Information Design in Cheap Talk

Qianjun Lyu*

First version: May 25, 2019

January 14, 2021

Abstract

An uninformed sender chooses a publicly observable experiment and sends a message to a receiver after privately learning the experimental outcome. To design the optimal experiment, the sender faces a tension between acquiring more information and alleviating the conflict of interest. With binary state space, the optimal experiment generates a conclusive signal about the state in which the two parties' interests coincide. When the choice of experiment is not publicly observable and the sender cannot commit to it, an informative equilibrium exists if and only if there is an equilibrium where the sender chooses to become perfectly informed.

Keywords: Cheap talk, information design, full communication, conclusive good news

JEL: D82, D83

*Faculty of Business and Economics, University of Hong Kong, lyu2014@hku.hk. I would like to thank my supervisor Wing Suen for his guidance. I would also like to thank Ricardo Alonso, Ettore Damiano, Jin Li, Matthew Mitchell, Xiaosheng Mu, Xianwen Shi, and Jianrong Tian for helpful discussions. All errors are mine.

1 Introduction

Starting from Crawford and Sobel (1982), many papers analyze how a biased sender (he) can gain from strategic communication with an uninformed receiver (she). Much of this literature assumes that the sender has perfect information about the state of the world; equivalently, an uninformed sender is presumed to conduct a fully-revealing experiment before transmitting information to the receiver. Would it make a difference if the sender can choose what to learn about the payoff-relevant state? If so, what would be the optimal experiment and how would the choice of experiment affect strategic information transmission? To what extent can the receiver benefit from the conveyed information given that the sender has full control over information acquisition? Specifically, we focus on arbitrary state-dependent preferences for both sender and receiver. We assume that, while the choice of experiment is publicly observable, the outcome of the experiment is privately learned by the sender, who sends a cheap talk message to the receiver, who then takes an action. Therefore, one can also interpret this game within the framework of Bayesian persuasion (Kamenica and Gentzkow (2011)), where we relax the sender's commitment to revealing the true information outcomes.

Communication loss is commonly incurred in cheap talk games where the sender's private information is exogenous. This is because the sender's bias incentivizes him to misreport his private information, resulting in the receiver rationally discounting his message. However we can recover the credibility of truthful report (full communication) if the sender can endogenously choose how much information to acquire. In particular, instead of acquiring more information and then coarsening it at communication, the sender can simply acquire the same amount of coarsened information and report it truthfully.[1]

Among all the experiments that are credible for full communication, which one is the best for the sender? Lipnowski and Ravid (2020) provides an elegant analysis of this problem when the sender's preferences depend solely on the receiver's action and are not affected by the state (the case of *transparent motives*), using the quasi-concave closure to characterize the highest achievable payoff for the sender.[2] Their analysis, however, does not generalize to situations where the sender's preferences over receiver's action depend on the state of the world—the case of *state-dependent preferences* being considered in this paper. Suppose, for example, a financial analyst (sender) prefers that his client (receiver) allocates all her resources to a risky asset if it is indeed

[1] Pei (2015), Deimen and Szalay (2019) also show that full communication can be generated, but they make restrictions on the payoff functions and technologies of information acquisition.

[2] Despite that Lipnowski and Ravid (2020) actually consider cheap talk where the sender privately knows the true state and they interpret the sender's reporting policy as the information structure. With transparent motives, their analysis is applicable to cases where an uninformed sender commits to an experiment and sends a message to a receiver after privately learning the experimental outcome.

profitable (state 0), but only allocate a small fraction of resources if the risky asset is unprofitable (state 1). The first role of information structure in this context is to provide the sender with valuable information about the payoff-relevant state, so that he can make more desirable recommendations. Furthermore, the degree of alignment over the two parties' interests may vary by state. In state 0, analyst and investor's interests are aligned in that they both want to direct all resources to the risky asset. In state 1, the analyst still prefers investing a non-zero fraction of resources into the risky asset for his own bias (e.g., contingent commissions), while the investor prefers zero investment. Therefore, the second role of information structure is to control the extent to which the sender and receiver's interests are aligned. While selecting the optimal choice of information structure, the sender faces a tension between acquiring more information and alleviating the conflict of interest.

In the main body of this paper, we focus on a specific group of preferences, *one-sided common interest*, such that the sender and the receiver share the same ideal action in one state (state 0). This type of preferences applies to situations where the sender—the consultant, salesperson, or project manager—is tempted to inflate the perceived value of a product or service. Therefore when the true valuation is indeed high, both parties prefer the same action. Except for this common interest, we allow for arbitrary preferences and action space. With binary state space, we show that it is optimal for the sender to design an experiment that generates a conclusive signal about the state where both parties' interests coincide. In the analyst-investor example, an experiment that fully reveals state 0 would generate a bad signal in state 1 with probability 1, and would generate a good signal or a bad signal in state 0. This means that a good signal can only be generated in state 0 with probability less than 1.[3] In other words, in spite of choosing an optimal experiment, and even if the risky asset is indeed profitable, there is still a positive chance that the investor is recommended *not* to direct all resources to the good risky asset.

This seemingly counter-intuitive result is driven by two forces: (1) a more informative experiment would allow the sender to make better use of the information; and (2) an experiment that fully reveals the common-interest state would align the two parties' interests ex post. More specifically, when the sender considers acquiring more information, he has to foresee the degree of conflict conditional on all possible future beliefs. For example, acquiring perfect information about the risky asset is suboptimal. Because with positive probability state 1 could be realized, and the conflict of interest in state 1 would restrain the sender from effectively using his information, e.g., he cannot persuade the receiver to take his ideal action when state 1 is fully revealed.[4]

[3]This holds as long as the optimal experiment does not fully reveal the bad state, which is usually true.

[4]See the example in Section 2 for further elaboration.

Building on this result of conclusive good news, it is easy to pin down the optimal experiment using a geometric approach. In the last part of this paper, we also discuss a model where the sender has no commitment power in the selection of experiments, which can be interpreted as the sender having full control over manipulating the probability of generating different signal realizations. We show that focusing on a fully-revealing experiment is without loss of generality; i.e., an informative equilibrium exists if only if there exists an informative equilibrium with a perfectly informed sender. This rationalizes the assumption in the literature on cheap talk, where the sender typically has perfect information about the true state.

Related Literature

There is a growing literature exploring information acquisition in strategic communication. Most of these papers assume specific utility functions and make restrictions on information structure. Argenziano et al. (2016) discuss strategic communication with costly information acquisition, where the information structure that the sender chooses is restricted to different numbers of binary trials. Deimen and Szalay (2019) assume the state space is two-dimensional. They find that when the sender can commit to a normal information structure, it is optimal for him to acquire information of equal use to both himself and the receiver, which endogenously aligns their interests. Kosenko (2018) and Strulovici (2017) discuss mediator problems in which, after the sender acquires information, a mediator with different preferences can garble the information. Krähmer (2020) studies arbitrary information disclosure from the receiver to the sender by introducing the idea of randomizing over information structures.

This paper is in line with the approach of Bayesian persuasion (Kamenica and Gentzkow (2011), Dworczak and Martini (2019)) and the recent literature that relaxes the full commitment assumption. Guo and Shmaya (2020) examine the situation where the sender can manipulate the information structure at a cost which is related to the nature of the distortion. Di Tillio et al. (2017) investigate the persuasion outcome if the researcher is able to manipulate the experiment with private information, such as in deliberate assignment of treatment and control groups. Lipnowski et al. (2019) discuss the situation where the sender has private type on ability to fake information outcomes, concluding that the receiver could benefit from a less credible sender in terms of attaining more precise information. Tan and Nguyen (2021) consider Bayesian persuasion where the sender publicly designs an experiment and reports a message to Receiver at a cost which depends on the experimental outcomes. Alonso and Camara (2018) study information design problems where information tempering is detectable by the receiver, and how the receiver's audit probability affects the sender's choice of experiment. In contrast to these papers, my paper discusses two extreme cases: (1) full commitment on information acquisition and no commitment on communication; and

(2) no commitment on information acquisition and no commitment on communication.

The paper closest to this one is Lipnowski and Ravid (2020), which considers information design when there is no commitment on communication for the case where sender has transparent motives. With transparent motives, the sender's preferences are the same regardless of his private information. Hence, if he strictly prefers one message under some belief, he also prefers the same message across all other beliefs. Their model therefore requires the sender to be indifferent among different messages. However, with state-dependent payoff functions, the sender strategically chooses an experiment to shape his future preferences, and hence it is possible to have non-binding incentive constraints in equilibrium. In Section 2, we provide a simple example to illustrate how state-dependent preferences influence communication and the optimal information structure.

This paper is also related to Pei (2015), which discusses a cheap talk game where the sender can acquire, at a cost, information that is unobserved by the receiver. He adopts linear quadratic utility and restricts the information structure to finite partitions. Besides, coarser information structure incurs smaller cost. The results show that, when the cost is sufficiently high, an upwardly biased sender conveys more precise information when recommending higher action. My paper establishes this result with costless information acquisition. The major difference is that Pei (2015) focuses on unobservable information structure, emphasizing the sender's trade-off of balancing the marginal benefit from a more informative information structure with the marginal cost. On the other hand, the purpose of my paper is to determine which information structure is better, conditional on the receiver's decision endogenously depending on sender's choice of experiment. In Section 6, we also discuss the case when information structure is unobservable to the receiver but we assume zero information cost.

The remainder of this paper is organized as follows. Section 2 describes a simple example. Section 3 discusses the model setup. Section 4 and 5 discuss the optimal experiment where the sender can commit to it. Section 6 characterizes the optimal experiment where the sender cannot commit to the experiment. Section 7 concludes.

2 An Example

To understand how state-dependent preferences influence communication, consider the analyst-investor example. The binary state is $\omega \in \{0, 1\}$, where state 0 (state 1) represents good (bad) state where the risky asset is profitable (unprofitable). Both analyst (sender) and investor (receiver) share a common prior $\mu_0(\omega = 1) = 0.5$. Here, the investor can choose from three different portfolios, $\{a_1, a_2, a_3\}$, where a_1 directs all her resources to the risky asset (full investment), a_2 allocates half of her resources to risky asset (partial investment), and a_3 represents zero risky investment. The analyst conducts a free experiment to acquire information about the risky asset's profitability, after which he sends a message to the investor, who then takes an action. For the investor, given her preference,[5] her decision rule can be reduced to a mapping from belief to action: if $\mu \leq 0.4$, she prefers full investment a_1; if $\mu \in (0.4, 0.8]$, she prefers partial investment a_2; and if $\mu > 0.8$, she optimally chooses zero investment a_3. In Figure 1, the horizontal axis represents both parties' beliefs (about the probability of the bad state occurring), and the action written below the solid line is investor's optimal choice conditional on her belief.

(1) Transparent motives

First, consider an analyst whose preferences are state-independent (transparent motives). He obtains the payoffs $(4, 3, 0)$ from actions $\{a_1, a_2, a_3\}$ (shown as the number written above the solid line in Figure 1). To shed some light on how information design affects communication and equilibrium outcome, consider two particular experiments, X: (0,0.8) and Y: (0.4,0.8), where each experiment is defined as a pair of posterior beliefs.[6]

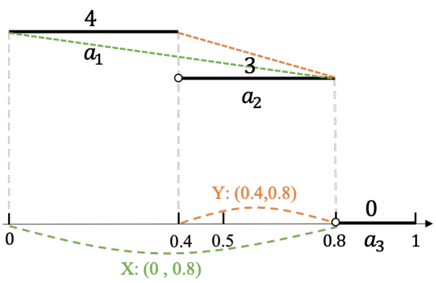

Figure 1: State-independent payoffs

If the analyst can commit to both information structure and truthfully reporting the information outcomes, the standard concavification argument suggests that Y is the optimal experiment, since the analyst can generate full investment (a_1) more often with Y than with X. However, in my setting where the sender's messages are cheap talk, neither of these experiments is incentive compatible for truthful report. With X, for example, the analyst would want to report a belief of 0 to induce the investor to choose full investment (a_1) regardless of the signal he obtains. However, the investor would

[5]For example, the investor's payoffs from actions $\{a_1, a_2, a_3\}$ are $(10, 8, 0)$ in state 0 and $(-5, -2, 0)$ in state 1. When relevant, the investor is assumed to break a tie in the sender's favor.

[6]Following the literature, an experiment can be defined by a distribution of posterior beliefs. With binary state space, the pair of posterior beliefs is enough to pin down an experiment. We can also define an experiment by the conditional distribution of signals over state space. Denote $\pi(s|\omega)$ as the probability that signal s is generated under state ω. In this example, we can equate the signal s to the posterior belief, and the signal distribution of X is: $\pi_X(0|0) = 0.75$, $\pi_X(0.8|0) = 0.25$, $\pi_X(0|1) = 0$ and $\pi_X(0.8|1) = 1$. The signal distribution of Y is: $\pi_Y(0.4|0) = 0.9$, $\pi_Y(0.8|0) = 0.1$, $\pi_Y(0.4|1) = 0.6$ and $\pi_Y(0.8|1) = 0.4$.

rationally infer from the analyst's temptation to misreport to ignore his message. Therefore, with transparent motives, if information transmission is cheap talk, the analyst cannot achieve a payoff higher than babbling.[7]

(2) State-dependent preferences

Now, consider an analyst whose preferences vary across states (shown in Figure 2). When the risky asset is indeed good (state 0), the analyst's preferences are the same as with transparent motives: the greater the risky investment, the better. When the risky asset is bad (state 1), zero risky investment is still least preferable to him, as it yields no commission. Yet, due to reputation concerns, he becomes more conservative on investment choice and prefers partial investment to full investment. His payoffs from different actions under different states are written above the solid lines.

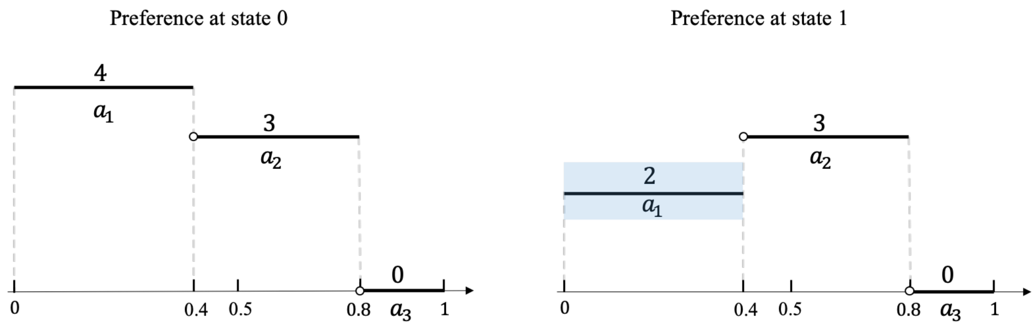

Figure 2: State-dependent payoffs

Building on these preferences, both experiments X and Y are incentive compatible for truthful report if the investor relies on the analyst's messages.[8] For example, when the analyst performs Y and obtains belief 0.8, he gets utility 3 from reporting the true information outcome, which is larger than the utility $0.8 \times 2 + 0.2 \times 4 = 2.4$ from mis-reporting the unrealized belief 0.4.

Note that when the state is good, the analyst benefits from getting full investment (a_1) more often. But when the state is bad, he prefers to have less frequent full investment. This is where the information value comes in. Performing a more informative experiment allows the sender to make better use of the information. In other words, he can make recommendations given a more precise belief. Recall that Y generates full investment more often than X under both states. Hence, the analyst prefers Y over X if the state is 0, while he prefers X over Y if the state is 1. In particular, conditional

[7] Lipnowski and Ravid (2020) show that the highest ex-ante payoff the sender can achieve is 3, which is the same as in the babbling equilibrium. However, they also show that there could be informative equilibrium. For example, if the receiver does not break a tie in sender's favor and takes a_2 given belief 0.4, then experiment Y is incentive compatible.

[8] As long as the sender's utility of a_1 conditional on state 1 is within [1.5,2.75], both experiments X and Y are incentive compatible for truth-telling.

on the state being 0, the sender's utility from X is $0.75 \times 4 + 0.25 \times 3 = 3.75$,[9] which is smaller than his utility from Y, $0.9 \times 4 + 0.1 \times 3 = 3.9$. In contrast, if the state is 1, the sender obtains utility 3 from X, which is larger than his utility $0.6 \times 2 + 0.4 \times 3 = 1.2$ from Y. In total, X outperforms Y. In fact, X turns out to be the optimal experiment in this example,[10] and it fully reveals state 0 with positive probability.

By performing X, the ex-ante probability of investor taking a_1 (0.375) is smaller than the ex-ante probability of good state (0.5). This seemingly counter-intuitive result for one thing is caused by the information value. Intuitively, by performing X, the analyst knows for sure that full investment is optimal if he obtains a belief of 0. In comparison, performing Y would yield a belief of 0.4, which would render the optimal recommendation of full investment to be inappropriate 40 percent of the time. For another thing, to ensure the optimality of conclusive good news (fully revealing state 0), there is one crucial assumption under this payoff structure: the full alignment of interests (*one-sided common interest*) in state 0 (i.e., the analyst and investor both prefer full investment when the risky asset is good). This ensures that, when the analyst indeed reveals state 0, that part of information can be effectively utilized since the investor's best response is also full investment. In other words, the experiment that fully reveals state 0 provides the analyst with both information value and a better alignment of interests. In contrast, fully revealing state 1 cannot be optimal: although it provides more information, the information value can be offset by the two parties' conflicting interests in state 1, since the investor's optimal choice in state 1 gives the worst payoff to the analyst.

3 Setting

This model describes a general cheap talk game where an uninformed sender (Sender) can first acquire information and then release information to a receiver (Receiver). In particular, we focus on situations where Sender has state-dependent preferences. The state space is finite with a typical element, $\omega \in \Omega$. Initially, Sender and Receiver share common prior belief $\mu_0 \in \Delta\Omega$. The game consists of two stages. In the first stage, Sender designs an information structure and conducts it with *no cost*. The information structure (or the experiment) $\{\pi(\cdot|\omega)\}_{\omega\in\Omega}$ is a distribution of signal realizations $s \in \mathcal{S}$ in each state. In the main part of this paper, π is assumed to be publicly observable, or equivalently, Sender can commit to π. In the second stage, Sender privately observes the signal realization and then sends a cheap talk message $m \in M$ to Receiver, after which Receiver chooses her action, $a \in A$. Note that Sender's reporting policy can be viewed as an information structure, $\{\delta(\cdot|s)\}_{s\in\mathcal{S}}$, which is a distribution of messages

[9] This is calculated given the signal distribution in footnote 5.

[10] X is the optimal experiment in this example since it is the most informative one that can credibly induce a_1 and a_2 in equilibrium.

for each signal realization. Therefore, a reporting policy can be regarded as a garbling matrix, which describes how the outcome of experiment π gets "garbled" in a way that is independent of the state (Blackwell and Girshick (1979), Börgers (2009)).

Let $u : \Delta A \times \Omega \to \mathbb{R}$ be Sender's utility and $u^r : A \times \Omega \to \mathbb{R}$ be that of Receiver.[11] After Sender chooses π and observes the signal realization, he chooses a reporting policy $\sigma : \Pi \times \mathcal{S} \to \Delta M$. Receiver then chooses a decision rule $\rho : \Pi \times M \to \Delta A$. The belief system is $\mu^r : \Pi \times M \to \Delta\Omega$. A perfect Bayesian equilibrium (PBE) consists of three maps:

1. $\sigma(\pi, s)$ is supported on $\arg\max_{m \in M} E_\omega[u(\rho(\pi, m), \omega)|\mu^s(\pi, s)]$, where $\mu^s(\pi, s)$ is Sender's posterior belief after observing signal s.

2. $\rho(\pi, m)$ is supported on $\arg\max_{a \in A} E_\omega[u^r(a, \omega)|\mu^r(\pi, m)]$.

3. $\mu^r(\pi, m)$ is obtained from μ_0, given π, σ, using Bayesian rule whenever possible.

To simplify the exposition, we omit the notation for π when it does not cause ambiguity. The belief system can be reduced to:

$$\mu^r(m) = \sum_{\mathcal{S}} \frac{\Pr(s)\delta(m|s)}{\sum_{\mathcal{S}} \Pr(s)\delta(m|s)} \mu^s(s) \tag{BC}$$

where $\Pr(s) = \sum_\Omega \pi(s|\omega)\mu_0(\omega)$ is the ex-ante probability that signal s is generated. Apparently, Receiver's posterior belief is a weighted average over Sender's posterior beliefs. We focus on Sender-preferred PBE,[12] and use π^* to denote the optimal information structure in this equilibrium. Usually, the message space is rich. It can contain recommendations of actions, and mixtures over actions or beliefs. Denote $v(\mu^s, \alpha(\mu^r))$ as Sender's expected payoff when his posterior belief is μ^s and Receiver's belief is μ^r.

$$v(\mu^s, \alpha(\mu^r)) := E_\omega[u(\alpha(\mu^r), \omega)|\mu^s]$$

$$\text{s.t.} \quad \alpha(\mu^r) \in \Delta(\arg\max_{a \in A} E_\omega[u^r(a, \omega)|\mu^r])$$

When Receiver is indifferent between two actions, she may need to employ a mixed strategy to ensure the existence of informative equilibrium (Lipnowski and Ravid, 2020). To gain tractability, let $\{\mu^r : \exists\ a \neq a' \text{ s.t. } a, a' \in \arg\max_{a \in A} E_\omega[u^r(a, \omega)|\mu^r]\}$ be finite. Note that when Sender's payoffs are state-dependent, the order of Sender's preferences over different actions may vary. Therefore when we say Receiver breaks indifference in Sender's favor, it means that she takes action $\hat{\alpha}(\mu)$, where

$$\hat{\alpha}(\mu) = \arg\max_{\alpha(\mu)} \ v(\mu, \alpha(\mu))$$

[11]The different domains for Sender and Receiver utility functions are solely to simplify later statements.

[12]In particular, Sender chooses the equilibrium that gives him the highest ex-ante payoff.

Denote $\hat{v}(\mu) := v(\mu, \hat{\alpha}(\mu))$ as Sender's expected payoff given that both parties' beliefs coincide and Receiver breaks ties in Sender's favor.[13] In strategic environment, though we are focusing on Sender-preferred PBE, assuming this particular tie-breaking rule is not without loss of generality: the restricted tie-breaking rule would shrink the set of incentive compatible experiments. Hence, to clarify, we state and prove our results by allowing arbitrary tie-breaking rule, and when we use $\hat{\alpha}$ instead of α, it is proved to be without loss.

Following the literature, we adopt a belief-based approach. An experiment π is equivalent to a distribution of posterior belief, $F \in \Delta(\Delta\Omega)$, and a reporting policy garbles the information structure F. Let $G \in \Delta(\Delta\Omega)$ be the distribution of Receiver's posterior belief. Hence, in equilibrium, G is a garbling of F. The belief-based approach allows us to focus on the equilibrium outcomes. Formally, an equilibrium outcome is a triplet, $(F, G, v) \in \Delta(\Delta\Omega) \times \Delta(\Delta\Omega) \times \mathbb{R}$. Claim 1 indicates that it can be narrowed down to (G, v),[14] as long as (G, v) is induced by an equilibrium.

Claim 1. *For any fixed σ and α, Sender's expected payoff can be directly pinned down by the distribution of Receiver's posterior belief G induced by (BC):*

$$\sum_{\mathcal{S}} \Pr(s) \sum_{M} \delta(m|s) v(\mu^s(s), \alpha(\mu^r(m))) = \sum_{supp(G)} \Pr(\mu^r) v(\mu^r, \alpha(\mu^r)) \tag{1}$$

where $\Pr(\mu^r)$ is the probability mass of Receiver's posterior belief. This result depends on the fact that Sender's expected payoff is linear in his belief for any fixed action of Receiver. A formal proof is given in the Appendix. Claim (1) implies that the sufficient statistic to calculate Sender's ex-ante payoff is the net information content that he delivers to Receiver, which is just the distribution of Receiver's posterior belief G. In other words, Sender can get the same ex-ante payoff by choosing a more informative experiment and meanwhile garbling it more, *as long as* the garbling (reporting) is incentive compatible and G remains the same.

Definition: If $F = G$ for some equilibrium, we call it a *full communication* (truth-telling) equilibrium. Otherwise, we call it a pooling equilibrium.

Lemma 1. *For any pooling equilibrium where $F \neq G$, there always exists a full communication equilibrium with $F' = G$ in which Sender obtains the same expected payoff.*

We obtain this Lemma by observing that experiment G is incentive compatible for truth-telling: In the original pooling equilibrium, each belief μ_j in the support of G is a weighted average over the set of beliefs $\{\mu_i\}$ in the support of F such that a Sender

[13] This is defined the same way as in Kamenica and Gentzkow (2011).

[14] Note that for (F, G, v), v is calculated by $v(\mu^s, \alpha(\mu^r))$. For (G, v), v is obtained by $v(\mu, \alpha(\mu))$.

of type μ_i would recommend μ_j with positive probability. Because preferences are linear in beliefs, if a type-μ_i Sender would weakly prefer to recommend μ_j, therefore a type-μ_j Sender would also weakly prefer to report μ_j. Thus, experiment G is credible for full communication. Given Claim 1, by performing experiment G and reporting the true information outcomes, Sender achieves the same ex-ante payoff.

Lemma 1 indicates the sufficiency to restrict our attention on full communication equilibrium that substantially reduce the complexity of equilibrium construction. One can see the crucial role that information design plays in cheap talk games: Sender endogenously designs an "efficient" experiment to align his interest with Receiver, so that he can always truthfully report the outcome he learns from the experiment. Communication loss is eliminated by the freedom of information acquisition.

Hence, Sender's optimization problem becomes the following.

$$\sup_{F,\alpha} \quad E_F[v(\mu, \alpha(\mu))]$$

$$\text{s.t.} \qquad \text{F is Bayes Plausible}^{15} \tag{BP}$$

$$v(\mu, \alpha(\mu)) \geq v(\mu, \alpha(\mu')), \quad \forall \mu, \mu' \in supp(F) \tag{IC}$$

We close this section by making some comments on full communication equilibrium. First, we can introduce a *lying cost* to obtain the uniqueness of full communication. Suppose the message is a recommendation of belief and Sender incurs a cost to report a belief different from the outcome he obtains from the experiment (Kartik (2009), Guo and Shmaya (2020)). Now if we think about a pooling equilibrium where Sender deviates from the true information outcomes but reports the belief that a rational Receiver would anticipate, then he suffers from the lying cost but his informational advantage turns out to be *not* valuable. Therefore, he would rather avoid knowing things that he does not want to reveal. Second, while discussing Bayesian persuasion with limited commitment, especially promotion game where Sender strictly prefers one action over the other (Guo and Shmaya, 2020), lying (miscalibration) cost is an intrinsic force to restrain Sender from misreporting. When the cost intensity varies, Sender can adjust his information structure (choosing messages that are more distant to each other) to reduce his incentive to miscalibrate and therefore approach "truth-telling". However, in our game, it is mainly the assumption of *state-dependent preferences* that grants us the existence of (strict) full communication equilibrium where Sender and Receiver's preferences are endogenously aligned by the experiment outcome.[16]

[15] We adopt the same definition as with Kamenica and Gentzkow (2011) such that the posterior can average back to the prior.

[16] By strict full communication equilibrium, we refer to the case where (IC) holds strictly.

4 One-sided Common Interest

In this section, we focus on binary state space $\omega \in \{0,1\}$ and let $\mu = \Pr(\omega = 1)$. Additionally, we discuss a particular group of preferences: Sender and Receiver have one-sided common interest in one state.

Definition: Sender has *one-sided common interest* with Receiver in state 0 if

$$v(0, \alpha(0)) \geq v(0, \alpha(\mu^r)), \ \forall \mu^r \in [0,1], \alpha(\cdot).$$

Without loss of generality, we focus on the case in which Sender and Receiver have one-sided common interest in state 0. The argument is symmetric if they have common interest in state 1. This definition says that, when Sender privately knows the true state is 0, he prefers Receiver to hold the same belief as his when making her decision, regardless of her tie-breaking rule. Essentially, if $\alpha(0)$ is unique, then when the true state is 0, Sender's ideal action is indeed Receiver's optimal decision. Note that when we assume one-sided common interest in state 0, we implicitly allow for a conflict of interests at any belief $\mu \neq 0$, (i.e., Sender may prefer Receiver's posterior to depart from his as long as he is not sure that the state is 0).

This assumption can broadly apply to situations where Sender is tempted to under-report (over-report) his private information, e.g., downward (upward) bias.[17] For example, a doctor may prefer to prescribe tests that are not necessary from the patient's point of view if the medical conditions are relatively mild, but both would prefer an invasive treatment when the patient's condition is severe. Furthermore, this assumption can also arise in problems with multi-dimensional attributes. For instance, the shareholder and manager of a corporation prefer the same project if one dominates the others in terms of both net present value and capital intensity; otherwise, the shareholder prefers the project with largest net present value and the manager prefers the one with highest capital intensity (Armstrong and Vickers, 2010). As another example, while discussing innovation activities in organization, if a R&D project is revealed to be good or bad, then there is no conflict of interest between the firm and agent conducting the project. However, when the state is unknown, the agent prefers to work on the project for a longer time (Guo, 2016).[18]

[17] Suppose the action space is $A = [0,1]$. If Receiver's ideal action is equal to the expected state, while Sender has quadratic loss utility and is biased towards an action lower than the expected state, then both Sender and Receiver would prefer action 0 when the true state is 0.

[18] We can also attain one-sided common interest in promotion games where Sender's preferences are state-independent, e.g., salesperson always prefers the consumer to buy, auto mechanic always prefers more expensive repair. However, since we do not discuss transparent motive in this paper, we can recover state-dependency by introducing competition in the market so that reputation matters. Alternatively, in the financial industry, a third party regulatory authority may exist to ensure the brokers make “suitable" recommendations to their customers (Inderst and Ottaviani, 2009).

For the main result of this paper, I make some regularity assumptions on Sender's preferences.

A 1. $\inf_\alpha v(\mu^s, \alpha(\mu^r))$ *and* $\sup_\alpha v(\mu^s, \alpha(\mu^r))$ *are weakly quasi-concave in* μ^r *for* $\mu^s \in [0,1]$.

A 1′. *If for some* μ^s *and* $\alpha(\cdot)$ *such that* $v(\mu^s, \alpha(\mu^r))$ *is a constant for all* $\mu^r \in [\mu', \mu'']$, *then* $\alpha(\mu^r)$ *is the same for all* $\mu^r \in [\mu', \mu'']$.

A 2. $\sup_\alpha v(\mu^s, \alpha(\mu^r))$ *is upper-semi continuous in* μ^r *and* $\inf_\alpha v(\mu^s, \alpha(\mu^r))$ *is lower-semi continuous in* μ^r. *Both have finitely many one-sided jumps.*

A1 and A1′ ensure that Sender's preference $v(\mu^s, \alpha(\cdot))$ is well-ordered on μ^r for all feasible tie-breaking rules. This is stronger than weak quasi-concavity, but weaker than strict quasi-concavity.[19] A2 is simply a regularity condition (Dworczak and Martini, 2019). Consider a linear quadratic model with $A = [0,1]$. Suppose Receiver's utility is $u^r(a, \omega) = -(\omega - a)^2$ and Sender's utility is $u(a, \omega) = -(\omega - a - b)^2$, where $b > 0$ means downward bias. All the assumptions are satisfied in this example.

Proposition 1. *(Conclusive good news) Suppose Sender and Receiver have one-sided common interest in state 0 and A1, A1′ and A2 are satisfied. If there exists an equilibrium in which Sender gets a payoff strictly higher than with babbling, then it is optimal to design an experiment that generates conclusive news about state 0:*

$$\pi^*(s = 0|\omega = 0) > 0, \ \pi^*(s = 1|\omega = 1) = 1$$

Proof. First note that it is without loss to consider binary signal structure. The reason is that if there exists a full communication equilibrium which induces $n > 2$ posterior beliefs, there are two beliefs in this set of n beliefs such that the convex combination of the payoffs from these two beliefs is weakly higher than the convex combination of the payoffs from the n beliefs. To prove the optimality of conclusive good news, we follow two steps. Step 1: If for some experiment $Y : (\mu_1, \mu_2)$, there exists a full communication equilibrium in which Sender gets strictly higher payoff than with babbling, we can construct a new experiment $X : (0, \mu_2)$ without violating Sender's incentive constraints. Step 2: We prove that X dominates Y. Hence, it is optimal to design an experiment that fully reveals state 0. Later on, we use π_X and π_Y if needed to call attention to the experiment's signal structure.

For Step 1, the following inequality shows that reporting μ_2 is incentive compatible when the outcome of experiment X is μ_2:

$$v(\mu_2, \alpha(\mu_2)) \geq (\mu_2, \alpha(\mu_1)) \geq v(\mu_2, \alpha(0)) \tag{2}$$

[19] Appendix B provides an analysis of the problem without assuming A1′. The main result still holds: an incentive compatible experiment that generates conclusive good news dominates any other incentive compatible experiment.

The first inequality in weak form follows from the incentive compatibility of the original experiment Y. However, if $v(\mu_2, \alpha(\mu_2)) = v(\mu_2, \alpha(\mu_1))$ and $v(\mu_2, \alpha(\mu))$ is a constant for all $\mu \in [\mu_1, \mu_2]$, A1′ implies that Receiver takes the same action when holding beliefs μ_1, μ_2, or $\mu_0 \in [\mu_1, \mu_2]$, and therefore experiment Y cannot be strictly more profitable than babbling. This shows that the first inequality of (2) is strict, or if it is weak, there exists a $\mu \in (\mu_1, \mu_2)$ such that $v(\mu_2, \alpha(\mu_2)) = v(\mu_2, \alpha(\mu_1)) < v(\mu_2, \alpha(\mu))$. The second inequality then follows from the first and from quasi-concavity of $v(\mu_2, \alpha(\cdot))$ (A1). Furthermore, one-sided common interest in state 0 implies

$$v(0, \alpha(0)) \geq v(0, \alpha(\mu_2)) \tag{3}$$

Equations (2) and (3) show that experiment X is incentive compatible.

Step 3 needs to show that X dominates Y in terms of Sender's ex-ante payoff. With slight abuse of notation, let y_k denote the signal in experiment π_Y and let x_j denote the signal in π_X. Note that, since $X : (0, \mu_2)$ is more informative than $Y : (\mu_1, \mu_2)$, we can then write down a garbling matrix B such that $\pi_X B = \pi_Y$. Matrix B describes how the outcome of experiment X is garbled to yield the outcome of experiment Y. In particular, each element b_{jk} represents the probability of generating y_k when obtaining x_j, and thus $\sum_k b_{jk} = 1$. Because Sender's posterior μ_2 does not change between obtaining x_2 and y_2, it must be the case that y_2 only contains information from x_2, and hence $b_{12} = 0$. Besides, $\mu(x_1) < \mu(y_1) < \mu(x_2)$ means that y_1 is generated by both x_1 and x_2 with strictly positive probability.

$$\begin{array}{cc} & \begin{array}{cc} y_1 & y_2 \end{array} \\ \begin{array}{c} x_1 \\ x_2 \end{array} & \begin{pmatrix} b_{11} = 1 & b_{12} = 0 \\ b_{21} > 0 & b_{22} > 0 \end{pmatrix} \end{array} \tag{B}$$

Denote $c_{jk} = b_{jk} \frac{\Pr(x_j)}{\Pr(y_k)}$. For each k, we have $\sum_j c_{jk} = 1$. Sender's posterior after observing y_k is a weighted average over the posteriors of x_j such that $b_{jk} > 0$, with the weight being c_{jk}.[20] Therefore, $\mu_1 = c_{11}(0) + c_{21}(\mu_2)$, and

$$\begin{aligned} v(\mu_1, \alpha(\mu_1)) &= c_{11} v(0, \alpha(\mu_1)) + c_{21} v(\mu_2, \alpha(\mu_1)) \\ &\leq c_{11} v(0, \alpha(0)) + c_{21} v(\mu_2, \alpha(\mu_2)) \end{aligned} \tag{4}$$

where the equality follows from the linearity of $v(\cdot, \alpha(\mu_1))$. The inequality comes from: (1) Sender and Receiver have one-sided common interest in state 0; and (2) reporting μ_2 rather than μ_1 is incentive compatible when μ_2 is observed under experiment Y. Furthermore, because $c_{12} = 0$ and $c_{22} = 1$, the following holds trivially:

$$v(\mu_2, \alpha(\mu_2)) = c_{12} v(0, \alpha(0)) + c_{22} v(\mu_2, \alpha(\mu_2)) \tag{5}$$

[20] For any state ω, we have $\Pr(\omega)\Pr(y_k|\omega) = \sum_j b_{jk} \Pr(x_j|\omega)\Pr(\omega)$. Dividing both sides by $\Pr(y_k)$ gives this result.

From (4) and (5), we have:

$$v(\mu(y_k), \alpha(\mu(y_k))) \leq \sum_j c_{jk} v(\mu(x_j), \alpha(\mu(x_j))), \;\; k = 1, 2$$

Multiplying both sides by $\Pr(y_k)$ and summing over k, we have that Sender receives weakly higher expected payoff from choosing X over Y.

$$\sum_k \Pr(y_k) v(\mu(y_k), \alpha(\mu(y_k))) \leq \sum_j \Pr(x_j) v(\mu(x_j), \alpha(\mu(x_j))) \qquad \square$$

In single-agent decision-making, the agent's value function is convex in his belief. Therefore, acquiring more information allows the agent to make better decisions. This is not generally true in strategic environments. The major issue is that Receiver's action could depend on Sender's choice of information structure. Therefore, a more informative experiment could make Sender worse off if Receiver takes an action unfavorable to Sender when more information is revealed. Hence, while determining the optimal information structure, Sender faces a tension between acquiring more information and alleviating the conflict of interest. On the one hand, Sender can take advantage of more information by making more precise recommendations (information value). On the other hand, the information structure is used to control the extent to which the two parties' interests are aligned. Essentially it could be geared coarsely (or precisely) to reduce conflict of interest.

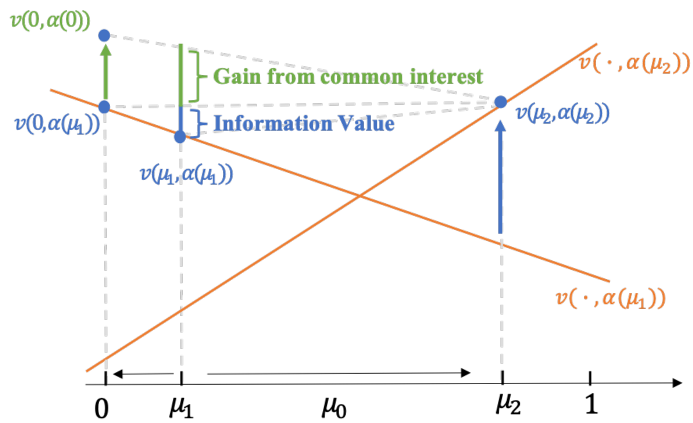

Figure 3: Conclusive good news

The key to the proof of Proposition 1 is inequality (4), for which we consider Figure 3. Loosely speaking, we construct X by splitting the belief μ_1 obtained under Y into a pair of more spread beliefs, 0 and μ_2, obtained under X. When Sender performs Y and induces the belief μ_1, Receiver takes action $\alpha(\mu_1)$. By splitting μ_1 into 0 and μ_2, Sender can induce Receiver to take $\alpha(\mu_2)$, which is a better action for Sender under his posterior belief μ_2, as shown by the blue arrow. In other words, having more information would allow Sender to make more precise recommendations in stead of pooling $\mu(x_1)$ and $\mu(x_2)$ together to make one single recommendation. Therefore, Sender will achieve a weakly positive payoff from obtaining more information while fixing Receiver's action space. Moreover, when Sender indeed conducts X, Receiver's best response also changes due to the new information structure. In particular, when conclusive good news is revealed, Receiver chooses $\alpha(0)$, which is precisely the action that Sender prefers when he confirms that the true state is 0. Therefore, by revealing the state of one-sided common interest, Sender would gain additional value, shown by the green arrow.

A symmetric argument, however, would not hold for conclusive news about state 1. For example, if we fix μ_1 and split μ_2 into μ_1 and 1, and the incentive constraints still hold, then Sender is still better off in terms of information value: when he holds belief μ_1, he can recommend a superior action $\alpha(\mu_1)$ in stead of inducing $\alpha(\mu_2)$. On the other hand, if state 1 were fully revealed, there could be a loss from conflict of interest in state 1. If $v(1, \alpha(1))$ is substantially lower than $v(1, \alpha(\mu_2))$, the experiment that fully reveals state 1 would be suboptimal.

Furthermore, note that Proposition 1 only implies that experiment $(0, \mu_2)$ is better than (μ_1, μ_2). It is not necessarily true that any (μ_1', μ_2) with $\mu_1' < \mu_1$ dominates (μ_1, μ_2). In other words, Sender's payoff is not necessarily monotone when having more information about state 0, although this monotonicity may be obtained with a stronger assumption, such as "downward bias."[21] Intuitively, when Sender has a consistent incentive to under-represent information, he wants to decrease Receiver's belief. The credible way to reduce Receiver's belief is to design an experiment that generates a more conclusive signal about the state 0.

4.1 Optimal experiment

After we prove the optimal experiment is conclusive about the state 0, we can pin it down using a geometric approach. Additionally, given that we are focusing on Sender-preferred PBE, it can be verified without loss of generality to let Receiver break ties in Sender's favor (formally proved in the Appendix).[22] Given one-sided common interest in state 0, we only need to focus on Sender's IC that

$$v(\mu_2, \hat{\alpha}(\mu_2)) \geq v(\mu_2, \hat{\alpha}(0)) \tag{6}$$

[21]Denote $\mu^*(\mu^s) = \arg\max_{\mu^r} \min_\alpha v(\mu^s, \alpha(\mu^r))$. Suppose that $\inf \mu^*(\mu^s) \leq \mu^s$ for all $\mu^s \in [0, 1]$ and $\inf \mu^*(\mu^s) = \mu^s$ implies $\mu^s \in \mu^*(\mu^s)$. We say that Sender has a "downward bias."

[22]This result relies on our assumptions on Sender and Receiver preferences and this is not generally true for arbitrary payoff functions.

Denote $B = \{\mu_2 \in (\mu_0, 1] : \text{inequality 6 hlods}\}$, which is not necessarily a convex set. We construct a new function: $\overline{v}(\mu) : [0, \max\{B\}] \to \mathbb{R}$.

$$\begin{cases} \overline{v}(\mu) = \hat{v}(\mu), & \text{if } \mu = 0 \text{ or } \mu \in B \\ \overline{v}(\mu) = \inf \hat{v}(\mu), & \text{if } \mu \neq 0 \text{ and } \mu \notin B \end{cases} \tag{7}$$

Define $\overline{V}(\cdot)$ as the concavification of $\overline{v}(\cdot)$.

$$\overline{V}(\mu) = \sup_{\mu', \mu'' \in [0, \max B],\ \gamma \in [0,1],\ s.t\ \gamma\mu' + (1-\gamma)\mu'' = \mu} \{\gamma \overline{v}(\mu') + (1-\gamma)\overline{v}(\mu'')\}$$

Proposition 2. *With A1, A1′ and A2, if Sender can achieve expected payoff strictly higher than with babbling, the highest ex-ante payoff he can achieve is $\overline{V}(\mu_0)$, whose solution $(0, \mu_2^*)$ is the optimal information structure where $\mu_2^* \in B$.*

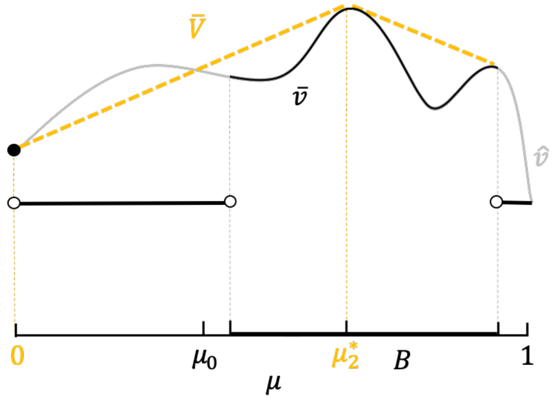

Figure 4: Optimal experiment

Figure 4 is an example showing how we construct $\overline{v}$ and $\overline{V}$. First, we find B, the set of μ_2 such that Sender weakly prefers to truthfully report information outcomes. Next, $\overline{v}(\mu)$ remains the same with $\hat{v}(\mu)$ if $\mu = 0$ or $\mu \in B$. Otherwise, we let $\overline{v}(\mu)$ be a sufficiently small number, e.g., $\inf \hat{v}(\mu)$, when $\mu \in [0, \max B]$. By this construction, we reduce Sender's expected payoff from those experiments that are not incentive compatible and to make them suboptimal. Last, we construct $\overline{V}$ as the concavifivation of $\overline{v}$. With the standard concavification argument, the incentive compatible experiment $(0, \mu_2^*)$ as the solution of $\overline{V}(\mu_0)$ is then optimal.

5 More than two states

Given Proposition 1, a natural extension might be that Sender optimally designs an information structure that fully reveals the common-interest state(s) when the state space contains more than two elements. However, this is not generally true. One reason is that revealing the common-interest state(s) may fail Sender's incentive constraint in the sense that Receiver's optimal action(s) to the common-interest state(s) are more preferable to Sender when other information comes in. Nevertheless, even if it survives Sender's incentive compatibility, it may still be sub-optimal.

To illustrate, we give a simple example with three states $\{\omega_1, \omega_2, \omega_3\}$ and three actions $\{a_1, a_2, a_3\}$. The prior belief is $\mu_0 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, each element corresponding to the probability of each state. Following the conventional visualization,[23] we use an equilateral triangular plane to depict all probability distributions over the state space. For example, each vertex of a triangle represents the distribution that put mass 1 on ω_i. In Figure 5(a), we depict Receiver's preference on her belief space. In particular, the simplex is divided into three regions where actions a_1, a_2, and a_3 are respectively optimal to Receiver. The intersection of any two regions are beliefs where Receiver is indifferent.[24] Sender obtains utility 1 from a_1 and utility 0 from a_3 regardless of the state. With a_2, his payoff is $(0, 0.6, 3)$ in state $\{\omega_1, \omega_2, \omega_3\}$ respectively. Obviously, a_3 is dominated for all beliefs, therefore in Figure 5(b), the simplex is divided to two regions by the orange line where Sender is indifferent between a_1 and a_2. Combining 5(a) and 5(b), we can easily obtain 5(c), in which the orange region characterizes the beliefs that both parties prefer the same action,[25] Note that ω_1 is the common interest state that Sender and Receiver both prefers a_1. The gray dot denotes the prior belief.

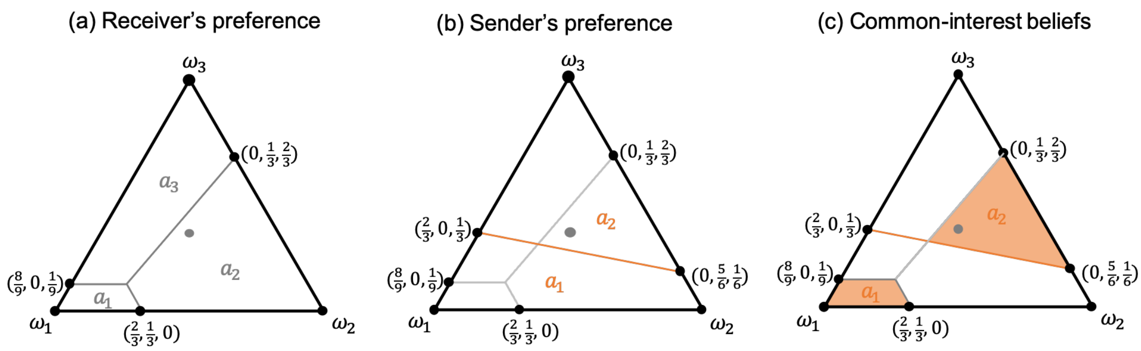

Figure 5: Common-interest beliefs

To give a brief explanation why revealing ω_1 is sub-optimal, consider experiment X and experiment Y in Figure 6(a) and 6(c) respectively. Experiment X is drawn as the green line in 6(a) such that it splits the prior belief into two posterior beliefs, $(0, \frac{1}{2}, \frac{1}{2})$ and $(1, 0, 0)$. Hence it reveals ω_1 with probability $\frac{1}{3}$. Experiment Y is drawn as the blue line in 6(c) such that it splits the prior belief into $(0, \frac{1}{3}, \frac{2}{3})$ and $(\frac{2}{3}, \frac{1}{3}, 0)$. Therefore Y does not give conclusive information about ω_1. Note that both experiments are incentive compatible for full communication since all posterior beliefs are of common-interest.

We want to show Y is strictly better than X. First, referring to Y, action a_1 is induced when the posterior $(\frac{2}{3}, \frac{1}{3}, 0)$ is realized. If we split $(\frac{2}{3}, \frac{1}{3}, 0)$ into $(1, 0, 0)$ and $(0, 1, 0)$ as in Figure 6(d), and *let* Receiver take a_1 at both beliefs,[26] then Sender obtains

[23] See Mas-Colell et al. (1995) page 169 and Bergemann et al. (2015).

[24] Given the exact indifferent beliefs in Figure 5(a), we can derive Receiver's cardinal utilities, e.g., her payoffs from actions $\{a_1, a_2, a_3\}$ are (1,0,0) in state ω_1, (1,3,0) in state ω_2 and (1,3,9) in state ω_3.

[25] We assume Sender-preferred tie breaking rule.

[26] a_1 is not Receiver's optimal action when $(0, 1, 0)$ is realized.

the same expected payoff with this new experiment $((1,0,0),(0,1,0),(0,\frac{1}{3},\frac{2}{3}))$ as with experiment Y. Next, consider experiment X. Receiver takes a_2 with posterior $(0,\frac{1}{2},\frac{1}{2})$, therefore if we split this into $(0,\frac{1}{3},\frac{2}{3})$ and $(0,1,0)$ (as in Figure 6(b)) with Receiver taking a_2 at both of them, Sender obtains the same expected payoff with this new experiment $((1,0,0),(0,1,0),(0,\frac{1}{3},\frac{2}{3}))$ as with experiment X. However, Sender strictly prefers a_1 over a_2 at $(0,1,0)$. Hence, conducting experiment $((1,0,0),(0,1,0),(0,\frac{1}{3},\frac{2}{3}))$ but inducing Receiver to take a_1 at $(0,1,0)$ gives Sender strictly higher payoff, which makes Y strictly better than X.[27]

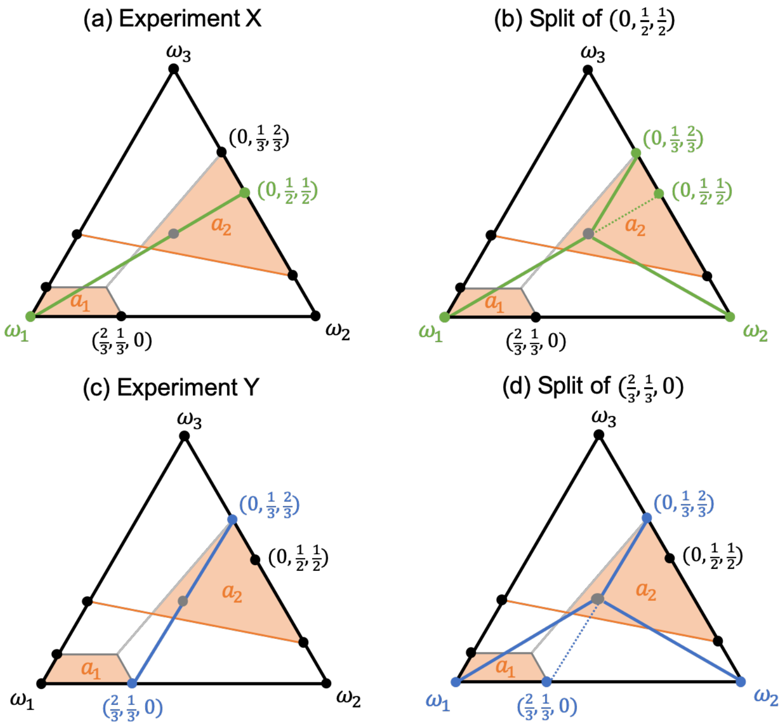

Figure 6: Sub-optimality to revealing ω_1

Therefore, when the state space contains more than two elements, we do not have a general property of the optimal information structure. It is mainly because there is a larger set of experiments that are not Blackwell-ordered when the probability distributions have more degrees of freedom.

Any Bayes plausible information structure, given which all posterior beliefs induce common-interest, is incentive compatible for a full communication equilibrium. However, the opposite is *not* true. For a full communication equilibrium to exist, we only require that Sender weakly prefers to report the true posterior over other posteriors

[27] Since in this example, only a_1 and a_2 can be induced in a valuable full communication equilibrium, it is then without loss to look at experiment that generates only two posteriors. However, the reader might note that we can also split a binary experiment into a more spreading experiment so that ω_1 can be revealed, e.g., the incentive compatible experiment $((\frac{5}{6},\frac{1}{6},0),(0,\frac{4}{9},\frac{5}{9}))$ can be split into another incentive compatible experiment $((1,0,0),(\frac{2}{3},\frac{1}{3},0),(0,\frac{4}{9},\frac{5}{9}))$ that reveals ω_1. But with similar argument, we can show such experiment is sub-optimal.

that the experiment might generate, *without* regarding to beliefs that cannot be realized by the experiment. Therefore, when the action space is large, it is also difficult to find the optimal experiment by digging into the set of common-interest beliefs.

6 When Sender Cannot Commit

In this section, we explore the possibility that Sender has no commitment power on both information structure and reporting policy. In particular, after Sender proposes an experiment to Receiver, he is at liberty to choose any signal generating process without being audited. This is an extreme case of limited commitment where Sender has the least credibility. Nevertheless, the experiment itself becomes part of the equilibrium. In particular, one more incentive constraint to check is whether Sender has incentive to deviate to another experiment.

Lemma 2. *For any finite state space, an experiment π is part of an (informative) full-communication equilibrium if and only if for each $\omega \in \Omega$,*

$$u(\alpha(\mu(s_i)), \omega) = \max_{s \in \mathcal{S}} u(\alpha(\mu(s)), \omega) \text{ for all } s_i \text{ such that } \pi(s_i|\omega) > 0 \tag{8}$$

Essentially, condition (8) requires that, for each state ω and for all s_i that are generated with positive probability in ω, reporting those signals is optimal in ω. When Sender has no commitment power on information structure, a full-communication equilibrium requires two sets of incentive constraints. (1) At stage 1: given Receiver's decision rule, Sender weakly prefers to perform the experiment that he proposes to Receiver. (2) At stage 2: given the information structure Sender chooses, he weakly prefers to report the true information outcomes. The first one is important when the proposed experiment is not perfectly informative. In this case, deviating to a more informative experiment would allow for (weakly) better use of the information in terms of which message to send. Condition (8) ensures that a more informative experiment does not give Sender strictly positive information value, and this condition turns out to be sufficiently strong for ensuring the incentive constraints in stage 2. To see this, consider an arbitrary experiment π shown as the matrix below.

$$\begin{array}{cc} & \begin{array}{ccc} s_1 & s_2 & s_3 \end{array} \\ \begin{array}{c} \omega_1 \\ \omega_2 \\ \omega_3 \end{array} & \begin{pmatrix} \pi_{11} & \pi_{12} & 0 \\ 0 & \pi_{22} & \pi_{23} \\ \pi_{31} & 0 & \pi_{33} \end{pmatrix} \end{array}$$

If this experiment induces a full communication equilibrium, then taking ω_2 as example, Sender has to be indifferent between reporting s_2 and s_3 when the true state is ω_2. Besides, both are optimal for Sender in ω_2. If this is not true, e.g., reporting s_3 is not optimal in ω_2, then Sender is strictly better off by deviating to a perfect information structure. In particular, by splitting s_3 into two signals that fully reveal ω_2 and ω_3 separately, then Sender can report the signal that gives him payoff higher than s_3

when ω_2 is revealed. This is strictly better than having non-degenerate belief $\mu(s_3)$ and reporting s_3 to Receiver. Furthermore, Sender's incentive constraints in stage 2 hold trivially given condition (8). Formal proof is in the Appendix.

Note that in this section, it is still without loss to focus on full communication equilibrium in terms of Sender's expected payoff. In particular, to maintain the existence of a pooling equilibrium, in each state ω, Sender has to be indifferent among all the signal realizations that are generated in that state, meaning that he is indifferent among all the messages he could report in ω given those realizations. By Lemma 2, this is strong enough to induce a full communication equilibrium with the new experiment, which is a garbling of the original experiment using the reporting strategy as the garbling matrix.

Proposition 3. *An (informative) full communication equilibrium exists if and only if there exists an (informative) equilibrium in which Sender is perfectly informed about the true state and his expected payoff remains the same. Specifically, there exists a distribution G and $\alpha(\cdot)$, such that for each ω,*

$$u(\alpha(\mu_j), \omega) = \max_{\mu \in supp(G)} u(\alpha(\mu), \omega) \text{ for all } \mu_j \text{ such that } \Pr(\omega|\mu_j) > 0 \tag{NC}$$

Proof. The if part is obvious. For the only if part, consider a full communication equilibrium with experiment π and condition (8) holds. Sender can replicate the same equilibrium outcome by performing a fully-revealing experiment and using π as the reporting policy, i.e., reporting $\mu(s_i)$ with probability $\pi(s_i|\omega)$ when ω is fully revealed. Condition (8) implies that this reporting policy is incentive compatible. Additionally, with this construction, the distribution of Receiver's posterior does not change and $supp(G) = \{\mu(s_i)\}$. Hence, by Claim 1, Sender obtains the same expected payoff. Therefore, we can modify condition (8) to (NC) by assuming Sender performing a fully-revealing experiment. After Sender privately learns the true state ω, he considers it's optimal to report any belief in the set of $\{\mu_j\}$ that $\Pr(\omega|\mu_j) > 0$. □

When Sender cannot commit to the experiment, though his preferences are state-dependent, Proposition 3 implies that endogenous choice of experiment *does not* provide Sender with strictly positive valuation. Hence if Sender has no commitment power, for any arbitrary payoff functions (transparent motives or state-dependent preferences), it is without loss to discuss cheap talk game in which Sender is privately informed about the true state. This rationalizes the cheap talk literature's assumption that Sender initially has perfect information.

7 Conclusion

This paper discusses cheap talk games where an uninformed sender can endogenously acquire information before communication. However, the freedom to acquire information is only valuable when the sender's preferences are state-dependent *and* the information structure is public observable. We restrict our attention to preferences that are of one-sided common interest, since it is relatively difficult to find a general property of the optimal information structure without inducing any structure to the payoff functions. But we believe it will be very interesting to discuss cheap talk with endogenous conflict of interest in a more general set-up, which will require different techniques and structures.

References

Alonso, Ricardo and Odilon Camara, "Tampering with Information," *working paper*, 2018.

Argenziano, Rossella, Sergei Severinov, and Francesco Squintani, "Strategic Information Acquisition and Transmission," *American Economic Journal: Microeconomics*, August 2016, *8* (3), 119–55.

Armstrong, Mark and John Vickers, "A Model of Delegated Project Choice," *Econometrica*, 2010, *78* (1), 213–244.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris, "The Limits of Price Discrimination," *American Economic Review*, March 2015, *105* (3), 921–57.

Blackwell, David and M. A. Girshick, *Theory of Games and Statistical Decisions* 1979.

Börgers, Tilman, "Notes on Blackwell's Comparison of Experiments," *http://www-personal.umich.edu/ tborgers/Blackwell-Girshick.pdf*, 2009.

Crawford, Vincent P. and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, *50* (6), 1431–1451.

Deimen, Inga and Dezső Szalay, "Delegated Expertise, Authority, and Communication," *American Economic Review*, April 2019, *109* (4), 1349–74.

Di Tillio, Alfredo, Marco Ottaviani, and Peter Norman Sørensen, "Persuasion bias in science: can economics help?," *The Economic Journal*, 2017, *127* (605), F266–F304.

Dworczak, Piotr and Giorgio Martini, "The Simple Economics of Optimal Persuasion," *Journal of Political Economy*, 2019, *127* (5), 1993–2048.

Guo, Yingni, "Dynamic Delegation of Experimentation," *The American Economic Review*, 2016, *106* (8), 1969–2008.

_ **and Eran Shmaya**, "Costly Miscalibration," *Forthcoming at Theoretical Economics*, 2020.

Inderst, Roman and Marco Ottaviani, "Misselling through Agents," *The American Economic Review*, 2009, *99* (3), 883–908.

Kamenica, Emir and Matthew Gentzkow, "Bayesian persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Kartik, Navin, “Strategic Communication with Lying Costs,” *The Review of Economic Studies*, 2009, *76* (4), 1359–1395.

Kosenko, Andrew, “Mediated Persuasion,” *Working Paper*, 2018.

Krähmer, Daniel, “Information Design and Strategic Communication,” *Forthcoming at American Economic Review: Insights*, 2020.

Lipnowski, Elliot and Doron Ravid, “Cheap Talk with Transparent Motives,” *Econometrica*, 2020, *88* (4), 1631–1660.

_ , _ , and Denis Shishkin, “Persuasion via Weak Institutions,” *working paper*, 2019.

Mas-Colell, Andreu, Michael Dennis Whinston, and Jerry R. Green, “Microeconomic theory,” *Oxford University Press*, 1995.

Pei, Di, “Communication with endogenous information acquisition,” *Journal of Economic Theory,*, 2015, *160*, 132–149.

Strulovici, Bruno, “Mediated Truth,” *Working Paper*, 2017.

Tan, Teck Yong and Anh Nguyen, “Bayesian Persuasion with Costly Messages,” *Working Paper*, 2021.

Appendix A

Proof of Claim 1

Though the message space is rich, it is without loss to assume $M = supp(G)$. To simplify the exposition, we omit the notion of α and we prove equation (1) from RHS to LHS.

$$\begin{aligned}\sum_{supp(G)} \Pr(\mu^r)v(\mu^r,\mu^r) &= \sum_{supp(G)} \Pr(\mu^r)\sum_{\Omega}\mu^r u(\mu^r,\omega)\\ &= \sum_{supp(G)} \Pr(\mu^r)\sum_{\Omega}\sum_{\mathcal{S}}\frac{\Pr(s)\delta(\mu^r|s)}{\Pr(\mu^r)}\mu^s(s)u(\mu^r,\omega)\\ &= \sum_{supp(G)}\sum_{\Omega}\sum_{\mathcal{S}}\Pr(s)\delta(\mu^r|s)\mu^s(s)u(\mu^r,\omega)\\ &= \sum_{\mathcal{S}}\Pr(s)\sum_{supp(G)}\delta(\mu^r|s)\sum_{\Omega}\mu^s(s)u(\mu^r,\omega)\\ &= \sum_{\mathcal{S}}\Pr(s)\sum_{supp(G)}\delta(\mu^r|s)v(\mu^s(s),\mu^r)\end{aligned}$$

Proof of Lemma 1

Let $\mu_i := \mu(s_i) \in supp(F)$ and $\mu_j \in supp(G)$. We want to show that experiment G is incentive compatible for a full communication equilibrium. In the original pooling equilibrium, for each belief μ_j, there is a set of beliefs $\{\mu_i\}_{i\in I}$ that a type-μ_i Sender recommends μ_j with probability $\delta(\mu_j|s_i) > 0$, indicating that $v(\mu_i, \alpha(\mu_j)) \geq v(\mu_i, \alpha(\mu_j'))$ for all $\mu_j, \mu_j' \in supp(G)$, and μ_j is a weighted average over beliefs $\{\mu_i\}_{i\in I}$. In particular, $\mu_j = \sum_I \frac{\Pr(s_i)\delta(\mu_j|s_i)}{\sum_I \Pr(s_i)\delta(\mu_j|s_i)}\mu_i$. Note that for fixed action of Receiver, Sender's expected payoff is linear in his belief. Hence,

$$v(\mu_j,\alpha(\mu_j)) = \sum_I \frac{\Pr(s_i)\delta(\mu_j|s_i)}{\sum_I \Pr(s_i)\delta(\mu_j|s_i)} v(\mu_i,\alpha(\mu_j))$$

If for all $i \in I$, type-μ_i Sender weakly prefers to report μ_j, then type-μ_j Sender weakly prefers to report μ_j. Hence a full communication equilibrium exists with experiment G, and Sender's expected payoff remains the same by Claim 1.

Proof of inequality (2)

Denote $\sup_\alpha v(\mu^s, \alpha(\mu^r)) := \overline{f}(\mu^r)$ and $\inf_\alpha v(\mu^s,\alpha(\mu^r)) := \underline{f}(\mu^r)$, where we omit the notion for μ^s. Denote $f(\mu^r) = \alpha(\mu^r)\overline{f}(\mu^r) + (1-\alpha(\mu^r))\underline{f}(\mu^r)$. We want to show that

$$f(\lambda\mu_1^r + (1-\lambda)\mu_2^r) \geq \min\{\overline{f}(\mu_1^r), \overline{f}(\mu_2^r)\} \tag{9}$$

Suppose it is not true, then $f(\lambda\mu_1^r+(1-\lambda)\mu_2^r) < \overline{f}(\mu_1^r)$ and $f(\lambda\mu_1^r+(1-\lambda)\mu_2^r) < \overline{f}(\mu_2^r)$. Since, $f(\lambda\mu_1^r + (1-\lambda)\mu_2^r) \geq \underline{f}(\lambda\mu_1^r + (1-\lambda)\mu_2^r)$, then by semi-continuity and payoff function is one-side jump, there exists a $\mu_3^r \to \lambda\mu_1^r + (1-\lambda)\mu_2^r$ such that μ_3^r is a convex combination of $\mu_1^r + (1-\lambda)\mu_2^r$ and μ_2^r. Then $\overline{f}(\mu_3^r) < \overline{f}(\mu_2^r)$ and $\overline{f}(\mu_3^r) < \overline{f}(\mu_1^r)$, which violates A1. Since inequality (9) is true for all $\alpha(\mu^r)$, then $v(\mu^s, \alpha(\mu^r))$ is quasi-concave in μ^r no matter what the mixing probability Receiver is using.

Proof of Proposition 2

Call an equilibrium that Sender gets payoff strictly higher than with babbling a *profitable* equilibrium. We want to check if it is without loss to assume Receiver breaks ties in Sender's favor. To show this, the set of incentive compatible experiments cannot shrink if we assume $\alpha = \hat{\alpha}$. Essentially, $v(\mu_2, \alpha(\mu_2)) \geq v(\mu_2, \alpha(0))$ and $v(0, \alpha(0)) \geq v(0, \alpha(\mu_2))$ imply that $v(\mu_2, \hat{\alpha}(\mu_2)) \geq v(\mu_2, \hat{\alpha}(0))$ and $v(0, \hat{\alpha}(0)) \geq v(0, \hat{\alpha}(\mu_2))$. Note that the last inequality is always true given common interest assumption in state 0. Hence, we only need to check if $v(\mu_2, \hat{\alpha}(\mu_2)) \geq v(\mu_2, \hat{\alpha}(0))$: (1) Suppose $v(\mu_2, \alpha(\mu_2)) > v(\mu_2, \alpha(0))$, then by quasi-concavity, $v(\mu_2, \alpha(\mu_2)) > v(\mu_2, \hat{\alpha}(0))$ and hence $v(\mu_2, \hat{\alpha}(\mu_2)) > v(\mu_2, \hat{\alpha}(0))$; (2) Suppose $v(\mu_2, \alpha(\mu_2)) = v(\mu_2, \alpha(0))$, then same argument validates if there exists a $\mu \in (0, \mu_2)$ such that $v(\mu_2, \alpha(\mu_2)) = v(\mu_2, \alpha(0)) < v(\mu_2, \alpha(\mu))$. Otherwise, there is no profitable equilibrium due to A1$'$ even if $v(\mu_2, \hat{\alpha}(\mu_2)) < v(\mu_2, \hat{\alpha}(0))$.

Proof of lemma 2

Let $\pi_{ni} := \pi(s_i|\omega_n)$. First, we prove the only if part. Suppose this is not true, then there exist a ω_n and s_i such that $u(\alpha(\mu(s_i)), \omega_n) < \max_{s \in \mathcal{S}} u(\alpha(\mu(s)), \omega_n)$. Then if Sender deviates to a perfect signal structure, for all n such that $\pi_{ni} > 0$,

$$v(\mu(s_i), \alpha(\mu(s_i))) = \sum_n c_{ni}\, u(\alpha(\mu(s_i)), \omega_n) < \sum_n c_{ni}[\max_s u(\alpha(\mu(s)), \omega_n)]$$

where $c_{ni} = \pi_{ni} \frac{\Pr(\omega_n)}{\Pr(s_i)}$. Then multiplying both LHS and RHS by $\Pr(s_i)$ and summing over i, we have

$$\sum_i \Pr(s_i) v(\mu(s_i), \alpha(\mu(s_i))) < \sum_n \Pr(\omega_n)[\max_s u(\alpha(\mu(s)), \omega_n)]$$

Second, the if part is simply implied by $u(\alpha(\mu(s_i)), \omega_n) = \max_{s \in \mathcal{S}} u(\alpha(\mu(s)), \omega_n)$ for all i such that $\pi_{ni} > 0$. In particular, conditional on s_i being realized, the true state must be one of the $\{\omega_n\}$ that $\pi_{ni} > 0$. But no matter which one the true state is, Sender weakly prefers to report s_i. Hence, the incentive constraints hold.

Appendix B

In this section, we are going to characterize the result without A1$'$.

Lemma 3. *Suppose there exists an incentive compatible experiment (μ_1, μ_2) where $\mu_0 \in (\mu_1, \mu_2)$ and $v(\mu_2, \alpha(\mu))$ is a constant for all $\mu \in [\mu_1, \mu_2]$. If there exists an experiment $(0, \mu_2')$ that is incentive compatible, then $(0, \mu_2')$ dominates (μ_1, μ_2).*

Proof. First, trivially, if $(0, \mu_2)$ validates IC, then we can use the same proof as in

Proposition 1 to show $(0, \mu_2)$ is better than (μ_1, μ_2). Hence, we only need to focus on cases that $(0, \mu_2)$ violates IC, in particular, $v(\mu_2, \alpha(0)) \geq v(\mu_2, \alpha(\mu))$ for all $\mu \in [\mu_1, \mu_2]$. Additionally, $v(0, \alpha(0)) \geq v(0, \alpha(\mu))$ for all $\mu \in [\mu_1, \mu_2]$, therefore $v(\mu', \alpha(0)) \geq v(\mu', \alpha(\mu))$ for all $\mu \in [\mu_1, \mu_2]$ and $\mu' \in [0, \mu_2]$. Hence if there exists an experiment $(0, \mu_2')$ that is incentive compatible, then $\mu_2' > \mu_2$.

Now we want to show that Sender's payoff from full communication equilibrium with (μ_1, μ_2) is weakly smaller than his payoff from conducting experiment $(0, \mu_2)$ and Receiver taking $\alpha(0)$, which is weakly smaller than that from full communication with experiment $(0, \mu_2')$. The first part is simple because Sender prefers to recommend belief 0 rather than reporting belief μ_1 or μ_2. The second part comes from two facts: (1) $(0, \mu_2')$ is a mean-preserving spread of $(0, \mu_2)$ and (2) $v(0, \alpha(0)) \geq v(0, \alpha(\mu_2'))$ and $v(\mu_2', \alpha(\mu_2')) \geq v(\mu_2', \alpha(0))$. Hence, by the similar proof with Proposition 1, we can have the second part. $\square$

$\mathcal{P} : \{[a, b] \ni \mu_0 : \forall x \in [a, b], v(b, \alpha(x))$ is a constant and $\exists x \neq x'$, s.t. $\alpha(x) \neq \alpha(x')\}$. Denote $\mathcal{U} : \{b > \mu_0 : \forall [a, b] \in \mathcal{P}\}$. Note that $|\mathcal{U}| \leq 2$. This is because, if for some $x \neq x'$, s.t. $\alpha(x) \neq \alpha(x')$ and $v(b, \alpha(x)) = v(b, \alpha(x'))$, then for all $b' \neq b$, $v(b, \alpha(x)) \neq v(b, \alpha(x'))$, otherwise, $\alpha(x) = \alpha(x')$. Denote $\underline{a}(b) : \{a < \mu_0 : \min a \text{ s.t. } [a, b] \in \mathcal{P}\}$. $\Pi^o : \{(\underline{a}(\mu_2), \mu_2) : \mu_2 \in \mathcal{U}\}$ is a set of experiments and contains at most two elements.

Proposition 4. *With A1, A2 and one-sided common interest in state 0, suppose there exists a equilibrium in which Sender gets payoff strictly higher than babbling. (1) If $B \neq \emptyset$, then the highest payoff sender can achieve is $\overline{V}(\mu_0)$, whose solution $(0, \mu_2^*)$ is the optimal experiment where $\mu_2^* \in B$; (2) If $B = \emptyset$, then the optimal experiment belongs to Π^o.*

Proof. The first part is straight forward given Lemma 3 and Proposition 1. For the second part, if $B = \emptyset$, then a profitable equilibrium exists when (a) $v(\mu_2, \alpha(\mu_2)) = v(\mu_2, \alpha(\mu_1))$ and $v(\mu_2, \alpha(\mu))$ is a constant for all $\mu \in [\mu_1, \mu_2]$; and (b) $v(\mu_1, \alpha(\mu_1)) \geq v(\mu_1, \alpha(\mu_2))$. Note that (b) can be implied by (a) for the reason that $v(\mu_2, \alpha(\mu_1)) = v(\mu_2, \alpha(\mu_2))$ and $v(0, \alpha(\mu_1)) \geq v(0, \alpha(\mu_2))$ (given the assumption of common interest in state 0 and quasi-concavity). Considering some fixed μ_2 and when (a) is satisfied, all experiments (μ, μ_2) where $\mu \in [\mu_1, \mu_0)$ are incentive compatible. Among those experiments, we want to show $(\underline{a}(\mu_2), \mu_2)$ is the best. Essentially, we want to show (μ, μ_2) dominates (μ', μ_2) when $\mu < \mu'$ and $\mu, \mu' \in [\mu_1, \mu_0)$: (1) $v(\mu, \alpha(\mu)) \geq v(\mu, \alpha(\mu'))$ is implied by $v(\mu_2, \alpha(\mu)) = v(\mu_2, \alpha(\mu'))$ and $v(0, \alpha(\mu)) \geq v(\mu_2, \alpha(\mu'))$; (2) (μ, μ_2) is a mean-preserving spread of (μ', μ_2). By (1) and (2), the optimal experiment belongs to Π^o for the same reason as in Proposition 1. Recall that the carnality of Π^o is at most two, hence it is easy to find the optimal experiment. $\square$