranties as a signal of quality. *The Canadian Journal of Economics*, 22(1): 50–61, 1989.
- Chirantan Ganguly and Indrajit Ray. Simple mediation in a cheap-talk game. *Mimeo*, 2011.
- George Georgiadis and Balász Szentes. Optimal monitoring design. *Econometrica*, 2020.
- Maria Goltsman, Johannes Hörner, Gregory Pavlov, and Francesco Squintani. Mediation, arbitration and negotiation. *Journal of Economic Theory*, 144(4):1397 – 1420, 2009.
- Olivier Gossner and Jean-François Mertens. The value of information in zero-sum games. *Mimeo*, July 2001.
- Ronen Gradwohl and Timothy Feddersen. Persuasion and transparency. *The Journal of Politics*, 80(3):903–915, 2018.
- Maxim Ivanov. Communication via a strategic mediator. *Journal of Economic Theory*, 145 (2):869 – 884, 2010.

- Morton I Kamien, Yair Tauman, and Shmuel Zamir. On the value of information in a strategic conflict. *Games and Economic Behavior*, 2(2):129 – 153, 1990.
- Deniz Kattwinkel and Alexander Winter. Optimal elicitation of information. *Mimeo*, 2021.
- Andrew Kephart and Vincent Conitzer. The revelation principle for mechanism design with reporting costs. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, page 85–102, 2016.
- Ehud Lehrer and Dinah Rosenberg. What restrictions do bayesian games impose on the value of information? *Journal of Mathematical Economics*, 42(3):343 – 357, 2006.
- Avi Lichtig and Ran Weksler. Information in evidence games. *Mimeo*, 2021.
- Bartosz Mackowiak, Filip Matejka, and Mirko Wiederhold. Rational inattention: A review. *Forthcoming, Journal of Economic Literature*, 2022.
- Bernard De Meyer, Ehud Lehrer, and Dinah Rosenberg. Evaluating information in zero-sum games with incomplete information on both sides. *Mathematics of Operations Research*, 35(4):851–863, 2010.
- Paul Milgrom and John Roberts. Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821, 1986.
- Roger B. Myerson. Multistage games with communication. *Econometrica*, 54(2):323–358, 1986.
- Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- Phillip Nelson. Advertising as information. *Journal of Political Economy*, 82(4):729–754, 1974.
- Marcin Peński. Comparison of information structures in zero-sum games. *Games and Economic Behavior*, 62(2):732 – 735, 2008.

- F. P. Ramsey. Weight or the value of knowledge. *The British Journal for the Philosophy of Science*, 41(1):1–4, 1990.
- Armin Rick. The benefits of miscommunication. *Mimeo*, November 2013.
- Andrés Salamanca. The value of mediated communication. *Journal of Economic Theory*, 2021.
- Michael Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 1973.
- Takashi Ui and Yasunori Yoshizawa. Characterizing social value of information. *Journal of Economic Theory*, 158:507 – 535, 2015. Symposium on Information, Coordination, and Market Frictions.
- Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.
- Ran Weksler and Boaz Zik. Informative tests in signaling environments. *Forthcoming, Theoretical Economics*, 2021.
- Asher Wolinsky. Eliciting information from multiple experts. *Games and Economic Behavior*, 41(1):141–160, 2002.