DISCUSSION PAPER SERIES

DP16801

Cross-verification and Persuasive Cheap Talk

Alp Atakan, Mehmet Ekmekci and Ludovic Renou

ORGANIZATIONAL ECONOMICS

ISSN 0265-8003

Cross-verification and Persuasive Cheap Talk

Alp Atakan, Mehmet Ekmekci and Ludovic Renou

Discussion Paper DP16801
Published 10 December 2021
Submitted 08 December 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Cross-verification and Persuasive Cheap Talk

Abstract

We study a cheap-talk game where two experts first choose what information to acquire and then offer advice to a decision-maker whose actions affect the wel- fare of all. The experts cannot commit to reporting strategies. Yet, we show that the decision-maker's ability to cross-verify the experts' advice acts as a commit- ment device for the experts. We prove the existence of an equilibrium, where an expert's equilibrium payoff is equal to what he would obtain if he could commit to truthfully revealing his information.

JEL Classification: N/A

Keywords: N/A

Alp Atakan - alpeatakan@gmail.com
QMUL

Mehmet Ekmekci - ekosistem@gmail.com
Boston College

Ludovic Renou - lrenou.econ@gmail.com
Queen Mary University of London and CEPR

Acknowledgements
We are grateful for the helpful comments of Navin Kartik, Elliot Lipnowski and Joel Sobel. Lu- dovic Renou gratefully acknowledges the support of the Agence Nationale pour la Recherche under grant ANR CIGNE (ANR-15-CE38-0007-01) and through the ORA Project "Ambiguity in Dynamic Environments" (ANR-18-ORAR-0005). Atakan's work on this project was supported by a grant from the European Research Council (ERC 681460, InformativePrices).

Cross-verification and Persuasive Cheap Talk*

Alp Atakan[†] & Mehmet Ekmekci[‡] & Ludovic Renou [§]

October 6, 2021

Abstract

We study a cheap-talk game where two experts first choose what information to acquire and then offer advice to a decision-maker whose actions affect the welfare of all. The experts cannot commit to reporting strategies. Yet, we show that the decision-maker's ability to cross-verify the experts' advice acts as a commitment device for the experts. We prove the existence of an equilibrium, where an expert's equilibrium payoff is equal to what he would obtain if he could commit to truthfully revealing his information.

Keywords: Bayesian persuasion, information design, commitment, cheap talk, multiple experts

JEL Classification Numbers: C73, D82

*We are grateful for the helpful comments of Navin Kartik, Elliot Lipnowski and Joel Sobel. Ludovic Renou gratefully acknowledges the support of the Agence Nationale pour la Recherche under grant ANR CIGNE (ANR-15-CE38-0007-01) and through the ORA Project "Ambiguity in Dynamic Environments" (ANR-18-ORAR-0005). Atakan's work on this project was supported by a grant from the European Research Council (ERC 681460, InformativePrices).

†Koc University and QMUL.

‡Boston College

§QMUL and CEPR

1 Introduction

Decision-makers routinely solicit advice from experts who have a vested interest in the decision at hand. Consulting multiple experts may allow a decision-maker to check the veracity of the advice that he receives by comparing one expert's recommendation with another's. We call this practice *cross-verification* and study how cross-verification affects communication.

Naturally, the effectiveness of cross-verification depends on the experts' information. If the experts have perfectly correlated information, then inconsistent recommendations definitively indicate untruthful, self-serving advice. Alternatively, if experts have uncorrelated information, then cross-verification cannot detect misleading advice. Thus, if the experts *strategically* acquire information, their choices will affect the scope for cross-verification. This paper sheds light on this interplay by analyzing a cheap-talk game, where experts independently acquire information before providing advice to a decision-maker. More precisely, we study the following game: two experts with identical preferences, first select statistical experiments that provide information about an unknown state of the world. The selected experiments are observed by the decision-maker. The experts privately observe their experiments' outcomes, and then offer private reports to the decision-maker. The decision-maker collects all the reports and chooses an action.

As a benchmark, suppose that an expert could commit to revealing his experiment's outcome truthfully. Following Kamenica and Gentzkow (2011), we call the experiment that this expert would optimally select the expert-optimal experiment. In our model, however, the experts *cannot* commit. Yet, we show that there exists an equilibrium, where both experts choose the expert-optimal experiment and truthfully report the outcomes of their experiments. In equilibrium, the experts optimally select perfectly correlated experiments, which enable cross-verification to be most effective. In turn, cross-verification facilitates truthful communication and allows the experts to receive their best possible payoff. In other words, cross-verification acts as a commitment device.

The existence of such an equilibrium relies on three essential properties. First, we assume that the experts are free to choose arbitrarily correlated statistical experiments (see Green and Stokey (1978) and Gentzkow and Kamenica (2016, 2017)). In fact in the equilibrium that we construct, they choose to correlate their experiments' outcomes perfectly and thus allow the decision-maker to cross-verify their reports perfectly. Second, suppose one expert deviates from reporting the experiment's outcome truthfully, while the other is truthful. In this case, the decision-maker detects a deviation as the two reports are inconsistent. However, the decision-maker cannot deduce the deviator's identity. Third, we show that a *uniform* punishment always exists. There is an action that punishes the experts for deviating from truthful reporting, irrespective of the experts' private information.

The existence of the aforementioned uniform punishment is key to our equilibrium construction since the decision-maker does not know the deviator's identity and, therefore, cannot condition the punishment on the deviator's information. Proving the existence of a uniform punishment is the main technical contribution of the paper. We stress that the uniform punishment is relative to the expert-optimal experiment. Arbitrary experiments do not necessarily admit uniform punishments, and therefore, cross-verification does not necessarily elicit honest advice when the experts choose arbitrary experiments.

Our main result, described above, also generalizes to situations where the experts have non-identical preferences, provided that a uniform punishment continues to exist. In particular, we show that there is a uniform punishment when the preferences of the second expert are a convex combination of the preferences of the first expert and the decision-maker. For example, this is the case in the quadratic utility example of Crawford and Sobel (1982) when the two experts have like-biases.

Finally, we also study cross-verification from the decision-maker's perspective. We show that there is an equilibrium where the decision-maker benefits from cross-verification if the expert-optimal experiment is informative at some prior belief. The intuition is as follows: The decision-maker benefits from any additional information, and even the expert-optimal experiment provides valuable information in many circumstances. If the expert-optimal experiment does not provide useful information to the decision-maker, we appropriately modify the expert-optimal experiment. The modified experiment offers valuable information for the decision-maker, and the experts can truthfully communicate this information in equilibrium. We also establish this result's converse: the decision-maker's unique equilibrium payoff is equal to his payoff at his prior belief if the expert-optimal experiment is uninformative at every prior belief. In other words, the decision-maker only benefits from cross-verification in situations where the experts also benefit.

Related literature. This paper is related to the literature on cheap talk pioneered by Crawford and Sobel (1982) and several papers in this literature study communication with multiple experts. In particular, Battaglini (2002) shows that the decision-maker can learn a multidimensional state by consulting experts about different dimensions. Krishna and Morgan (2001a) and Ambrus and Takahashi (2008) are the papers most closely related to ours. These papers study two-expert cheap-talk games where the experts have quadratic preferences, perfectly observe the state and simultaneously send messages to the decision-maker. Krishna and Morgan (2001a) assume that the state space is unidimensional and show that there is an equilibrium where the experts truthfully reveal the state if the experts' preferences are identical.[1] Ambrus and Takahashi (2008) study the case where the state is multidimensional and show that there is an equilibrium where the experts truthfully reveal the state if the experts are biased in the same direction. Their equilibrium construction is similar to ours and uses cross-verification to elicit truth-telling: if an expert deviates from truthfully revealing the state, then the decision-maker chooses a uniform punishment and they show that such a uniform punishment exists if the experts are biased in the same direction. The survey by Sobel (2013) also discusses how cross-verification ensures truth-telling in the context of multi-sender cheap-talk games if there is an arbitrarily harsh exogenously-given punishment.[2]

Our work differs from these articles in several respects: Foremost, our main result shows that the experts obtain their commitment payoff. In contrast, the cheap-talk literature is predominantly interested in full information revelation. In other words, our emphasis is on the experts' perspective while the cheap-talk literature focuses on the decision-maker's perspective. Second, we assume that the experts choose what kind of information to acquire, while past works typically assume that the experts perfectly know the state.[3] This is an important distinction since the experts' information affects the scope for cross-verification. Third, the literature on cheap talk focuses on agents with single-peaked preferences and frequently assumes that all agents have quadratic utility. In contrast, we put no restrictions on the utility functions. With quadratic utility, the expert-optimal information structure coincides with the decision-maker's and entails choosing an experiment that perfectly reveals the state. Therefore, as in Krishna and Morgan (2001a) and Ambrus and Takahashi (2008), our result also implies that full information revelation is an equilibrium in this particular case. However, with other utility specifications, the expert-optimal and decision-maker optimal information structures need not coincide.

Our paper is also related to the following works that focus on single-expert cheap-talk games. Lyu (2020) characterizes the equilibrium set in a model where the expert acquires information before providing advice. Lipnowski (2020) shows that an expert

[1] In contrast, Krishna and Morgan (2001b) prove that such an equilibrium does not exist if the experts send messages sequentially.

[2] Also, see Wolinsky (2002), and Gilligan and Krehbiel (1989) for related work on multi-sender cheap-talk games.

[3] Ambrus and Lu (2014) and Mylovanov and Zapechelnyuk (2013) are notable exceptions. Ambrus and Lu (2014) show that if the state space is large enough, there are equilibrium outcomes of multi-sender cheap-talk games that are arbitrarily close to full revelation when the senders observe the state with noise, as the noise converges to zero. Mylovanov and Zapechelnyuk (2013) argue that the receiver can induce the senders to reveal commonly-known events if the receiver can commit to a decision rule instead of best replying to his belief as in our paper.

can obtain his commitment payoff if the expert's value function is continuous.[4] Instead, we focus on a model with multiple experts and show that the experts receive their commitment payoff, without making any assumptions on their payoff functions.

Finally, this paper is closely related to the literature on Bayesian persuasion (Kamenica and Gentzkow (2011)). A number of articles that include Au and Kawai (2020), Gentzkow and Kamenica (2016, 2017), Koessler et al. (2018) and Li and Norman (2018, 2020) study persuasion with multiple experts. In all of these papers, the experts can commit to revealing their information truthfully. In contrast, we assume that the experts' recommendations are cheap-talk, i.e., we require sequential rationality at every stage of the game. Our result shows that the experts can achieve their commitment payoff even though they cannot commit to revealing their information. For a recent survey of the literature on Bayesian persuasion, we refer to Kamenica (2019).

2 The Model

We study a cheap-talk game between two experts, labelled 1 and 2, and a decision-maker. The experts provide information to the decision-maker about a payoff-relevant state $\omega \in \Omega$, who then chooses an action $a \in A$. The sets A and Ω are finite. The experts have identical preferences. An expert's payoff is $u(a, \omega)$ when the decision-maker chooses action a and the state is ω. (We relax the assumption of identical preferences in the next section.) The decision-maker's payoff is $v(a, \omega)$. Initially, neither the experts nor the decision-maker knows the state. The common prior probability that the state is ω is $\pi^\circ(\omega)$.

We first provide an informal description of the cheap-talk game. The game has three stages. In the first stage, the two experts simultaneously choose statistical experiments. The selected experiments are publicly observed. In the second stage, each expert privately observes his experiment's outcome and then sends a message to the decision-maker. In the third stage, the decision-maker observes the experts' messages and chooses an action.

We now provide a formal description. To model the choice of statistical experiments, we follow Gentzkow and Kamenica (2016, 2017). These authors define a statistical experiment σ as a partition of $\Omega \times [0, 1]$ into finitely many (Lebesgue) measurable subsets. A signal s is an element of the partition σ, i.e., a measurable subset of $\Omega \times [0, 1]$. The probability of signal $s \in \sigma$ conditional on ω is the (Lebesgue) measure of the set

[4] The value function describes the expert's highest expected payoff at a given belief conditional on the decision-maker choosing a best-reply to that belief. Continuity of the value function is a strong assumption. E.g., with two states and two actions, it requires the expert to be indifferent between the two actions whenever the decision-maker is.

$\{x \in [0,1] : (\omega, x) \in s\}$.[5] Throughout, we omit the dependence on the experiment σ, and write λ_s for the probability of the signal s and π_s for the posterior probability. We denote the set of experiments that the experts can choose from by Σ.

In the first stage, expert i thus chooses an experiment $\sigma_i \in \Sigma$. The chosen experiments (σ_1, σ_2) are publicly observed. In the second stage, expert i privately observes the realization $s_i \in \sigma_i$ and sends a private message $m_i \in M_i$ to the decision-maker. We assume that the sets of messages are rich enough to communicate any signal realizations. Finally, the decision-maker observes the messages (m_1, m_2) (but not the realized signals (s_1, s_2)) and chooses an action a. We denote $\Gamma(\pi^\circ, u, v)$ the cheap-talk game. Note that different extensive-form games are consistent with our description. Throughout, we assume that the state $(\omega, x) \in \Omega \times [0,1]$ is chosen by Nature according to the probability distribution $\pi^\circ \times U[0,1]$ after the experts have chosen their experiments, where $U[0,1]$ denotes the uniform distribution on the unit interval. Thus, we have a proper sub-game after each choice of statistical experiments (σ_1, σ_2).

A strategy for expert i is a pair (σ_i, τ_i), where $\sigma_i \in \Sigma$ and $\tau_i(\sigma_i, \sigma_j, s_i) \in \Delta(M_i)$ for all $(\sigma_i, \sigma_j, s_i)$ with $s_i \in \sigma_i$. A strategy for the decision-maker specifies a mixed action $\alpha(\sigma_i, \sigma_j, m_i, m_j) \in \Delta(A)$ for all $(\sigma_i, \sigma_j, m_i, m_j)$.[6] The solution concept is weak perfect Bayesian equilibrium. We stress that this requires the beliefs to be consistent with the chosen experiments (σ_1, σ_2) even if these experiments are off the equilibrium path.

Few remarks are worth making. First, as in classical cheap-talk games, none of the experts can commit to reporting strategies.

Second, if the experiments are (σ_1, σ_2), then the joint probability of $(s_1, s_2) \in \sigma_1 \times \sigma_2$ conditional on ω is the measure of the set $\{x : (\omega, x) \in s_1 \cap s_2\}$. Thus, if both experts choose the same experiment σ, then the probability of $(s, s') \in \sigma \times \sigma$ is zero, whenever $s \neq s'$. (To see this, note that if $s \neq s'$, then $s \cap s' = \emptyset$ since σ is a partition.) In words, if both experts choose the same experiment, their realized signals are perfectly correlated. This property will turn out to be crucial. An alternative modeling approach would be to assume that there is a fixed set of statistical experiments and let the experts observe the realization of the experiment of their choices. This alternative modeling also implies that if the two experts choose to observe the same experiment's realization, their observations are identical. For instance, this would be the case if the experts source their information from the same provider.

[5]It is usual to model statistical experiments as probability kernels $\sigma^* : \Omega \to \Delta(S)$, where S is the (finite) set of signals. This implies the formulation that we use in this paper: for each ω, we can partition $[0,1]$ into $|S|$ non-empty and disjoint intervals such that the length of the s-th interval is $\sigma^*(s|\omega)$ when the state is ω. With a slight abuse of notation, we identify the probability kernel σ^* with that particular partition of $\Omega \times [0,1]$.

[6]To ease exposition, we do not explicitly consider randomizations over the choices of experiments. This does not affect any of our results.

Third, we can allow for the set of experiments to also include identical and independent experiments without affecting our results. To do so, it suffices to define an experiment as a finite partition of $\Omega \times [0,1] \times [0,1]$, with (ω, x, y) distributed according to $\pi^\circ \times U([0,1]) \times U([0,1])$. Intuitively, if the experts condition their random observations on x, they are correlated, while they are independent if one expert conditions on x and the other on y. Importantly, though, the option of choosing a perfectly correlated experiment remains available to the experts. For example, consider two statisticians learning about the state from a population sample. If they use the same sample x and perform identical statistical tests on this sample, then the statisticians will have perfectly correlated signals – the same statistics. In contrast, if they apply the same statistical test to two different samples x and y, they will have signals, which are independent conditional on the state (that is, the underlying population). Similarly, if two doctors use the same set of lab results when making a recommendation to a patient, they have perfectly correlated signals.

We now introduce some additional notation. We denote by $v(\alpha, \pi)$ the decision-maker's expected payoff when he chooses the mixed action α and his belief about ω is given by $\pi \in \Delta(\Omega)$, by $BR(\pi) := \{\alpha \in \Delta(A) : v(\alpha, \pi) \geq v(\alpha', \pi), \forall \alpha' \in \Delta(A)\}$ the decision-maker's best-replies at π, and by $br(\pi) \subset A$ the decision-maker's pure best-replies at π. Similarly, we write $u(\alpha, \pi)$ for an expert's expected payoff. Note that payoffs are linear in π.

Throughout, **truthful equilibria**, in which the two experts choose the same experiment in the first stage and truthfully report the common signal realization in the second stage, play an important role.

3 The Main Result

In this section, we show that the ability of the decision-maker to cross-verify information serves as a commitment device for the experts. More precisely, we show that there exists an equilibrium of the cheap-talk game, in which the experts obtain their *commitment value.*

We define the *commitment value* as the highest payoff an expert can obtain when he commits to truthfully disclose the realized signal, as in games of Bayesian persuasion. Formally, consider the persuasion game, where an expert first chooses a statistical experiment $\sigma : \Omega \to \Delta(S)$, *commits* to truthfully reveal the realized signal s to the decision-maker, who then makes a decision. Kamenica and Gentzkow (2011) prove that the best equilibrium payoff for the expert in this game is given by cav $\overline{u}(\pi^\circ)$, where cav $\overline{u}$ is the concavification of $\overline{u}$ and $\overline{u}(\pi) := \max_{\alpha \in BR(\pi)} u(\alpha, \pi)$ for all π. (See also

Aumann and Maschler, 1995.) For later reference, we write $(\lambda_s^*, \pi_s^*)_{s\in S}$ for an optimal splitting of the prior π°, that is, $\sum_{s\in S} \lambda_s^* \overline{u}(\pi_s^*) = \text{cav}\, \overline{u}(\pi^\circ)$ and $\sum_{s\in S} \lambda_s^* \pi_s^* = \pi^\circ$. We write Π^* for $\{\pi_s^* : s \in S\}$, $\text{co}\, \Pi^*$ for the convex hull of Π^*, and Δ^* for the set of all probability distributions over Π^*. The corresponding optimal experiment is denoted σ^*.

Theorem 1. *There exists a truthful equilibrium of the cheap-talk game, where both experts obtain their commitment value cav* $\overline{u}(\pi^\circ)$.

Before proving Theorem 1, we explain our result's logic with the help of a simple example. There are two states, ω_0 and ω_1, and four actions, a_0, a_L, a_R and a_1. The preferences are depicted in Figure 1. Throughout the example, probabilities refer to the probability of ω_1. Assume that $\pi^\circ = 0.45$.

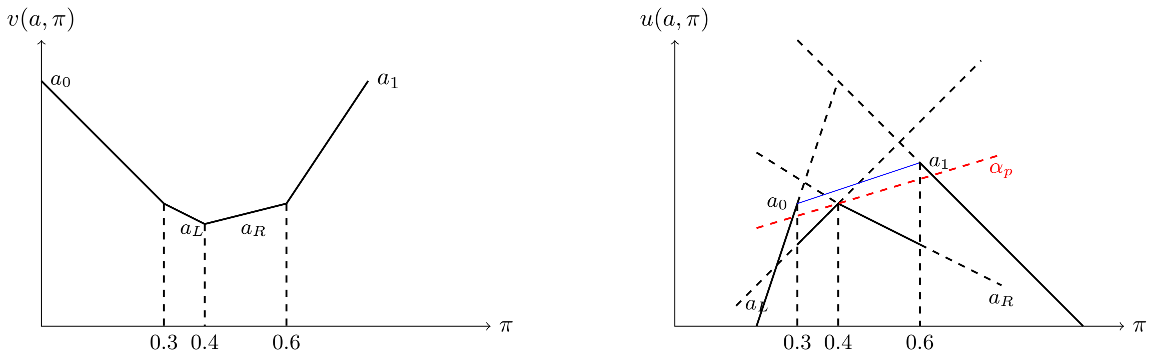

(a) DM's Preferences. This figure depicts $v(a,\pi)$ as a function of $\pi = \Pr[\omega = \omega_1]$. Action a_0 is optimal for the DM for $\pi \in [0, 0.3]$, action a_L is optimal for $\pi \in [0.3, 0.4]$, a_R is optimal for $[0.4, 0, 6]$, and a_1 is optimal for $[0.6, 1]$.

(b) Expert's Preferences. The black dashed lines depict $u(a,\pi)$ as a function of π. The solid black lines depict $\overline{u}(\pi)$, the solid blue line depicts cav $\overline{u}(\pi)$ for $\pi \in [0, 3, 0.6]$, and cav $\overline{u}(\pi)$ coincides with $\overline{u}(\pi)$ for $\pi \notin [0, 3, 0.6]$. The red dashed line depicts the payoff to the mixed action $\alpha_p \in \Delta(\{a_L, a_R\})$. The mixed action α_p is the uniform punishment; it is a best response for the DM to belief $\pi_p = 0.4$.

Figure 1: Uniform Punishment and Cross-verification.

We first note that the optimal experiment σ^* consists in splitting the prior into the posteriors posteriors $\pi_{s_0}^* = 0.3$ and $\pi_{s_1}^* = 0.6$.[7] We also note that $u(a_0, \pi_{s_0}) < u(a_1, \pi_{s_0})$ and $u(a_1, \pi_{s_1}) < u(a_0, \pi_{s_1})$, that is, the experts have an incentive to mis-report the realized signals. Thus, if there was a single expert, choosing the experiment σ^* and truthfully reporting the realized signal would not be an equilibrium. More generally, no equilibrium would give the expert his commitment value.

[7] We have $\lambda_{s_0}^* = \lambda_{s_1}^* = 1/2$. The experiment is given by: $\sigma^*(s_0|\omega_0) = 0.64$ and $\sigma^*(s_1|\omega_1) = 0.67$.

Matters are different if the decision-maker chooses to consult another expert. To see this, suppose that the two experts choose the experiment σ^* and truthfully report the outcome of the experiment. The decision-maker then holds belief 0.3 (resp., 0.6) and plays action a_0 (resp., a_1) after observing two matching messages equal to s_0 (resp., s_1). Off the equilibrium path, i.e., when the decision-maker observes two contradictory messages, assume that he holds belief $\pi_p = 0.4$ and plays action $\alpha_p \in \Delta(\{a_L, a_R\}) = BR(0.4)$.

The key observation to make is that the mixed strategy $\alpha_p \in BR(0.4)$ is a *uniform punishment*, that is, $u(\alpha_p, \pi_{s_0}) < u(a_0, \pi_{s_0})$ and $u(\alpha_p, \pi_{s_1}) < u(a_1, \pi_{s_1})$. (See Figure 1.) In words, regardless of the realized signal, an expert is punished for deviating from truth-telling. All the decision-maker needs to know is that a deviation has occurred, and the presence of the second expert indeed guarantees that deviations are detected. The experts thus benefit from the decision-maker cross-verifying their information. (Naturally, there are other equilibria, where the decision-maker benefits from cross-verification. See the next section.)

We conclude with two additional remarks. First, if the experts choose the perfectly informative experiment, truthful reporting does not constitute an equilibrium. This is because the actions that are best for the decision-maker at beliefs $\pi_{s_0} = 0$ and $\pi_{s_1} = 1$ are the worst for the experts at those beliefs. Second, for any two experiments σ_1 and σ_2, there is an equilibrium, where experts 1 and 2 choose experiments σ_1 and σ_2, respectively, and a babbling equilibrium of the ensuing sub-game is played.

We now turn to the proof of Theorem 1. The proof rests on three essential properties. First, as already explained, if the two experts choose the same experiment, their signals' realizations are *perfectly correlated.* Second, if the two experts choose the same experiment, the decision-maker *detects* any deviation from truth-telling. This is because the decision-maker receives contradicting messages after any deviation. However, he cannot identify the deviator and, thus, cannot infer the true signal's realization. Therefore, to deter deviations, the decision-maker must be able to punish the two experts simultaneously. The third property is the existence of such a *uniform punishment* whenever the experiment is expert optimal. The following lemma states this property.

Lemma 1 (Uniform punishment). *Let $(\lambda^*_s, \pi^*_s)_{s \in S}$ be an optimal splitting. There exist $\pi_p \in co\ \Pi^*$ and $\alpha_p \in BR(\pi_p)$ such that $u(\alpha_p, \pi^*_s) \leq \overline{u}(\pi^*_s)$ for all $\pi^*_s \in \Pi^*$.*

Lemma 1 is our main technical contribution. We postpone its proof to the end of this section and now show how to construct an equilibrium of the cheap-talk game with a payoff of $\operatorname{cav} \overline{u}(\pi^\circ)$ to the experts.

Proof of Theorem 1. Let $(\lambda^*_s, \pi^*_s)_{s \in S}$ be an optimal splitting inducing the payoff $\operatorname{cav} \overline{u}(\pi^\circ)$. Let σ^* be the optimal experiment associated with that splitting. Recall that $\Pi^* :=$

$\{\pi_s^* : s \in S\}$. From Lemma 1, there exist $\pi_p \in \operatorname{co}\Pi^*$ and $\alpha_p \in BR(\pi_p)$ such that for all $\pi_s^* \in \Pi^*$, $u(\alpha_p, \pi_s^*) - \overline{u}(\pi_s^*) \leq 0$.

We construct a truthful equilibrium as follows. The experts choose the optimal experiment σ^* and truthfully report the realized signal. Following the choice of σ^*, the decision-maker chooses $\alpha \in BR(\pi_s)$, with $u(\alpha, \pi_s) = \overline{u}(\pi_s)$, when he observes two identical messages equal to s. Alternatively, if the decision-maker receives two conflicting messages, he chooses α_p (sustained by the belief π_p). Finally, following the choice of any other statistical experiment, an equilibrium of the continuation game, which exists by finiteness, is played. It is routine to check that we indeed have an equilibrium. □

We now offer a series of remarks.

Remark 1. Throughout the paper, we have assumed that the set of actions is finite. We have verified that Lemma 1 and Theorem 1 continue to hold if A is a compact subset of $\mathbb{R}^d$, and the functions u and v are continuous in a. A proof is available upon request.

Remark 2. So far, we have assumed that the two experts share the same preferences. If the preferences of one expert, say the second expert, are a convex combination of the preferences of the first expert and the decision-maker, i.e., $\beta u(a, \omega) + (1 - \beta) v(a, \omega)$ for some $\beta \in [0, 1]$, then we can still construct a truthful equilibrium, where the first expert continues to obtain his commitment value. To see this, let a_s be such that $u(a_s, \pi_s^*) = \overline{u}(\pi_s^*)$ and note that $v(\alpha_p, \pi_s^*) \leq v(a_s, \pi_s^*) = \max_{\tilde{\alpha}} v(\tilde{\alpha}, \pi_s^*)$ for all s, where α_p is the uniform punishment, which exists by Lemma 1. This implies that $\beta u(\alpha_p, \pi_s^*) + (1 - \beta) v(\alpha_p, \pi_s^*) \leq \beta u(a_s, \pi_s^*) + (1 - \beta) v(a_s, \pi_s^*)$ for all $\pi_s^* \in \Pi^*$, i.e., α_p is also a uniform punishment for the second expert. We illustrate this remark with a simple example. As in Crawford and Sobel (1982), assume that the decision-maker obtains the payoff $-(\alpha - \omega)^2$, when he chooses $\alpha \in [0, 1]$ and the state is ω. The payoff of the two experts are $-(\alpha - \omega - b)^2$ and $-(\alpha - \omega - \beta b)^2$, with $\beta \in [0, 1]$ and $b > 0$, respectively. The second expert is (weakly) less biased than the first expert. Observe that, up to a constant, the payoff of the less biased expert is a convex combination of the payoff of the most biased expert and the decision-maker, that is:

$$-\left[(1 - \beta)(\alpha - \omega)^2 + \beta(\alpha - \omega - b)^2\right] = -(\alpha - \omega - \beta b)^2 - b^2 \beta (1 - \beta).$$

Therefore, there exists an equilibrium, which gives the most biased expert his commitment value.[8]

[8]In the quadratic example, the payoff $\overline{u}(\pi)$ to the most biased expert is $-(\mathbb{V}_\pi[\omega] + b^2)$, with $\mathbb{V}_\pi[\omega]$ the variance of ω with respect to the distribution π. Since the variance of a real-valued random variable is concave in its distribution, full information disclosure attains the commitment value.

Remark 3. Lemma 1 can also be used to identify a class of preferences under which a fully revealing equilibrium exists in a cheap-talk game played between two *perfectly informed* experts and a decision-maker. Such a cheap-talk game corresponds to stages 2 and 3 of our model, when the experts choose the fully revealing experiment in the first stage. In particular, suppose that the preferences of the first expert are given by u, the optimal experiment with respect to u is fully revealing, and the preferences of the second expert are given by:

$$u_2(a,\omega) = f_\omega(\beta_\omega u(a,\omega) + (1-\beta_\omega)v(a,\omega)),$$

for all (a,ω), for some concave and non-decreasing functions f_ω and scalars $\beta_\omega \in [0,1]$. Let α_p be the uniform punishment with respect to u, which exists by Lemma 1. Under these assumptions, there is a fully revealing equilibrium of the cheap-talk game because $u_2(\alpha_p,\omega) \le u_2(a_\omega,\omega)$ for all ω, with $u(a_\omega,\omega) = \overline{u}(\omega)$, i.e., α_p is also a uniform punishment with respect to u_2. To see this, note that

$$\begin{aligned} u_2(\alpha_p,\omega) &= \sum_{a\in A}\alpha_p(a) f_\omega(\beta_\omega u(a,\omega) + (1-\beta_\omega)v(a,\omega)) \\ &\le f_\omega(\beta_\omega u(\alpha_p,\omega) + (1-\beta_\omega)v(\alpha_p,\omega)) \\ &\le f_\omega(\beta_\omega u(a_\omega,\omega) + (1-\beta_\omega)v(a_\omega,\omega)) = u_2(a_\omega,\omega), \end{aligned}$$

for each $\omega \in \Omega$, where the first inequality follows from the concavity of f_ω and the second from the non-decreasingness.

A prominent example of a cheap-talk game with multiple experts, where the optimal experiment is fully revealing, is Ambrus and Takahashi (2008). In Ambrus and Takahashi (2008), the decision-maker's preferences are given by $-||a-\omega||^2$, where $||\cdot||$ is the Euclidian norm, while the experts' preferences are given by $-||a-\omega-b||^2$ and $-||a-\omega-\beta_\omega b||^2$, respectively, with $\beta_\omega \in [0,1]$ and $b \in \mathbb{R}^d$. The second expert is (weakly) less biased in the same direction as the first expert. The set Ω is a finite subset of $\mathbb{R}^d$ and A is a subset of $\mathbb{R}^d$ that contains the convex hull of Ω.

In this game, the optimal experiment for the most biased expert is well-known to be fully revealing. Intuitively, the expert aims to minimize the expected variance of the state, which is attained by fully disclosing the state.[9] Moreover, at each state, the payoff of the less biased expert is a convex combination of the payoff of the most biased

[9]To see that the optimal experiment for the most biased expert is fully revealing, note that the payoff $\overline{u}(\pi)$ to the most biased expert is $-(\sum_{j=1}^d \mathbb{V}_\pi[\omega^j] + \sum_{j=1}^d (b^j)^2)$, where ω^j and b^j are the jth components of the vectors ω and b, respectively, and $\mathbb{V}_\pi[\omega^j]$ is the variance of ω^j with respect to the distribution π. Since the variance of a real-valued random variable is concave in its distribution, $\overline{u}(\pi)$ is a convex function.

expert and the decision-maker up to a constant, that is:

$$-\left[(1-\beta_\omega)||a-\omega||^2+\beta_\omega||a-\omega-b||^2\right] = -\sum_{j=1}^{d}(\alpha^j-\omega^j-\beta_\omega b^j)^2+(b^j)^2\beta_\omega(1-\beta_\omega)$$
$$= -||a-\omega-\beta_\omega b||^2+\sum_{j=1}^{d}(b^j)^2\beta_\omega(1-\beta_\omega),$$

where a^j, ω^j, and b^j are the j-th components of the vectors a, ω and b, respectively. Therefore, the above argument implies the existence of a uniform punishment and this uniform punishment can be used to construct a fully revealing equilibrium. The existence of a fully revealing equilibrium already appears in Ambrus and Takahashi (2008).

Remark 4. We have assumed that the choice of experiments is publicly observed. If the decision-maker does not observe the experiments chosen by the two experts, but if the experts observe each other's experiment choice, then again there is a truthful equilibrium, where the optimal experiment σ^* is chosen as in Theorem 1. In this equilibrium, play on the equilibrium path unfolds as in Theorem 1. If any expert deviates and chooses another experiment $\sigma \neq \sigma^*$, then the two experts send the message m_0, where m_0 is a message that is never sent on the equilibrium path. If the decision-maker observes two messages that do not match or observes a message equal to m_0 from either of the two experts, then he best responds to belief π_p by playing action α_p.

Remark 5. Similarly, if we assume that the experts do not observe each other's choice of experiments, but the decision-maker does, then our result continues to hold. To see this, we construct an equilibrium as follows. In the first stage, the experts choose the optimal experiment. In the second stage, an expert truthfully reports his signal if he has chosen the optimal experiment in the first stage. (The strategies are left unspecified in other contingencies.) If the decision-maker observes the experts choosing the optimal experiment, he follows the same strategy as in our main proof. If the decision-maker observes only one expert choosing the optimal experiment, he plays a best-reply to the message sent by that expert. (The strategies are left unspecified in all other contingencies.) On path, the experts receive their commitment value. If an expert chooses another experiment, the decision-maker observes the deviation but not the other expert, who continues to truthfully reveal the signal. Hence, the deviation does not change the expert's payoff.

Remark 6. We have assumed that the two experts choose experiments simultaneously. This assumption is again not required for our result. Suppose instead that one expert, say the first expert, chooses an experiment $\sigma : \Omega \to \Delta(S_1 \times S_2)$, with expert i privately observing the signal's realization s_i. As before, after observing their signals, the experts send messages to the decision-maker, who then chooses an action. Yet again, we have a truthful equilibrium, where the equilibrium payoff of the two experts is $\mathrm{cav}\,\overline{u}(\pi^\circ)$ as in Theorem 1. In this equilibrium, the first expert chooses the optimal experiment and perfectly correlates the second expert's signal with his own.

Remark 7. We used weak perfect Bayesian equilibrium as our solution concept. If we restrict attention to a finite set of experiments, which contains σ^*, then we can strengthen the solution concept to sequential equilibrium. We only need a slight modification of Lemma 1 that ensures that the decision-maker believes the realized signal is either s or s' after observing report (s, s'). A minor adaptation of the proof of Lemma 1 shows that there is a belief $\pi_{s,s'} \in \Delta(\{\pi^*_s, \pi^*_{s'}\})$ and a mixed action $\alpha_{s,s'} \in BR(\pi_{s,s'})$ such that $u(\alpha_{s,s'}, \pi^*_{\tilde{s}}) - \overline{u}(\pi^*_{\tilde{s}}) \leq 0$ for all $\tilde{s} \in \{s, s'\}$.

Proof of Lemma 1. We first establish two intermediate claims, then we use these two claims to establish the lemma. Let $(\lambda^*_s, \pi^*_s)_{s\in S}$ be an optimal splitting. Recall that $\Pi^* := \{\pi^*_s : s \in S\}$ and Δ^* is the set of all probability distributions over Π^*.

Claim 1: For any $\lambda \in \Delta^*$, $\overline{u}(\sum_s \lambda_s \pi^*_s) \leq \sum_s \lambda_s \overline{u}(\pi^*_s)$.

Proof of Claim 1: Consider the convex hull of the graph of $\overline{u}$, i.e., $\mathrm{co}\,\{(\pi, r) \in \Delta(\Omega) \times \mathbb{R} : r = \overline{u}(\pi)\}$. By construction, the point $(\pi^\circ, \mathrm{cav}\,\overline{u}(\pi^\circ)) = (\sum_s \lambda^*_s \pi^*_s, \sum_s \lambda^*_s \overline{u}(\pi^*_s))$ is on the boundary of the convex hull. From the supporting hyperplane theorem, there exists a hyperplane $h \in \mathbb{R}^{|\Omega|} \times \mathbb{R}$ supporting $\mathrm{co}\,\{(\pi, r) \in \Delta(\Omega) \times \mathbb{R} : r = \overline{u}(\pi)\}$ at $(\pi^\circ, \mathrm{cav}\,\overline{u}(\pi^\circ))$ such that the graph of $\overline{u}$ lies below h. For all $s \in S$, the point $(\pi^*_s, \overline{u}(\pi^*_s))$ also lies on the hyperplane h. Consequently, the point $(\sum_s \lambda_s \pi^*_s, \sum_s \lambda_s \overline{u}(\pi^*_s))$, must also lies on the hyperplane. Therefore, $\overline{u}(\sum_s \lambda_s \pi^*_s) \leq \sum_s \lambda_s \overline{u}(\pi^*_s)$ as required. ■

Claim 2: Choose any non-empty subset $B \subset A$ and $\varepsilon > 0$. If $\max_{s\in S}[u(\alpha, \pi^*_s) - \overline{u}(\pi^*_s)] \geq \varepsilon$ for each $\alpha \in \Delta(B)$, then there exists $\hat{\lambda} \in \Delta^*$ such that
$\min_{\alpha\in\Delta(B)} u(\alpha, \sum_s \hat{\lambda}_s \pi^*_s) \geq \sum_s \hat{\lambda}_s \overline{u}(\pi^*_s) + \varepsilon$.

Proof of Claim 2: The claim follows from duality. Consider the following linear program:

$$\min_{(x,\alpha)\in\mathbb{R}\times\Delta(B)} x$$

subject to: for all $s \in S$,

$$\sum_{a\in B} \alpha(a) \left[u(a, \pi^*_s) - \overline{u}(\pi^*_s)\right] \leq x.$$

This minimization problem has a solution $\hat{x}$. Our hypothesis implies that $\hat{x} \geq \varepsilon$. The dual program is given by

$$\max_{(y,\lambda)\in\mathbb{R}\times\Delta(\Pi^*)} y$$

subject to: for all $a \in B$,

$$\sum_{s\in S} \lambda_s \left[u(a,\pi_s^*) - \overline{u}(\pi_s^*)\right] \geq y.$$

Since the primal linear program has a solution, the dual program also has a solution $(\hat{y},\hat{\lambda})$. No duality gap further implies that $\hat{y} = \hat{x} \geq \varepsilon$. (See Section 4.2 of Luenberger and Ye, 2008.) Therefore, for all $a \in B$,

$$u(a, \sum_s \hat{\lambda}_s \pi_s^*) = \sum_{s\in S} \hat{\lambda}_s u(a,\pi_s^*) \geq \varepsilon + \sum_{s\in S} \hat{\lambda}_s \overline{u}(\pi_s^*)$$

Hence, $u(\alpha, \sum_s \hat{\lambda}_s \pi_s^*) \geq \sum_s \hat{\lambda}_s \overline{u}(\pi_s^*) + \varepsilon$ for all $\alpha \in \Delta(B)$, as required. ■

We now use Claims 1 and 2 to complete the proof. Recall that $br(\pi) \subset A$ is the decision-maker's pure best-replies to belief π.

By contradiction, assume that there does not exist $\pi_p \in \operatorname{co}\Pi^*$ and $\alpha_p \in BR(\pi_p)$ such that $u(\alpha_p, \pi_s^*) - \overline{u}(\pi_s^*) \leq 0$ for all $\pi_s^* \in \Pi^*$. Note that $\pi \in \operatorname{co}\Pi^*$ if and only if $\pi = \sum_s \lambda_s \pi_s^*$ for some $\lambda \in \Delta^*$. Hence, our contradiction hypothesis can be restated as follows: for each $\lambda \in \Delta^*$, there exists $\varepsilon(\lambda) > 0$ such that $\max_{s\in S}[u(\alpha,\pi_s^*) - \overline{u}(\pi_s^*)] \geq \varepsilon(\lambda)$ for each $\alpha \in \Delta(br(\sum_s \lambda_s \pi_s^*)) = BR(\sum_s \lambda_s \pi_s^*)$. Let $\varepsilon := \min_{\lambda\in\Delta^*} \varepsilon(\lambda)$. Note that $\varepsilon > 0$ because $\varepsilon(\lambda)$ depends only on the finite set $br(\sum_s \lambda_s \pi_s^*)$, and there are finitely many such subsets of A.

Define the correspondence $F : \Delta^* \to \Delta^*$, with

$$F(\lambda) := \Big\{\lambda' \in \Delta^* : \min_{\alpha \in BR(\sum_s \lambda_s \pi_s^*)} \sum_s \lambda'_s \Big(u(\alpha,\pi_s^*) - \overline{u}(\pi_s^*)\Big) \geq \varepsilon\Big\}.$$

We can readily check that this correspondence is convex and compact valued. We argue below that it is non-empty valued and lower hemi-continuous. Hence, the correspondence has a fixed point $\overline{\lambda} \in F(\overline{\lambda})$ by Theorem 15.4 in Border (1990). Noting that $\sum_s \overline{\lambda}(s) u(\alpha,\pi_s^*) = u(\alpha, \sum_s \overline{\lambda}(s)\pi_s^*)$, we find

$$\min_{\alpha\in BR(\sum_s \overline{\lambda}(s)\pi_s^*)} \Big(u(\alpha, \sum_s \overline{\lambda}(s)\pi_s^*) - \sum_s \overline{\lambda}(s)\overline{u}(\pi_s^*)\Big) \geq \varepsilon$$

for $\overline{\lambda} \in \Delta^*$ contradicting Claim 1 and establishing the result.

We now show that the correspondence is non-empty valued. Pick any $\lambda \in \Delta^*$. The contradiction hypothesis states that $\max_{s\in S}[u(\alpha,\pi_s^*) - \overline{u}(\pi_s^*)] \geq \varepsilon$ for each $\alpha \in$

$BR(\sum_s \lambda_s \pi_s^*)$. Claim 2 then implies that there exists $\hat{\lambda} \in \Delta^*$ such that

$$\min_{\alpha \in BR(\sum_s \lambda_s \pi_s^*)} \sum_s \hat{\lambda}_s (u(\alpha, \pi_s^*) - \overline{u}(\pi_s^*)) \geq \varepsilon,$$

i.e., the correspondence is non-empty valued.

Finally, we prove lower hemi-continuity. Pick an open set $O \subseteq \Delta^*$ such that $F(\lambda) \cap O \neq \emptyset$. Since BR is upper hemi-continuous (by the maximum principle) and A is finite, there exists a neighborhood O' of λ such that $BR(\sum_s \lambda'_s \pi_s^*) \subseteq BR(\sum_s \lambda_s \pi_s^*)$ for all $\lambda' \in O'$. Therefore, for all $\lambda' \in O'$,

$$\min_{\alpha \in BR(\sum_s \lambda'_s \pi_s^*)} \sum_s \lambda''_s \Big(u(\alpha, \pi_s^*) - \overline{u}(\pi_s^*)\Big) \geq \min_{\alpha \in BR(\sum_s \lambda_s \pi_s^*)} \sum_s \lambda''_s \Big(u(\alpha, \pi_s^*) - \overline{u}(\pi_s^*)\Big) \geq \varepsilon$$

for any $\lambda'' \in F(\lambda) \cap O$ because $BR(\sum_s \lambda'_s \pi_s^*) \subseteq BR(\sum_s \lambda_s \pi_s^*)$, i.e., $\lambda'' \in F(\lambda')$. Hence, $F(\lambda') \cap O \neq \emptyset$ for all $\lambda' \in O'$, which proves the lower hemi-continuity of F (Definition 11.3 in Border (1990)). $\square$

4 The Decision-maker and Cross-verification

The previous section showed that the experts benefit from the decision-maker cross-verifying their information. This section explores whether the decision-maker can also benefit from cross-verification.

We begin with some definitions. Fix a cheap-talk game $\Gamma(\pi^\circ, u, v)$. We say that the experts benefit from persuasion if $\text{cav}\, \overline{u}(\pi^\circ) > \overline{u}(\pi^\circ)$. Similarly, we say that the decision-maker benefits from cross-verification if there exists an equilibrium of the cheap-talk game, where the decision-maker's payoff exceeds the ex-ante payoff $\max_{a \in A} v(a, \pi^\circ)$. Notice that if the decision-maker benefits from cross-verification, the experts must reveal some information to the decision-maker.

Define $\hat{A} := \{a \in A : \exists \pi \in \Delta(\Omega) \text{ s.t. } a \in BR(\pi)\}$ and $\overline{v}(\pi) := \max_\alpha v(\alpha, \pi)$ for all $\pi \in \Delta(\Omega)$. We say that there are *no redundant actions* for the decision-maker if for any non-empty $B \subset \hat{A}$ and $B \neq \hat{A}$, there exists $\pi \in \Delta(\Omega)$ such that $\overline{v}(\pi) > \max_{a \in B} v(a, \pi)$. There are no redundant actions for the experts if there are no two distinct actions a and a' such that $u(a, \omega) = u(a', \omega)$ for all $\omega \in \Omega$.

Remark 8. The conditions of non-redundancy are generic. Moreover, the condition of no redundant actions for the decision-maker does not preclude strictly dominated actions. Two important implications of that condition are as follows: (i) the set $BR^{-1}(a) := \{\pi \in \Delta(\Omega) : v(a, \pi) = \overline{v}(\pi)\}$ has full dimension (as a subset of the simplex of dimension

$|\Omega| - 1)$, and (ii) no action other than a is optimal in the relative interior of $BR^{-1}(a)$, denoted by $\text{int}\, BR^{-1}(a)$.

Theorem 1 showed that the experts benefit from cross-verification in games where they benefit from persuasion. The following proposition further establishes that the decision-maker also benefits from cross-verification in such games.

Proposition 1. *Assume that there are no redundant actions for the decision-maker in the game $\Gamma(\pi^\circ, u, v)$. At almost all priors π°, if the experts benefit from persuasion, then the decision-maker benefits from cross-verification.*

We first illustrate the logic of the proposition with the help of a simple example. There are two states, ω_0 and ω_1, and three actions, a_0, a_1 and a_p. Throughout the example, probabilities refer to the probability of ω_1. The prior is $\pi^\circ = 0.45$. The payoffs are illustrated in Figure 2. The optimal experiment consists in splitting the prior into the posteriors $\pi^*_{s_0} = 0.3$ and $\pi^*_{s_1} = 0.6$. The experts strictly benefit from persuasion. From Theorem 1, there exists a truthful equilibrium, where an expert's payoff is his commitment value. Action a_p is the uniform punishment sustaining the equilibrium. Note that a_p is uniquely optimal at the prior and also optimal at the two posteriors. Consequently, the decision-maker does not benefit from cross-verification at the equilibrium. Yet, we can construct another equilibrium, where the decision-maker benefits from cross-verification. To see this, consider the splitting of the prior into $\pi_{s_0} = 0.2$ and $\pi_{s_1} = 0.8$. At π_{s_0} (resp., π_{s_1}), the decision-maker plays a_0 (resp., a_1). To sustain this splitting as an equilibrium, the decision-maker punishes the experts with a_p. The decision-maker strictly benefits from this more informative experiment.

We prove that the logic of the example generalizes to almost all priors. That is, for all priors, but for a subset with Lebesgue measure zero, we can always construct an equilibrium of the cheap-talk game, where the decision-maker benefits from cross-verification if the experts benefit from persuasion. More precisely, we prove that the proposition holds at all interior priors, where the decision-maker has at most two best-replies, a generic condition.

The need for non-redundancy is clear. If the decision-maker is indifferent between all his actions, the decision-maker cannot benefit from cross-verification, while the experts can benefit from persuasion. We now turn to the proof.

Proof of Proposition 1. Consider an optimal splitting $(\lambda^*_s, \pi^*_s)_{s \in S}$ of π°, which induces the value $\text{cav}\, \overline{u}(\pi^\circ)$, where $\text{cav}\, \overline{u}(\pi^\circ) > \overline{u}(\pi^\circ)$. Without loss of generality, assume that $\lambda^*_s > 0$ for all $s \in S$. Let $\overline{v}(\pi) := \max_\alpha v(\alpha, \pi)$ for all $\pi \in \Delta(\Omega)$.

If the decision-maker benefits from the statistical experiment, there is nothing to prove. So, assume that the decision-maker does not benefit from the statistical ex-

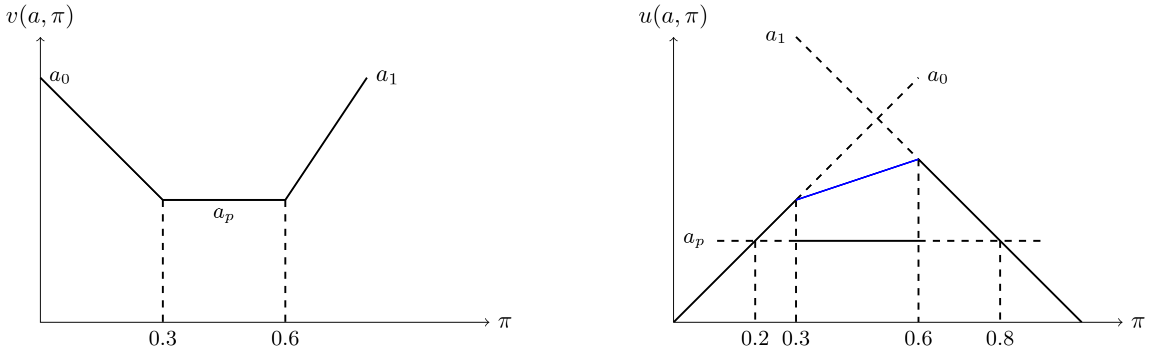

(a) DM's Preferences. Action a_0 is optimal for the DM for $\pi \in [0, 0.3]$, action a_p is optimal for $\pi \in [0.3, 0.6]$, and a_1 is optimal for $[0.6, 1]$.

(b) Expert's Preferences. The solid black lines depict $\overline{u}(\pi)$, the solid blue line depicts cav $\overline{u}(\pi)$ for $\pi \in [0, 3, 0.6]$, and cav $\overline{u}(\pi)$ coincides with $\overline{u}(\pi)$ for $\pi \notin [0, 3, 0.6]$.

Figure 2: DM Benefits from Cross-verification.

periment, i.e., $\sum_s \lambda_s^* \overline{v}(\pi_s^*) = \overline{v}(\pi^\circ)$. We construct another equilibrium at which the decision-maker benefits from cross-verification.

We first claim that for all $a \in br(\pi^\circ)$, $a \in br(\pi)$ for all $\pi \in \text{co}\{\pi_s^* : s \in S\}$. To see this, consider any $a \in br(\pi^\circ)$ and observe that

$$\sum_s \lambda_s^* \overline{v}(\pi_s^*) = \overline{v}(\pi^\circ) = v(a, \pi^\circ) = v\left(a, \sum_s \lambda_s^* \pi_s^*\right) = \sum_s \lambda_s^* v(a, \pi_s^*).$$

It follows that

$$\sum_s \underbrace{\lambda_s^*}_{>0} (\underbrace{\overline{v}(\pi_s^*) - v(a, \pi_s^*)}_{\geq 0}) = 0.$$

If there exists s such that $\overline{v}(\pi_s^*) > v(a, \pi_s^*)$, we have a contradiction. Hence, $a \in br(\pi_s^*)$ for all s and, consequently, $a \in br(\pi)$ for all $\pi \in \text{co}\{\pi_s^* : s \in S\}$.

From the definition of $\overline{u}$, we have that $u(a, \pi_s^*) \leq \overline{u}(\pi_s^*)$ for all s, for all $a \in br(\pi^\circ)$, since $BR(\pi^\circ) \subseteq BR(\pi_s^*)$ for all s. We now argue that for all $a \in br(\pi^\circ)$, there exists $s_a \in S$ such that $u(a, \pi_{s_a}^*) < \overline{u}(\pi_{s_a}^*)$. Choose any $a \in br(\pi^\circ)$. To the contrary, assume that $u(a, \pi_s^*) = \overline{u}(\pi_s^*)$ for all s. We then have

$$\text{cav}\, \overline{u}(\pi^\circ) = \sum_s \lambda_s^* \overline{u}(\pi_s^*) = \sum_s \lambda_s^* u(a, \pi_s^*) = u(a, \pi^\circ) \leq \overline{u}(\pi^\circ) \leq \text{cav}\, \overline{u}(\pi^\circ),$$

a contradiction with the expert benefiting from the experiment.

To sum up, we have (i) $BR(\pi^\circ) \subseteq BR(\pi)$ for all $\pi \in \text{co}\{\pi_s^* : s \in S\}$, and (ii) for each $a \in br(\pi^\circ)$, there exists s_a such that $u(a^*_{s_a}, \pi^*_{s_a}) > u(a, \pi^*_{s_a})$ with $a^*_{s_a} \in br(\pi^*_{s_a})$ satisfying $u(a^*_{s_a}, \pi^*_{s_a}) = \overline{u}(\pi^*_{s_a})$.

The fact that u is continuous in π and $u(a, \pi^*_{s_a}) < u(a^*_{s_a}, \pi^*_{s_a})$ for each $a \in br(\pi^\circ)$ together imply that there exists $\epsilon > 0$ and an open ball $\mathcal{O} = \{\pi \in \Delta(\Omega) : ||\pi - \pi^*_{s_a}|| < \varepsilon\}$ such that $u(a, \pi) < u(a^*_{s_a}, \pi)$ for all $a \in br(\pi^\circ)$ and all $\pi \in \mathcal{O}$.

We claim that $\mathcal{O}$ intersects the relative interior of $br^{-1}(a^*_{s_a})$. To see this, note that $\mathcal{O} \cap br^{-1}(a^*_{s_a}) \neq \emptyset$ since $\pi^*_{s_a}$ is an element of both $\mathcal{O}$ and $br^{-1}(a^*_{s_a})$. Moreover, it follows from the non-redundancy of A that $\pi^*_{s_a}$ is not in the relative interior of $br^{-1}(a^*_{s_a})$ since any $a \in br(\pi^\circ)$ is also optimal at $\pi^*_{s_a}$. Since the relative interior of $br^{-1}(a^*_{s_a})$ is non-empty, there exists π^{**} in the relative interior such that the half-open line segment $[\pi^{**}, \pi^*_{s_a})$ is contained in the relative interior. (See Theorem 2.1.3 and Lemma 2.1.6 in Hiriart-Urruty and Lemaréchal.) Therefore, there exists $\overline{\pi}_a$ in the intersection of the relative interior of $br^{-1}(a^*_{s_a})$ and $\mathcal{O}$, i.e., such that $u(a, \overline{\pi}_a) < u(a^*_s, \overline{\pi}_a) = \overline{u}(\overline{\pi}_a)$. Note that $v(a^*_{s_a}, \overline{\pi}_a) > v(a, \overline{\pi}_a)$ since $a^*_{s_a}$ is uniquely optimal at $\overline{\pi}$. In other words, there is an element of $br(\pi^\circ)$, namely a, which is not an element of $br(\overline{\pi}_a)$.

The last step consists in showing that there exists $a \in br(\pi^\circ)$ and $\underline{\pi}_a \in br^{-1}(a)$ such that the open segment $(\underline{\pi}_a, \overline{\pi}_a)$ includes π°. Indeed, if such an open segment exists, we have a splitting $(\underline{\pi}_a, \overline{\pi}_a)$ of π° such that $\overline{u}(\underline{\pi}_a) \geq u(a, \underline{\pi}_a)$, $\overline{u}(\overline{\pi}_a) = u(a^*_{s_a}, \overline{\pi}_a) > u(a, \overline{\pi}_a)$. This splitting can be supported as a truthful equilibrium (with a as the punishment at belief π°). Moreover, since $v(a^*_{s_a}, \overline{\pi}_a) > v(a, \overline{\pi}_a)$, the decision-maker strictly benefits, the desired contradiction.

Finally, suppose that π° is in the interior of the simplex. If $br(\pi^\circ) = \{a\}$, then π° is in the relative interior of $br^{-1}(a)$. Thus, we can trivially find a segment with the required property.

If $br(\pi^\circ) = \{a, b\}$ and $s_a = s_b$, then the same arguments apply, since the open segment will intersect either $br^{-1}(a)$ or $br^{-1}(b)$. If $s_a \neq s_b$, choose $\overline{\pi}_{s_a}$ such that b is uniquely optimal at $\overline{\pi}_{s_a}$. Such $\overline{\pi}_a$ exists since $v(b, \pi_{s_a}) = \max_{a' \in br(\pi_{s_a})} v(a', \pi_{s_a})$ (if not $s_a = s_b$). As before, the open segment intersects either $br^{-1}(a)$ or $br^{-1}(b)$. However, it cannot be $br^{-1}(b)$. If it were, b would be uniquely optimal at $\overline{\pi}_{s_a}$ and optimal at π° and $\underline{\pi}_{s_a}$, which is not possible since $br^{-1}(b)$ is convex.

Since the set of interior priors with at most two best-replies is generic, the proof is complete. □

Proposition 1 does not generalize to all priors. For a counter-example, consider Figure 3. There are three states, ω_0, ω_1 and ω_2, and two actions, a and b. The action a (resp., b) is optimal in the left triangle marked "a" (resp., in the right triangle marked "b"). At the prior π°, the action a is the unique best-reply of the decision-maker.

Assume that $u(b, \omega_1) > u(a, \omega_1)$. Thus, if the experts truthfully reveal the state, they benefit from persuasion, while the decision-maker does not.[10]

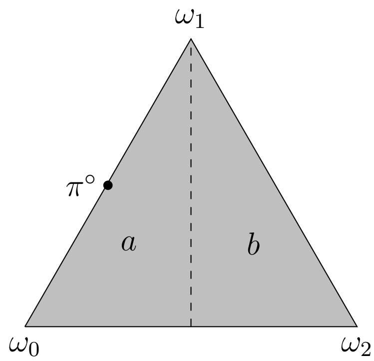

Figure 3: A counter-example

Proposition 1 proved that the decision-maker benefits from cross-verification whenever the experts benefit from persuasion. We now show a partial converse, that is, the decision-maker benefits from cross-verification only when the experts benefit from persuasion.

Proposition 2. *Assume that there are no redundant actions for the experts and the decision-maker in the game $\Gamma(\pi^\circ, u, v)$. If $\overline{u}$ is a concave function, then the decision-maker does not benefit from cross-verification. That is, in all equilibria of $\Gamma(\pi^\circ, u, v)$, the decision-maker's payoff is $\overline{v}(\pi^\circ)$.*

To understand Proposition 2, assume that the experts and the decision-maker have opposing preferences, that is, $u = -v$. In this case, what is best for the decision-maker is worst for the experts, and therefore, $\overline{v} = -\overline{u}$. Moreover, if either of the experts, say expert 1, chooses a uninformative experiment, an expert's payoff is $u(\pi^\circ)$ in all equilibria of the ensuing game. This is because if expert 2's experiment produces two signals s and s' such that the set of best-replies at π_s differs from the set of best-replies at s', then expert 2 has an incentive to misreport one of the two signals, if not both. The experts cannot credibly communicate any information. Therefore, no expert can obtain less than $\overline{u}(\pi^\circ)$ in equilibrium. Experts cannot obtain more than $\overline{u}(\pi^\circ)$ either. Indeed, for every on-path posterior π, the decision-maker chooses a best-reply in equilibrium, hence an expert's payoff is minimized at π, i.e., an expert's payoff is $u^{\min}(\pi) := \min_a u(a, \pi)$. The result then follows from the concavity of $u^{\min}$. Proposition 2 does not require opposing preferences; the logic outlined above extends to all games, where $\overline{u}$ is concave.

[10] If there are two states, Proposition 1 generalizes to all interior priors. In this case, non-redundancy of the decision-maker's payoff implies that the decision-maker has at most two best-replies at each belief, where we know that Proposition 1 holds. In general, however, we do not know whether the proposition generalizes to all interior priors.

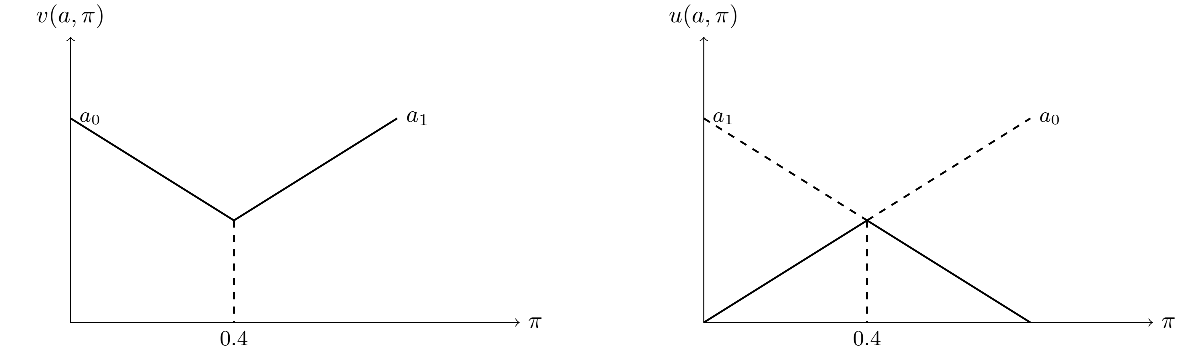

(a) DM's Preferences. Action a_0 is optimal for the DM for $\pi \in [0, 0.4]$ and a_1 is optimal for $[0.4, 1]$.

(b) Expert's Preferences. The solid black lines depict the concave function $\overline{u}(\pi)$.

Figure 4: No Benefit from Cross-verification or Persuasion.

To further illustrate Proposition 2, consider Figure 4. For the decision-maker to benefit from cross-verification, the experts would need to choose an experiment, which induces the decision-maker to play different actions after receiving different signals. However, we cannot sustain such a choice as an equilibrium. An expert would always have an incentive to misreport the realized signal. This is because any action other than the one chosen by the decision-maker improves an expert's payoff, i.e., there is no uniform punishment.

The need for the non-redundancy of the experts' actions is again clear. If the experts are totally indifferent, they cannot benefit from persuasion but can provide the decision-maker with perfectly informative signals. It remains to prove Proposition 2. We do so by establishing a series of lemmata under the assumptions of Proposition 2. The following lemma shows that the conflict of interest between the experts and the decision-maker is maximal when the experts cannot benefit from persuasion; that is, the decision-maker's best-replies at belief π minimizes the experts' expected payoff. Recall that $\hat{A}$ is the set of actions that are a best response for the decision-maker to some belief.

Lemma 2. *For every* $\pi \in \Delta(\Omega)$, $BR(\pi) = \arg\min_{\alpha' \in \Delta(\hat{A})} u(\alpha', \pi)$.

Proof of Lemma 2. We start by proving the following claim.

Claim 1. $a \in \hat{A}$ *and* $\pi \in \operatorname{int} br^{-1}(a)$ *implies* $\{a\} = \arg\min_{\alpha' \in \Delta(\hat{A})} u(\alpha', \pi)$.

Proof of Claim 1. Fix $a \in \hat{A}$ and $\pi \in \operatorname{int} br^{-1}(a)$. We first argue that there does not exist $a' \in \hat{A}$ such that such that $u(a', \pi) < u(a, \pi)$. To the contrary, suppose such a' exists. Pick an arbitrary $\pi' \in \operatorname{int} br^{-1}(a')$. There exists $\pi'' \in \operatorname{int} br^{-1}(a')$ and $\lambda \in (0, 1)$

such that $\pi'' = \lambda\pi + (1-\lambda)\pi'$. We obtain

$$\begin{aligned} u(a',\pi'') &= \overline{u}(\pi'') \\ &\geq \lambda\overline{u}(\pi) + (1-\lambda)\overline{u}(\pi') \\ &= \lambda u(a,\pi) + (1-\lambda)u(a',\pi') \\ &> \lambda u(a',\pi) + (1-\lambda)u(a',\pi') \\ &= u(a',\pi''), \end{aligned}$$

where the first inequality follows from the concavity of $\overline{u}$, the desired contradiction.

We now argue that there does not exist $a' \in \hat{A}$ such that $u(a',\pi) = u(a,\pi)$. From the above, for all $\pi_n \in \operatorname{int} br^{-1}(a)$, $u(a',\pi_n) \geq u(a,\pi_n)$. Consider any convex combination $(\lambda_n,\pi_n)_n$ satisfying $\sum_n \lambda_n\pi_n = \pi$, $\pi_n \in \operatorname{int} br^{-1}(a)$ for all n, $\lambda_n > 0$ for all n, and the π_n being linearly independent. Such a convex combination exists since $br^{-1}(a)$ has full dimension. If $u(a',\pi) = u(a,\pi)$, then

$$u(a',\pi) = \sum_n \lambda_n u(a',\pi_n) \geq \sum_n \lambda_n u(a,\pi_n) = u(a,\pi) = u(a',\pi),$$

i.e., $u(a',\pi_n) = u(a,\pi_n)$ for all n, a contradiction with the condition of no redundant actions for the experts. Therefore, for all $a' \neq a$, $u(a',\pi) > u(a,\pi)$, which completes the proof of the claim. □

From Claim 1, the statement is true for all π such that $\pi \in \operatorname{int} br^{-1}(a)$ for some $a \in \hat{A}$. Since BR and $\arg\min_{\alpha' \in \hat{A}} u(\alpha',\pi)$ are upper hemi-continuous correspondences, which coincide almost everywhere (in Lebesque measure), they coincide everywhere. □

We now derive an immediate implication of Lemma 2. We first introduce some additional notation. Recall that following the choice of experiments (σ_1,σ_2), we have a proper sub-game. We are interested in analyzing the play in these sub-games. To ease notation, we drop the dependence on (σ_1,σ_2) and write $\pi(m_1,m_2) \in \Delta(\Omega)$ for the decision-maker's belief after observing the messages (m_1,m_2). Similarly, we write $\alpha(m_1,m_2)$ for the decision-maker's equilibrium reply. Notice that $\alpha(m_1,m_2) \in \Delta(\hat{A})$ because this action is a best response to belief $\pi(m_1,m_2)$. Finally, let $\mathbb{P}$ denote the probability distribution over signals, messages and actions induced by the prior and the strategy profile, conditional on the experiments (σ_1,σ_2). At an equilibrium, sequential rationality requires the decision-maker to choose a best-reply to his belief. Fix an equilibrium, an on-path profile of messages (m_1,m_2), and its associated belief $\pi(m_1,m_2)$. Since all best-replies of the decision-maker to $\pi(m_1,m_2)$ minimize the experts' payoffs, no expert must be able to induce the decision-maker to choose an action outside $BR(\pi(m_1,m_2))$ by changing his message to m_1'.

Lemma 3. *If $\mathbb{P}(m_i, m_j) > 0$, then for all m_i', $\alpha(m_i', m_j) \in BR(\pi(m_i, m_j))$.*

Proof of Lemma 3. Without loss of generality, let $i = 1$, $j = 2$. The proof is by contradiction. Assume that there exists m_1, m_1', m_2' such that $\alpha(m_1', m_2') \notin BR(\pi(m_1, m_2'))$.

From Lemma 2, $u(\alpha(m_1', m_2'), \pi(m_1, m_2')) > u(\alpha(m_1, m_2'), \pi(m_1, m_2'))$. The equilibrium payoff to expert 1 is

$$\sum_{(\tilde{m}_1, \tilde{m}_2)} \mathbb{P}(\tilde{m}_1, \tilde{m}_2) u(\alpha(\tilde{m}_1, \tilde{m}_2), \pi(\tilde{m}_1, \tilde{m}_2)).$$

If expert 1 deviates by always sending the message m_1', his expected payoff is:

$$\sum_{(\tilde{m}_1, \tilde{m}_2)} \mathbb{P}(\tilde{m}_1, \tilde{m}_2) u(\alpha(m_1', \tilde{m}_2), \pi(\tilde{m}_1, \tilde{m}_2)).$$

We now argue that the deviation is profitable, the required contradiction.

From Lemma 2, we have that $u(\alpha(\tilde{m}_1, \tilde{m}_2), \pi(\tilde{m}_1, \tilde{m}_2)) \leq u(\alpha(m_1', \tilde{m}_2), \pi(\tilde{m}_1, \tilde{m}_2))$ for all $(\tilde{m}_1, \tilde{m}_2)$. Moreover, there exists (m_1, m_2) such that the inequality is strict and $\mathbb{P}(m_1, m_2) > 0$. Thus, the deviation is profitable. □

The next lemma shows that if any expert chooses an uninformative experiment, then the experts' and the decision-maker's payoff in the ensuing equilibrium is equal to their payoff at their prior belief.

Lemma 4. *Let (σ_1, σ_2) be a profile of experiments. If either σ_1 or σ_2 is an uninformative experiment, then the experts' equilibrium payoff is $\overline{u}(\pi^\circ)$ and the decision-maker's equilibrium payoff is $\overline{v}(\pi^\circ)$ in the ensuing sub-game.*

Proof of Lemma 4. Without loss of generality, assume that σ_2 is uninformative. Since the experiments are observed by the decision-maker, this implies that $\pi(m_1, m_2)$ is independent of m_2. (Recall that we require the beliefs to be consistent with the experiments.) To ease the notation, we drop the dependence on m_2.

Together with Lemma 3, this implies that for all (m_1, m_2) such that $\mathbb{P}(m_1, m_2) > 0$, $\alpha(m_1', m_2) \in BR(\pi(m_1))$ for all m_1'. That is, $\alpha(m_1', m_2)$ is a best-reply to all posterior beliefs $\pi(m_1)$. Note that since $\mathbb{P}(m_1, m_2) > 0$, the message m_1 has strictly positive probability. It follows that $\alpha(m_1', m_2)$ is a best-reply to π° (as the prior is a convex combinations of the posteriors). Since it is true for all (m_1', m_2), the decision-maker payoff is $\overline{v}(\pi^\circ)$.

Finally, since Lemma 2 states that the experts are indifferent among all best-replies of the decision-makers, an expert's payoff is $\overline{u}(\pi^\circ)$. □

We now conclude the proof.

Lemma 5. *In any equilibrium of the cheap-talk game, the experts' payoff is $\overline{u}(\pi^\circ)$, and the decision-maker's payoff is $\overline{v}(\pi^\circ)$.*

Proof of Lemma 5. Fix any equilibrium of the cheap-talk game. From Lemma 4, the payoff to any expert must at least be $\overline{u}(\pi^\circ)$. We now argue that it cannot be higher. If (σ_1^*, σ_2^*) are the experiments chosen at the first stage, then in the ensuing sub-game, an expert's payoff is:

$$\begin{aligned}\sum_{(m_1,m_2)} \mathbb{P}(m_1,m_2)u(\alpha(m_1,m_2),\pi(m_1,m_2)) &= \sum_{(m_1,m_2)} \mathbb{P}(m_1,m_2) \min_{a\in A} u(a,\pi(m_1,m_2)) \\ &\leq \min_{a\in A} u\left(a, \sum_{(m_1,m_2)} \mathbb{P}(m_1,m_2)\pi(m_1,m_2)\right) \\ &= \min_{a\in A} u\left(a,\pi^\circ\right) = \bar{u}(\pi^\circ).\end{aligned}$$

(Recall that $\mathbb{P}$, α and π depend on (σ_1^*, σ_2^*), but to ease notation, we do not explicitly write the dependence.)

Finally, we argue that the decision-maker cannot get a payoff higher than $\overline{v}(\pi^\circ)$ either. Indeed, for the decision-maker to obtain a higher payoff, there must exist an action $a \in br(\pi^\circ)$ and a message profile (m_1, m_2) such that $\mathbb{P}(m_1, m_2) > 0$ and $a \notin br(\pi(m_1, m_2))$. This, however, would imply that an expert's equilibrium payoff is strictly less than $u(a, \pi^\circ)$, a contradiction with an expert's equilibrium payoff being equal to $\bar{u}(\pi^\circ) = \min_{a'\in\hat{A}} u\left(a', \pi^\circ\right)$.

The latter assertion follows from Lemma 3, which states that $u(a, \pi(m_1, m_2)) > u(\alpha(m_1, m_2), \pi(m_1, m_2))$ and $u(a, \pi(m_1', m_2')) \geq u(\alpha(m_1', m_2'), \pi(m_1', m_2'))$ for all pairs of messages (m_1', m_2') with $\mathbb{P}(m_1', m_2') > 0$. □

5 Conclusion

In this paper, we studied the effects of cross-verification on the decision-maker's and experts' payoffs. Clearly, cross-verification is not the sole reason for soliciting advice from multiple experts. Consulting a diverse set of experts with different opinions, specializations, preferences can provide a decision-maker with insights about the merits of different aspects of an issue. In fact, a decision-maker may be able to perfectly learn a multidimensional state by consulting experts about different dimensions. However, consulting experts that have information about different dimensions of a decision reduces the scope for cross-verification since cross-verification is most effective when experts' information is highly correlated. Moreover, as we demonstrated in this paper, the experts have an incentive to facilitate cross-verification by acquiring correlated information. This points to an interesting tension that can inform future research on committee design.

References

Ambrus, A. and S. E. Lu (2014): "Almost Fully Revealing Cheap Talk with Imperfectly Informed Senders," *Games and Economic Behavior*, 88, 174–189.

Ambrus, A. and S. Takahashi (2008): "Multi-sender cheap talk with restricted state spaces," *Theoretical Economics*, 3, 1–27.

Au, P. H. and K. Kawai (2020): "Competitive Information Disclosure by Multiple Senders," *Games and Economic Behavior*, 119, 56–78.

Aumann, R. J. and M. B. Maschler (1995): *Repeated Games with Incomplete Information*, Cambridge, MA: MIT Press.

Battaglini, M. (2002): "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 70, 1379–1401.

Border, K. C. (1990): "Fixed Point Theorems with Applications to Economics and Game Theory," *Cambridge Books*.

Crawford, V. P. and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 1431–1451.

Gentzkow, M. and E. Kamenica (2016): "Competition in Persuasion," *The Review of Economic Studies*, 84, 300–322.

——— (2017): "Bayesian Persuasion with Multiple Senders and Rich Signal Spaces," *Games and Economic Behavior*, 104, 411–429.

Gilligan, T. W. and K. Krehbiel (1989): "Asymmetric Information and Legislative Rules with a Heterogeneous Committee," *American Journal of Political Science*, 33, 459–490.

Green, J. and N. Stokey (1978): *Two Representations of Information Structures and their Comparisons*, Institute for Mathematical Studies in the Social Sciences.

Kamenica, E. (2019): "Bayesian Persuasion and Information Design," *Annual Review of Economics*, 11, 249–272.

Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Koessler, F., M. Laclau, and T. Tomala (2018): "Interactive Information Design," *HEC Paris Research Paper No. ECO/SCD-2018-1260.*

Krishna, V. and J. Morgan (2001a): "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review*, 435–452.

——— (2001b): "A Model of Expertise," *The Quarterly Journal of Economics*, 116, 747–775.

Li, F. and P. Norman (2018): "On Bayesian Persuasion with Multiple Senders," *Economics Letters*, 170, 66–70.

——— (2020): "Sequential Persuasion," *Theoretical Economics*, forthcoming.

Lipnowski, E. (2020): "Equivalence of Cheap Talk and Bayesian Persuasion in a Finite Continuous Model," Working paper.

Luenberger, D. G. and Y. Ye (2008): *Linear and Nonlinear Programming*, Springer.

Lyu, Q. (2020): "Information Design in Cheap Talk," *Available at SSRN 3579297.*

Mylovanov, T. and A. Zapechelnyuk (2013): "Decision Rules Revealing Commonly Known Events," *Economics Letters*, 119, 8–10.

Sobel, J. (2013): "Giving and Receiving Advice," *Advances in Economics and Econometrics*, 1, 305–341.

Wolinsky, A. (2002): "Eliciting Information from Multiple Experts," *Games and Economic Behavior*, 41, 141–160.