

Sequential persuasion

FEI LI

Department of Economics, University of North Carolina Chapel Hill

PETER NORMAN

Department of Economics, University of North Carolina Chapel Hill

This paper studies sequential Bayesian persuasion games with multiple senders. We provide a tractable characterization of equilibrium outcomes. We apply the model to study how the structure of consultations affects information revelation. Adding a sender who moves first cannot reduce informativeness in equilibrium and results in a more informative equilibrium in the case of two states. Moreover, with the exception of the first sender, it is without loss of generality to let each sender move only once. Sequential persuasion cannot generate a more informative equilibrium than simultaneous persuasion and is always less informative when there are only two states.

KEYWORDS. Bayesian persuasion, communication, competition in persuasion, multiple senders, sequential persuasion.

JEL CLASSIFICATION. D82, D83.

1. INTRODUCTION

This paper studies a canonical model of Bayesian persuasion with multiple senders in which senders disclose information sequentially. An uninformed decision maker seeks to maximize her state-dependent payoff. Also, many senders move in sequence, each constructing an experiment with a precision ranging from no information to full revelation of the state. Each sender observes the experiments designed by previous players when moving.

Decision makers often must rely on outside experts to take informed actions. Sometimes multiple experts are consulted and then consultations are often sequential. For

Fei Li: lifei@email.unc.edu

Peter Norman: normanp@email.unc.edu

We thank three anonymous referees for detailed comments. We also thank Yu Awaya, Gary Biglaiser, Yeon-Koo Che, Navin Kartik, Keiichi Kawai, Kyungmin (Teddy) Kim, Alexey Kushnir, Elliot Lipnowski, Shuo Liu, Giuseppe Lopomo, Laurent Mathevet, Stephen Morris, Vendela Norman, Luca Rigotti, Joel Sobel, and seminar participants at University of Hawaii at Manoa, the Einaudi Institute for Economics and Finance, Rome, and Collegio Carlo Alberto, Turin, as well as participants at the 2017 CETC Conference in Vancouver, the 2017 International Conference on Game Theory at Stony Brook, the 2017 Midwest Economic Theory Conference in Dallas, the 2018 Southern Economic Association Annual Meeting, the 2018 NSF/CEME Decentralization Conference at Duke, and the 2018 NSF/NBER/CEME Mathematical Economics Conference in Chicago for helpful comments. The usual disclaimer applies.

example, in a recent lawsuit, Students for Fair Admissions claims that Harvard intentionally discriminates against Asian–American applicants.¹ Each party used an economist expert witness to analyze Harvard’s admissions data and testify in court. Despite using the same data, the conclusions reached by the expert witnesses on each side were vastly different due to different statistical models. This example fits the Bayesian persuasion model well because experts were symmetrically informed and designed their own experiments. Furthermore, the consultations were truly sequential. Throughout the process, the expert on each side sequentially released rebuttals to reports made by the other side. Our model aims to understand how strategic considerations among experts shape information revelation in such settings.

Instead of relying on the concavification approach popularized by [Aumann and Maschler \(1995\)](#) and [Kamenica and Gentzkow \(2011\)](#), we characterize equilibrium outcomes using linear algebra techniques. Equilibrium conditions are expressed as incentive compatibility constraints and share a similar flavor as in [Bergemann and Morris \(2016\)](#).

The first step in the equilibrium construction is to show that every subgame perfect equilibrium outcome can be supported using one-step equilibrium strategies. In a one-step equilibrium, the only player who provides information is the first sender to move. The preferences of the other senders matter, but instead of actually refining the information on the path, their preferences restrict what the first sender does through incentive compatibility constraints. This works also off the equilibrium path, so any equilibrium can be replicated by strategies that are one-step on and off the equilibrium path.

Our second simplifying step is to show that only a finite set of vertex beliefs matter for the analysis. We assume a finite set of states and actions, so, in belief space, the optimal choice rule of the decision maker can be characterized as intersections of upper half spaces, or convex polytopes. Each polytope defines a set of beliefs for which an action is optimal and is spanned by a finite set of vertices. We demonstrate that it is without loss of generality for every sender to provide only information that generate beliefs on these vertices.

Focusing on one-step strategies with support on a finite set of vertices, we use backward induction to construct equilibria, which are Markov. We also use the fact that one-step equilibria on a finite set of vertices fully characterize the set of equilibrium outcomes so as to demonstrate that for a set of preferences of full measure, there is a unique equilibrium distribution over states and outcomes.

Equilibrium distributions are recursively defined as stable vertex beliefs. In concrete terms, for the truncated game starting with the last sender, a stable belief is a probability distribution over the state space that the last sender has no incentive to further refine. Moreover, it is without loss of generality to consider only the vertices of the polytopes that define optimal actions for the decision maker, which we denote by X . For a persuasion game with n senders, let $X_n \subseteq X$ be the stable vertex beliefs in the single-sender persuasion game with sender n only. The penultimate sender, $n - 1$, understands that

¹Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll. (Harvard Corp.), Civil Action 14-cv-14176-ADB, 2019 U.S. Dist. LEXIS 170309 (D. Mass. Sep. 30, 2019).

any belief not in X_n will be split onto X_n , so he may as well consider only beliefs in this set. However, for some beliefs in X_n , he may be better off by creating a mean-preserving spread over other beliefs in X_n , so the set of stable beliefs in the sequential-persuasion game starting with sender $n - 1$, X_{n-1} , is a subset of X_n . The set of stable beliefs for the full game is constructed recursively from this idea, and it shrinks for each step of the backward induction process.

By studying these stable beliefs, we find that adding a sender who moves first cannot reduce the informativeness. In contrast, strategic considerations may reduce information disclosure if a sender is added later in the game.

Next, we ask whether multiple counterarguments can make equilibria more informative in our model. The answer is mainly negative. We prove that the set of stable beliefs is unchanged if a sender is given an additional chance to provide information that precedes the last time that the sender moves. Hence, there is no loss of generality in considering an extensive form in which each sender moves only once when characterizing the set of stable beliefs. However, the first sender can choose the distribution over stable beliefs, and different senders may prefer different distributions. Hence, having all senders except possibly the first moving only once is without loss of generality. This may seem counterintuitive in the context of debates or legal proceedings, but our model lacks natural constraints such as limitations on the amount of information that can be transmitted using a single argument.

We also compare sequential and simultaneous persuasion. We find that sequential persuasion can never generate a more informative equilibrium than simultaneous persuasion. Finally, we provide a simple and easy to interpret sufficient condition for when full revelation is the unique equilibrium, which is invariant on the order of moves.

Literature. Our paper relates to a large body of work on information disclosure, but is most directly connected to the literature on Bayesian persuasion started by [Kamenica and Gentzkow \(2011\)](#) and [Rayo and Segal \(2010\)](#). This literature has recently been extended to incorporate multiple senders by [Gentzkow and Kamenica \(2017a, 2017b\)](#), [Boleslavsky and Cotton \(2015, 2018\)](#), [Au and Kawai \(2019, 2020\)](#), [Hwang et al. \(2019\)](#), and others. However, none of these papers deals with sequential moves by the senders. In a companion paper, [Li and Norman \(2018\)](#) provide some examples to show that adding new senders may reduce information revelation in multi-sender persuasion settings.

[Wu \(2018\)](#) considers a sequential Bayesian persuasion model similar to ours. He develops a recursive concavification approach based on [Harris \(1985\)](#) and [Kamenica and Gentzkow \(2011\)](#) to establish equilibrium existence, and he independently constructs a one-step equilibrium (referred to as a silent equilibrium). Our paper differs from [Wu \(2018\)](#) in the following aspects. First, our methodologies are different. Thanks to the assumption of finite action space, we can apply primitive tools such as backward induction, convex polytope analysis, and linear programming to transparently characterize the equilibrium. Second, our model clarifies how senders' experiments are combined. This enables us to transparently compare equilibria for different extensive forms.

A growing body of work embeds persuasion into dynamic models (see [Ely et al. 2015](#) and [Ely 2017](#)), but the paper closest in spirit to ours is [Board and Lu \(2018\)](#), which incorporates Bayesian persuasion into a search model. However, [Board and Lu \(2018\)](#) consider payoff functions that are more restrictive than ours, and the decision maker in their

paper faces an optimal stopping problem. In contrast, the decision maker has no influence on the precision of her information in our model. Our formal analysis has some similarities with that of [Lipnowski and Mathevet \(2017, 2018\)](#), which focus on single-sender persuasion games.

Multi-sender information provision has been studied in other frameworks. [Glazer and Rubinstein \(2001\)](#) study a finite horizon sequential-persuasion model with limitations on the amount of information that can be revealed in each stage. There are also papers in the cheap talk and disclosure literature that ask what the implications of multiple senders are. See [Ambrus and Takahashi \(2008\)](#), [Battaglini \(2002\)](#), [Kawai \(2015\)](#), [Krishna and Morgan \(2001\)](#), [Kartik et al. \(2019, 2017\)](#), [Bhattacharya and Mukherjee \(2013\)](#), and [Milgrom and Roberts \(1986\)](#). [Hu and Sobel \(2019\)](#) compare simultaneous and sequential information disclosure in a setting where senders decide which set of facts to disclose and where the focus is on equilibria surviving iterated elimination of weakly dominated strategies.

With different applications in mind, these papers introduce frictions on information transmission such as asymmetric information, limited information process ability, restricted forms of signals, etc. Instead, our framework eliminates all such frictions and focuses solely on the strategic interaction among senders. It thus serves as a natural benchmark for identifying sources of communication inefficiency.

Organization. The remainder of this paper is organized as follows. [Section 2](#) describes the model. [Section 3](#) characterizes the set of equilibria, shows that every equilibrium outcome is supported as a one-step equilibrium with finite support, that equilibria exist, and that the equilibrium outcome is generically unique. In [Section 4](#), we apply the equilibrium characterization to discuss effects of changes in the extensive form. [Appendix A](#) collects omitted proofs and some examples are collected in [Appendix B](#).

2. THE MODEL

Players. Consider an environment with senders $i = 1, \dots, n$ and a decision maker d . Player $i = 1, \dots, n, d$ has a utility function $u_i : A \times \Omega \rightarrow \mathbb{R}$, where A is a finite set of actions and Ω is a finite state space. Payoff functions are common knowledge and players evaluate lotteries using expected utilities. Players hold a common prior belief $\mu_0 \in \Delta(\Omega)$. Fixing a belief μ and an action a , we define player i 's expected payoff as

$$v_i(a, \mu) \equiv \sum_{\omega \in \Omega} u_i(a, \omega) \mu(\omega) \quad \text{for } i = 1, \dots, n, d.$$

Experiments. Players are uninformed about the state of the world, but a sender may provide information to the decision maker by creating an *experiment*. We use the *partition representation* of experiments from [Green and Stokey \(1978\)](#) because combining multiple experiments becomes very intuitive under this representation.²

Under the partition representation, an experiment is given by a partition of $[0, 1] \times \Omega$, where, for each state ω , $\{\pi(s|\omega)\}_{s \in S}$ are disjoint sets such that $\bigcup_{s \in S} \pi(s|\omega) = [0, 1]$ and S

²This also allows us to easily compare our sequential framework with the simultaneous move model in [Gentzkow and Kamenica \(2017b\)](#).




FIGURE 1. There are two states ω_0 and ω_1 , and two senders $i = 1, 2$. Sender 1's signal space contains two signals: s_1 and s'_1 . Sender 2's experiment has two possible signals $\{s_2, s'_2\}$. The combination of two experiments $\hat{\pi}_2 = \pi_1 \vee \pi_2$ has three possible signals $\{\hat{s}_2, \hat{s}'_2, \hat{s}''_2\}$, and it is finer than π_1 and π_2 .

indexes the sets in partitions. Given experiment π , one can interpret each s as a *signal* by assigning state-contingent probabilities to each s according to the Lebesgue measure of each $\pi(s|\omega)$. In doing so, experiment π induces a state-contingent distribution over signals $p_\pi : \Omega \rightarrow \Delta(S)$. Letting $\lambda(\cdot)$ denote the Lebesgue measure, the probability of signal $s \in S$ being realized conditional on state ω is

$$p_\pi(s|\omega) = \lambda(\pi(s|\omega)), \tag{1}$$

where $\sum_s p_\pi(s|\omega) = 1$ for each $\omega \in \Omega$ because $\{\pi(s|\omega)\}_{s \in S}$ is a partition of the unit interval. With a slight abuse of notation, we use s both as a generic indexing set and the corresponding subset of $[0, 1] \times \Omega$ in the discussion below and in Figure 1.

Given two experiments π, π' , players combine the information into a joint experiment that we denote by $\pi \vee \pi'$, which consists of the set of all intersections of the sets in π and π' . Since each set in the joint experiment is an intersection of a set in the partition π with a set in the partition π' , it is immediate that $\pi \vee \pi'$ is finer than both π and π' . This, in turn, implies that the combined experiment $\pi \vee \pi'$ is more informative in Blackwell's sense than either of the two underlying experiments.³

Extensive form. Let Π denote the set of all experiments. Senders $1, \dots, n$ move sequentially to post experiments π_1, \dots, π_n in order of their index, where $\pi_i \in \Pi$ for every i and where each sender observes all previous senders' experiments. Then nature draws ω . Finally, the decision maker observes (π_1, \dots, π_n) and a joint realization $s = (s_1, \dots, s_n)$ according to the corresponding state-contingent probability $p_{\bigvee_i \pi_i}(s|\omega) = \lambda(\bigvee_i \pi_i(s|\omega))$ for $i = 1, \dots, n$, and takes an action $a \in A$.

As illustrated in Figure 1, combining sender 2's experiment with the experiment of sender 1 generates a finer joint experiment than either underlying experiment. Each

³Assume that π is finer than π' , and let p_π and $p_{\pi'}$ denote the corresponding state-contingent distributions over signals they generate. Then p_π is more informative in the sense of Blackwell (1953) than $p_{\pi'}$. See Green and Stokey (1978) for a proof.

signal in π_1 may be further partitioned and provides an example in which s_1 , but not s'_1 , is refined when combined with the experiment played by player 2. A sender, therefore, acts as if he observes and responds to the signal realizations of previous senders' experiments, despite the fact that the formal model assumes that the joint signal realization is drawn at the end. Generating joint experiments by taking intersections is, therefore, without loss of generality in our model because senders move sequentially.

Strategies and equilibrium. A pure strategy for sender i is a map $\sigma_i : \Pi^{i-1} \rightarrow \Pi$, where Π^0 is the trivial null history. That is, given a history $\{\pi_1, \dots, \pi_{i-1}\}$, sender i chooses π_i that results in a finer experiment $\bigvee_{k=1}^i \pi_k$. A history for the decision maker is a vector $(\pi_1, \dots, \pi_n, s_1, \dots, s_n)$. Let \mathcal{H}_d be the set of all histories for the decision maker and let $\sigma_d : \mathcal{H}_d \rightarrow A$ denote her strategy. There is uncertainty about the state, but information is symmetric and there is, therefore, never any point in the game in which any player needs to update the beliefs about the type of other players. Hence, subgame perfection is applicable.

3. EQUILIBRIUM CHARACTERIZATION

In this section, we first prove a result similar to the revelation principle. Without loss of generality, we may focus on *one-step* equilibria, which are equilibria where only the first sender discloses nontrivial information on the equilibrium path. Preferences of other senders enter much like incentive compatibility constraints in such equilibria. Then we construct an equilibrium, and show that the game has a unique equilibrium distribution over states and actions for a set of payoff functions with full Lebesgue measure.

3.1 Simplifying the problem

Players ultimately care only about the distribution over actions and states, which motivates the following definition.

DEFINITION 1. Two strategy profiles are *outcome equivalent* if they generate identical joint distributions over $\Omega \times A$.

There are often multiple outcome-equivalent equilibrium information structures, but all players are indifferent across all such equilibria. We, therefore, consider them equivalent even if they are Blackwell comparable, because ultimately players care only about probability distributions over $\Omega \times A$.

Next, we define strategy profiles in which only the first sender provides any information.

DEFINITION 2. Consider a strategy profile σ' and let h'_i denote the implied outcome path before the move by sender i . We say that σ' is *one-step* if $\bigvee_{i=1}^n \sigma'_i(h'_i) = \sigma'_1$.

We are now ready to present the first result.

PROPOSITION 1. *For any subgame perfect equilibrium, there exists an outcome-equivalent subgame perfect equilibrium in which senders play a one-step continuation strategy profile after any history of play.*

The idea behind **Proposition 1** is similar to the revelation principle. Consider an arbitrary subgame perfect equilibrium σ^* and let $\{\pi_1^*, \dots, \pi_n^*\}$ be the individual experiments on the equilibrium path that generate a joint experiment $\pi^* = \bigvee_{i=1}^n \pi_i^*$. To construct a one-step equilibrium, let sender 1 play π^* and assume that on the equilibrium path players $i = 2, \dots, n$ provide only redundant information. It then follows that the decision maker may as well generate the same distribution over $A \times \Omega$ as in the initial equilibrium after observing the one-step path history. Moreover, because π^* is finer than π_i^* for each $i < n$, any deviation that is feasible from the one-step outcome path is feasible also in the original equilibrium, so it is possible to replicate continuation play following deviations from the one-step equilibrium from the original equilibrium just like in the proof of the revelation principle. Off the equilibrium path, we can follow the original equilibrium strategies.⁴

For the one-step equilibrium characterization to be a significant simplification, it is important that it applies not only on the equilibrium path, but also off the path. The same logic as on the equilibrium path generalizes to any continuation equilibrium following an arbitrary history of play.

Proposition 1 implies that solving for an equilibrium of a sequential-persuasion game is equivalent to solving a static single-sender persuasion game disciplined by additional recursively defined incentive compatible constraints. After stage 1, no sender has an incentive to provide further information given the threat of subsequent senders' best responses.

3.2 Equilibrium construction

Now we explicitly construct a one-step equilibrium. The construction is essential for the rest of our analysis because several concepts critical to understanding the equilibrium structure and the effect of competition in persuasion are introduced through the process.

An equilibrium is constructed by backward induction. We begin with the decision maker's problem. As in standard persuasion models, what matters for the decision maker is her posterior belief about the state. Moreover, a key simplification is that we without loss may restrict attention to a finite set of *vertex beliefs*. Combined with the one-step equilibrium characterization, this allows us to construct equilibria recursively by checking which *stable vertex beliefs* in the continuation game are weakly better for the current sender than every mean-preserving spread over the stable vertex beliefs in the continuation game.

⁴The proof of **Proposition 1** is drastically simplified by two slightly unconventional modelling decisions. First, the partition representation makes it very easy to describe how individual experiments combine into a joint experiment. Second, having the uncertainty being resolved after all senders have moved implies that a history is a sequence of successively finer partitions. Hence, we can avoid senders having to condition on realized signals, which is a big simplification of the proof.

Decision maker's problem. Suppose that the decision maker observes a history of experiments $\{\pi_j\}_{j=1}^n$, which induces a joint experiment $\bigvee_{j=1}^n \pi_j$, as well as a signal realization s . Using $\bigvee_{j=1}^n \pi_j$ and s , the decision maker updates her belief about the state that summarizes all payoff-relevant aspects of the history. The posterior probability of state $\omega \in \Omega$ is

$$\mu(\omega|s) = \frac{p(s|\omega)\mu_0(\omega)}{\sum_{\omega' \in \Omega} p(s|\omega')\mu_0(\omega')}, \tag{2}$$

where we have dropped the subscript of $p(s|\omega)$ defined in (1). Denoting the unconditional probability of s by $p(s) = \sum_{\omega' \in \Omega} p(s|\omega')\mu_0(\omega')$, we note that an experiment π induces a distribution of posterior beliefs that satisfies the Bayes plausibility constraint

$$\sum_{s \in \pi} \mu(\omega|s)p(s) = \mu_0(\omega).$$

To characterize the optimal actions for the decision maker, we note that for any distinct pair $a, a' \in A$, the set

$$H(a \succeq a') \equiv \left\{ \mu \in \Delta(\Omega) \mid \sum_{\omega \in \Omega} \mu(\omega)[u_d(a, \omega) - u_d(a', \omega)] \geq 0 \right\}$$

defines the posterior beliefs such that the decision maker weakly prefers a to a' . It follows that the set of beliefs such that $a \in A$ is optimal is given by

$$M(a) = \bigcap_{a' \in A} H(a \succeq a'),$$

which is a finite convex polytope. See Figure 2 for a simple illustration.




FIGURE 2. $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and $A = \{a_1, a_2, a_3, a_4\}$.

Interim beliefs. A history $h_i = \{\pi_j\}_{j=1}^{i-1}$ induces a joint experiment $\pi^{i-1} = \bigvee_{j=1}^{i-1} \pi_j$. For each signal s of π^{i-1} , the corresponding belief $\mu(\omega|s)$ is given by (2). This is the decision maker's posterior belief if senders $i + 1, \dots, n$ do not add any information in the continuation game and s is realized. We call such a belief an interim belief. Each joint experiment $\pi^{i-1} \in \Pi$ generates a distribution of interim posterior beliefs τ^{i-1} , and we let $\Delta(\Delta(\Omega))$ denote the set of distributions of (interim or posterior) beliefs.

Given a joint experiment π^{i-1} that induces a joint belief distribution τ^{i-1} , sender i can refine the information into any partition that is finer than π^{i-1} . Using Theorem 1 in Green and Stokey (1978) together with the characterization in Gentzkow and Kamenica (2017a), we know that any mean-preserving spread of τ^{i-1} can be induced by some refined partitioning of π^{i-1} . Every feasible experiment for sender i therefore corresponds to a mean-preserving spread of each interim belief in the support of τ^{i-1} . Hence, sender i 's problem separates into finding an optimal mean-preserving spread belief by belief from the distribution induced by previous senders.

Sender n 's problem. Next, we consider the last sender's problem. The construction of $\{M(a)\}$ implies that we may consider optimal strategies for the decision maker that map posterior beliefs to actions. We abuse notation and denote such a map by $\sigma_d(\mu) \in \{a : \mu \in M(a)\}$. To guarantee that sender n 's problem is well defined, we assume that the decision maker always breaks ties in favor of sender n . If there are multiple such rules, we arbitrarily pick one of them. Given an interim belief μ and decision rule σ_d , sender n 's program can be written as

$$\begin{aligned}
 V_n(\mu) &= \max_{\tau \in \Delta(\Delta(\Omega))} \sum_{\mu'} v_n(\sigma_d(\mu'), \mu') \tau(\mu') \\
 \text{s.t. } &\sum_{\mu'} \mu' \tau(\mu') = \mu,
 \end{aligned}
 \tag{3}$$

and a solution is a mean-preserving spread of μ , denoted by $\tau_n(\cdot|\mu)$.

By construction, the beliefs for which a is optimal for the decision maker, $M(a)$, is a finite convex polytope for each $a \in \mathcal{A}$. Such a convex polytope has a finite set of $J(a)$ vertices $\{\mu_j^a\}_{j=1}^{J(a)}$ and these vertices span $M(a)$ so that every $\mu \in M(a)$ can be represented as a convex combination of the vectors $\{\mu_j^a\}_{j=1}^{J(a)}$.⁵ Denote

$$X = \bigcup_{a \in \mathcal{A}} \{\mu_j^a\}_{j=1}^{J(a)}
 \tag{4}$$

as the set of all vertices that defines the optimal actions for the decision maker, which is finite because both Ω and \mathcal{A} are finite.

LEMMA 1. *Program (3) has a solution $\tau \in \Delta(X)$.*

Hence, while there may be optimal solutions to (3) with support on a larger (even infinite) set, we can always find a solution in $\Delta(X)$. The idea is that each $M(a)$ is spanned

⁵See Grünbaum et al. (1967).

by its vertices. Hence the sender can replace any belief μ that is not one of the vertices with a convex combination over the vertices. There are then two possibilities. The first is that the action $\sigma_d(\mu)$ is taken on all the vertices in the convex combination. In this case, the sender is indifferent between μ and the convex combination over the vertices of $M(a)$. The second possibility is that a different action is taken on one or more of the vertices. Because the tie-breaking favors the sender, he is either indifferent or strictly better off by using the convex combination. Hence, restricting to $\Delta(X)$ generates a utility at least as great as (3). But $\Delta(X)$ is a subset of the feasible set in (3), so the two problems must have the same value.

Figure 2 provides an illustration. It depicts a feasible solution in which belief μ in the interior of $M(A_2)$ is played with positive probability. Replacing μ with the mean-preserving spread onto $\{\mu_j^{a_2}\}_{j=1,2,3}$ can be no worse for n because the decision maker breaks ties in favor of n at $\{\mu_1^{a_2}\}$ and $\{\mu_2^{a_2}\}$.

Lemma 1 suggests that we may characterize the optimal mean-preserving spread of every sender in terms of a finite optimization problem. The general idea is that if the last sender always uses a best response with support on the vertex beliefs X , then previous senders may as well use strategies limited to the same set of vertices, since the final sender will undo any attempt to generate any other beliefs by splitting them onto X .

Stable beliefs. To proceed further, we recursively define a set of *stable (vertex) beliefs*. Let X_n denote the set of vertex beliefs where sender n has no incentive to provide further information, i.e.,

$$X_n \equiv \{\mu \in X : v_n(\sigma_d(\mu), \mu) = V_n(\mu)\}. \tag{5}$$

Then we recursively define $\{X_i\}_{i=1}^n$ such that

$$X_i \equiv \{\mu \in X_{i+1} : v_i(\sigma_d(\mu), \mu) = \tilde{V}_i(\mu)\}, \tag{6}$$

where

$$\begin{aligned} \tilde{V}_i(\mu) &= \max_{\tau \in \Delta(X_{i+1})} \sum v_i(\sigma_d(\mu'), \mu') \tau(\mu' | \mu) \\ \text{s.t.} \quad &\sum_{\mu' \in X_{i+1}} \mu' \tau(\mu' | \mu) = \mu. \end{aligned} \tag{7}$$

Notice that (i) a solution to the auxiliary program (7) exists, (ii) $X_i \subseteq X_{i+1}$, and (iii) $X_1 \neq \emptyset$. In the auxiliary problem (7), sender i is restricted to use experiments that induce vertex beliefs only in X_{i+1} , and he believes that senders $i + 1, \dots, n$ will not add any information. Because $X_i \subseteq X_j \forall j > i$, sender i 's belief is indeed justified.⁶

DEFINITION 3. A belief is *stable* if $\mu \in X_i$ which is recursively defined by (5) and (6) for $i = 1, \dots, n$.

⁶It would be natural to define stable beliefs not just on the vertices. However, it is without loss of generality to consider equilibria with support on the vertices, and we avoid tedious repetitions of “stable vertex beliefs” by having the definition apply to vertices only.

By construction, no sender has an incentive to refine a stable belief. Therefore, one can recursively construct a one-step equilibrium where the resulting posterior belief is distributed only on the set of stable beliefs. On the path of play, if $\mu_0 \in X_1$, no sender sends a nontrivial signal; if $\mu_0 \notin X_1$, only sender 1 posts an informative experiment and the other senders provide no information. Off the equilibrium path, if one of sender i 's interim beliefs is $\mu_{i-1} \notin X_i$, he posts an experiment that “splits” the beliefs only in X_i and the subsequent senders do not add further information.

A key step in the construction is to make sure that best responses on the vertices exist for each sender. This is done by using strategies that split any nonvertex belief onto vertices and, which is crucial, never refine a stable vertex belief.⁷ Together, these two restrictions on continuation play imply that each player effectively has a finite choice set. This does not rely on making value functions continuous (or upper semicontinuous) in beliefs. For further details the reader may consult [Appendix A](#).

PROPOSITION 2. *There exists a one-step equilibrium.*

Notice that the equilibrium is *Markov* in the following sense. The decision maker's equilibrium strategy σ_d depends on the history only through the posterior belief, and for each $i = 1, 2, \dots, n$, for every experiment profile π_1, \dots, π_{i-1} and possible signal profile (s_1, \dots, s_{i-1}) that induce the same interim belief, the mean-preserving spread τ_i induced by sender i 's equilibrium strategy is identical.

3.3 Outcome uniqueness

Our third result regards the uniqueness of the equilibrium outcome, formally stated as follows.

PROPOSITION 3. *All subgame perfect equilibria are outcome equivalent for a set of payoff function profiles with full Lebesgue measure.*

Proposition 3 says that for generic preferences, there is an essentially unique equilibrium. Together with the fact that we can always construct a Markov equilibrium, this implies that restricting attention to Markov strategies is almost always without loss of generality. The case that may create multiple equilibrium outcomes is if one sender is indifferent between some vertex $\mu \in X_i$ and some mean-preserving spread over X_i while some other players are not indifferent. However, such preferences are knife edge and have probability 0.

The proof is relegated to [Appendix A](#). For intuition, first notice that the one-step equilibrium we construct in [Section 3.2](#) induces vertex beliefs only. A key intermediate result, [Lemma 2](#) below, establishes that this is without loss of generality

⁷Players may be indifferent between refining and not refining a stable vertex belief, and using a best response in which a stable belief is refined could make the best response problem of a previous mover ill-defined

LEMMA 2. *For every subgame perfect equilibrium, there exists an outcome-equivalent subgame perfect equilibrium in which senders play one-step strategies with implied beliefs with support on X after every history of play.*

The basic idea is much like Proposition 1, but the proof has to deal with on and off equilibrium path histories, and is, therefore, notationally more cumbersome. Lemma 2 is crucial because not only can we restrict attention to an equilibrium experiment that is restricted to vertex beliefs, but, additionally, it is without loss of generality to check one-step deviations to vertices. Therefore, if two continuation equilibria that are not outcome equivalent exist, some sender must be indifferent between some $\mu \in X$ and a mean-preserving spread with support on X .

There are two cases in which a sender is indifferent to splitting a belief to X . The first case is when a mean-preserving spread always induces the same action as the original belief. Such indeterminacy is irrelevant, as the distribution over $\mathcal{A} \times \Omega$ is unchanged. Any failure of essential uniqueness therefore corresponds to indifferences over mean-preserving spreads that induce distinct actions. However, this requires nongeneric preferences. Since X is a finite set, there exists a finite number of affinely independent sets of belief vectors and indifference between any two such sets can hold for a measure zero set of preferences. There is a finite set of pairs to consider, and it follows that essential uniqueness can fail for at most a measure zero set of preferences.

4. APPLICATIONS

This section discusses some applications. The aim is to shed light on some issues relevant for the design of a communication protocol. Specifically, to maximize the amount of information disclosure, the decision maker can structure the communication by selecting experts, organizing the order of consultations, deciding what information to share with experts, etc. As a first step, we examine some key aspects that affect the incentives for information revelation, including the number of senders, the order of the senders' moves, and the information shared among senders. Thanks to the stable belief characterization of equilibrium outcomes, this becomes relatively straightforward, as we can focus on how changes in the extensive form affect the set of stable beliefs.

Our goal is to derive some principles guiding the design of how to structure consultations. We focus on results that hold for arbitrary preferences. The justification for this is that results that do not depend on specific assumptions about preferences are more robust and may also be of value for real-world applications when preferences are not observable.

4.1 Information criteria

We begin with defining the criteria to evaluate information revelation. A unique equilibrium outcome makes comparisons more straightforward and transparent. Unfortunately, when senders move simultaneously, the only possibility to have such uniqueness

is when full revelation is the unique equilibrium. In general, one must use setwise comparisons. In contrast, the sequential model has a unique outcome for generic preferences. In the rest of the paper, we focus on the generic case with an essentially unique equilibrium distribution over states and outcomes in the sequential model.

It is easy to construct examples with multiple equilibrium belief systems that can be ranked according to the Blackwell order, but where the differences in informativeness are irrelevant because all equilibria induce the same joint distribution over $A \times \Omega$. We, therefore, treat π and π' as equivalent in terms of the information content provided that they are outcome equivalent.

DEFINITION 4 (Essential Blackwell order). For any given decision correspondence, we say that π is *essentially less informative* than π' if there exists an experiment that is outcome equivalent to π' and more informative than any experiment that is outcome equivalent to π in the Blackwell order.

First, note that this is a well defined partial order. If π'' is outcome equivalent to π' and more informative than any experiment that is outcome equivalent to π , there exists no experiment outcome equivalent to π that is strictly more informative than π'' , so antisymmetry holds. Transitivity is equally obvious.

Next, note that it is sufficient to compare experiments with support on vertex beliefs, as there exists an outcome-equivalent mean-preserving spread onto the vertices for any experiment involving at least one signal that is not on a vertex. Hence, consider an experiment in which belief μ in Figure 2 has a positive probability. When comparing the informativeness of this experiment to another experiment, we first replace μ with the mean-preserving spread onto $\{\mu_1^{a_2}, \mu_2^{a_2}, \mu_3^{a_2}\}$. In this example, the relevant mean-preserving spread is unique, which is not always true. However, by Proposition 3, for generic preferences, there is a unique mean-preserving spread on the vertices of $M(a)$ for every $\mu \in M(a)$ and then the order compares the finest experiment outcome equivalent with π to the finest experiment outcome equivalent with π' .

Finally, note that outcome equivalence can only be defined given the decision maker's preference. Hence, our essential Blackwell order is not purely based on informativeness, which is different from individual sufficiency in Bergemann and Morris (2016) and other conventional information criteria applying to all preferences. The advantage of our order is to allow us to focus on the comparison of outcome-relevant information of two information structures.

4.2 Adding senders in sequential persuasion

In this section, we examine the effect of adding senders in a sequential move Bayesian persuasion game and derive some general results. Intuition suggests that the added competition from an increase in the number of experts should increase the amount of information revealed in the market. This view may even be seen as an intellectual foundation for freedom of speech, a free press, the English common law system, and many other institutions. While the literature provides a somewhat mixed support for this view,

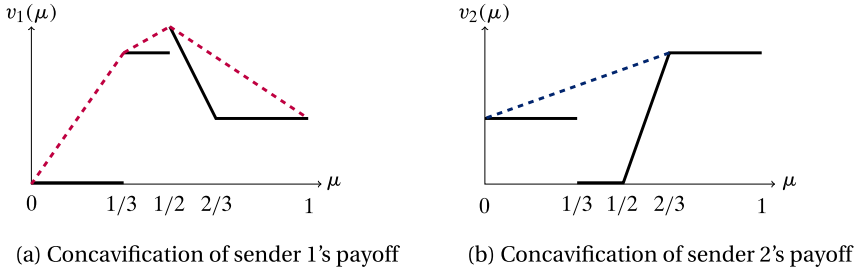


FIGURE 3. Continuation payoffs and the order of moves. The solid lines are the senders' payoffs as a function of decision maker beliefs, while the dashed line represents the concavified payoff when 1 (2) is the only sender. In (a), sender 1 splits $\mu < 1/3$ onto $\{0, 1/3\}$, $\mu \in [1/3, 1/2]$ onto $\{1/3, 1/2\}$, and $\mu > 1/2$ onto $\{1/2, 1\}$. In (b), sender 2 splits $\mu < 2/3$ onto $\{0, 2/3\}$ and $\mu > 2/3$ onto $\{2/3, 1\}$.

Gentzkow and Kamenica (2017a, 2017b) provide sufficient conditions under which additional senders do not reduce the amount of information revealed in simultaneous move games. Sequential moves further weakens the argument for additional experts generating more information, because the order of moves matters.

Consider an example with two states and two senders. Figure 3 depicts the preferences over the beliefs of the decision maker for sender 1 and 2 (in their single-sender persuasion games), respectively.⁸ Notice in particular that in a single-sender persuasion problem with $\mu > 2/3$, beliefs are split onto $\{1/2, 1\}$ when the experiment is constructed by sender 1.

When the two senders move in sequence, full revelation is the unique equilibrium if sender 1 is the last mover. In contrast, full revelation is *not* an equilibrium when sender 2 is the last mover, and for priors exceeding $2/3$, the equilibrium is less informative than the experiment constructed when sender 1 is the single sender. The difference between the two cases is that sender 1 is unable to commit to not splitting $\mu = 2/3$. Anticipating this, sender 2 provides full information. When the order is reversed, the commitment issue is gone.

To understand this, suppose that sender 1 is the last mover. Note that the tie is broken in favor of the action corresponding to $[1/3, 1/2]$ at $\mu = 1/3$ and the action corresponding to $[1/2, 2/3]$ at $\mu = 1/2$. Any belief in $(0, 1)$ is thus split by sender 1 in such a way that sender 2 gets the lowest possible payoff except when μ is 0 or 1. The unique best response for sender 2 is, therefore, to fully reveal the state.

In contrast, if sender 2 is the last mover, any μ in $[0, 2/3]$ is split onto $\{0, 2/3\}$. It follows that if the prior exceeds $2/3$, the (finest) best response for sender 1 is to split the beliefs onto $\{2/3, 1\}$, which results in no further refinement by sender 2. Hence, the order of moves matters for the equilibrium outcome. Moreover, for a prior larger than $2/3$, the equilibrium is less informative than the single-sender equilibrium with sender 1, which is to split the prior onto $\{1/2, 1\}$.

⁸We can generate such preferences if the decision maker has four actions available.

In the example above, the equilibrium is more informative when the new sender is added as a first mover. This is not quite general due to the incompleteness of the Blackwell ordering, but we can establish an analogue of the result for simultaneous move games.

PROPOSITION 4. *For generic preferences, if a sender is added who moves before all other senders, there is no equilibrium with $n + 1$ senders that is essentially less informative than the equilibrium in the original game.*

PROOF. Let X_1^n be the set of stable beliefs in the game with n senders and let X_1^{n+1} be the set of stable beliefs in the game with $n + 1$ senders. Because sender $n + 1$ is added to move before senders $1, \dots, n$ and the set of stable beliefs is defined backwardly, we have that

$$X_1^{n+1} \subseteq X_1^n. \quad (8)$$

Fix the prior belief μ_0 , let $X_1^n(\mu_0)$ be the support of the equilibrium in the game with n senders and let $X_1^{n+1}(\mu_0)$ be the support of the equilibrium in the game with $n + 1$ senders. As discussed in Section 4.1, when introducing the essential Blackwell ordering, it is without loss to assume that these beliefs are vertices and stable, i.e., $X_1^j(\mu_0) \subseteq X_1^j$ for $j = n, n + 1$.

For contradiction, suppose that the game with $n + 1$ senders has an equilibrium that is essentially less informative than the equilibrium in the original game with n senders. Then there exists at least one belief $\mu' \in X_1^{n+1}(\mu_0)$ such that μ' is in the convex hull of $X_1^n(\mu_0)$, but $\mu' \notin X_1^n(\mu_0)$. Because preferences are generic, in the original n -sender game, some sender has a strict incentive to split μ' onto X_1^n . Hence, $\mu' \notin X_1^n$, which contradicts (8). \square

The proposition says that when a new sender is added to move before all previous senders, the equilibrium cannot sustain more uncertainty regardless of the preference profile of the senders. The idea is simple. If a belief is induced by an equilibrium, it must be stable. Recall that the set of stable beliefs is constructed backwardly. Adding a new sender who moves first can only reduce the set of stable beliefs. As a result, such a change cannot make the outcome essentially less informative unless there are multiple equilibrium outcomes, which is ruled out by the restriction to generic preferences. In contrast, the counterexample in Figure 3 with sender 2 added at the end is robust in the sense that the qualitative features of the example are robust to perturbations in sender and decision maker payoff functions.

In the special case where there are only two states, the incompleteness of the essential Blackwell order no longer matters and we obtain a stronger result.

PROPOSITION 5. *Suppose that $\Omega = \{\omega_0, \omega_1\}$. If a sender is added who moves before all other senders, every equilibrium with $n + 1$ senders is weakly essentially more informative in the Blackwell ordering.*

PROOF. Without loss of generality, consider a one-step equilibrium with support on X , which contains beliefs where the decision maker is indifferent between two actions together with 0 and 1. Let $X_1^j(\mu_0)$ be the support of the equilibrium in the game with j senders where $j = n, n + 1$. For contradiction, assume that there are $\{\mu_L, \mu_M, \mu_H\}$ such that $\mu_M \in X_1^{n+1}(\mu_0)$, $\{\mu_L, \mu_H\} \in X_1^n(\mu_0)$, and $\mu_M \in (\mu_L, \mu_H)$. Without loss we can assume that there are at least two distinct actions that are taken at beliefs $\{\mu_L, \mu_M, \mu_H\}$, as otherwise μ_M would not be on the set of vertices X . But for μ_M to be stable with $n + 1$ players, every sender $i \in \{1, \dots, n + 1\}$ must be weakly better off at μ_M than at the unique mean-preserving spread onto $\{\mu_L, \mu_H\}$. This implies that transferring probability from $\{\mu_L, \mu_H\}$ to μ_M is consistent with equilibrium in the model with n senders, contradicting uniqueness with n senders. \square

The difference between Propositions 4 and 5 can be illustrated in Figure 4. The left panel of Figure 4 visualizes a case with three states. The support of the finest equilibrium is $X_n^1(\mu_0) = \{\mu_1, \mu_2, \mu_3\}$ in the original n -sender game. When a new sender is added to speak before other senders, the support of the finest equilibrium becomes $X_{n+1}^1(\mu_0) = \{\mu_1, \mu_2, \mu_4\}$. Proposition 4 leaves the possibility that two equilibria are non-comparable in the sense of Blackwell. To the contrary, when there are only two states, the support of the finest equilibrium contains at most two stable beliefs for generic preferences. Proposition 4 implies that $\mu_L^{n+1} \leq \mu_L^n \leq \mu_H^n \leq \mu_H^{n+1}$, which is visualized on the right panel of Figure 4.

When senders are added at any place except as a first mover, there is nothing that can be said in general about how the informativeness is affected. We know from Li and Norman (2018) that adding a sender at the end may strictly reduce the information revealed, and the example in Figure 3 is another example of that. To see that the same possibility exists when senders are added in the middle, assume that there is a sender 3 who has a preference such that splitting any vertex beliefs makes him worse off. Adding this sender (or multiple versions of him) at the end of any game with one or two senders leaves the equilibrium unchanged. Hence, we can use any example in which adding a sender at the end reduces information relative the single-sender problem to create an example where adding a sender in the middle reduces the information relative to a persuasion game with two (or multiple) senders. To construct examples where adding a




FIGURE 4. The left panel represents an example with $|\Omega| = 3$, while the right panel represents an example with $|\Omega| = 2$.

sender in the middle adds information is even easier. For example, one can just add a sender who prefers full revelation at any position in a game that does not fully reveal the state before the addition of the sender.

4.3 Multiple moves by the same sender

Our second application considers the communication protocol for a given set of senders. Up to this point, we have allowed each player to move only once. This is without loss of generality for results having to do with the characterization, existence, and uniqueness of equilibria, because we can always add multiple players with identical preferences. However, we now ask whether it is useful for the decision maker to allow multiple counterarguments or whether a sender is better off by moving more than once.

This exercise is relevant because senders who speak at late stages can respond to early movers' arguments, i.e., disclosing information conditional on the signals sent by previous senders. Then it is natural to ask if there is any value in letting senders respond to counterarguments from other senders? If so, what is the source of the value?

Our model offers a frictionless benchmark to identify the conditions needed to rationalize multiple rounds of rebuttals and counterarguments. Preferences are common knowledge and a sender can provide as much information as he wants in a single round of disclosure. Hence, the only constraints on communication are strategic considerations. Our results imply that these strategic considerations are per se insufficient to justify multiple rounds of communication, except that moving twice may be useful for the first sender who moves.

Formally, we let $i \in \{1, \dots, n\}$ denote the set of senders and we let the stage when senders move be denoted by $t = 1, \dots, T$ with $n \leq T$.

PROPOSITION 6. *Consider any sequential-persuasion game with n senders and finite horizon $n \leq T$. Then the set of stable beliefs is the same as in the sequential game with n senders and n periods in which, for each sender i , every move except the last one is eliminated.*

Proposition 6 says that for any sequential-persuasion game where senders move multiple times, to pin down its stable beliefs, it is sufficient to examine a reduced form game where each sender moves only once. For example, consider a game with three senders $i = 1, 2, 3$ and five stages. Exactly one sender moves at each stage and the order of moves is $1 \rightarrow 2 \rightarrow 3 \rightarrow 3 \rightarrow 2$. In words, sender 1 moves at the first stage, sender 2 moves at the second stage, sender 3 moves at the third and fourth stages, and then sender 2 moves again at the fifth stage. By **Proposition 6**, the game has the same set of stable beliefs as the game with three stages and the order of moves is $1 \rightarrow 3 \rightarrow 2$. The intuition is very simple. Consider the incentive of a sender who can speak at stages t_1 and t_2 , where $t_2 > t_1$. He may prefer to gradually disclose at multiple stages for two reasons. First, he may want to withhold information at t_1 but release it at t_2 to avoid triggering undesirable disclosure of his opponents who move in between. Second, he may want to respond to the experiments of some senders, which are only observed at t_2 . However, neither of these concerns is sufficient to rationalize gradual information disclosure

in our model. The first concern is inconsistent with the concept of Nash equilibrium. When it comes to the second one, whatever the sender can disclose at early stages can also be disclosed at the last stage, making it redundant to speak multiple times. This is due to the fact that a sender can deliver as much information to the decision maker as he wants.

Proposition 6 implies that if we begin with a game with n rounds of persuasion and n senders moving in the order $1, \dots, n$ and add a move for sender i that precedes his move in the initial game, then the set of stable beliefs is unaffected. In contrast, if the additional move comes after player $i + 1$, then the stable beliefs could change. However, in this case we can remove the move in the initial game, so the number of moves is irrelevant for the set of stable beliefs, whereas the order of moves matters.

However, there is one case in which multiple moves can be useful. Suppose that we start with a game in which $1 \rightarrow 2 \rightarrow 3$, so that each player moves only once. Change the game to $2 \rightarrow 1 \rightarrow 2 \rightarrow 3$, so that player 2 now moves first and third. By **Proposition 6**, the two games have the same set of stable beliefs. However, the two games may generate different equilibrium outcomes because the first mover in the game can choose a Bayes plausible distribution of stable beliefs. Hence, in a spirit similar to the literature on agenda-setting in political economy (Romer and Rosenthal (1978), McKelvey (1976), Chen and Eraslan (2017), and others), having the right sender speak first can be useful for the decision maker.

If the prior belief is stable, this choice does not matter, as any first mover is happy to not provide any information. If there are only two states, it is also irrelevant. This is because for any distribution of beliefs τ that is not finer than τ' , there is some μ in the convex hull of the support to τ' , a property that fails with more than two states. However, in general, it can be strictly better to be the first mover.

Notice that the claim is that adding a first move without giving up the existing turn is what is advantageous, whereas swapping a move from later in the game to position 1 may be disadvantageous, because then the relevant order of play changes, which may affect the set of stable beliefs. A simple example illustrating this first-mover advantage is provided in **Appendix B.2**.

Proposition 6 may seem at odds with some real-world institutions that allow for multiple rounds of counterarguments. However, information transmission is frictionless in our model, whereas constraints on the complexity of what can be communicated in a single argument seem likely to matter in many real-world settings. We believe that to justify multiple rounds of counterarguments, which are ubiquitous in legal settings and debates, one has to look beyond purely strategic considerations and consider information asymmetries or constraints on the complexity of what can be communicated in a single argument.

4.4 *Simultaneous versus sequential persuasion*

Now we fix the set of senders and the order of consultation. When the decision maker receives disclosures from senders sequentially, she can decide to what extent (if any)

to share the received information with subsequent senders. On the one hand, revealing this information disciplines subsequent senders' strategic information manipulation in a certain manner. On the other hand, as long as the decision maker's information remains imperfect, revealing this information allows subsequent senders to make targeted opportunistic disclosures. A natural starting point to study this question is to compare two extreme cases: the one in which each sender observes all suggestions made by previous senders and the one where a sender observes no suggestions by other senders. The Bayesian persuasion game of the first policy corresponds to our baseline model, whereas the second policy corresponds to [Gentzkow and Kamenica \(2017a\)](#), where senders choose their experiments simultaneously, and where each sender may make their experiment arbitrarily correlated with any other experiment. We conclude that information revealed in the simultaneous game cannot be essentially less informative than in the sequential game.

Suppose that $\tau \in \Delta(\Delta(\Omega))$ is an equilibrium distribution of beliefs in a simultaneous move persuasion game. By Proposition 2 in [Gentzkow and Kamenica \(2017a\)](#), we know that this is true if and only if for each μ in the support of τ and for each player i , the payoff from μ is weakly higher than for any mean-preserving spread τ' of μ . Additionally, we use the same reasoning as in the sequential setup to prove that we may restrict attention to distributions with support on X .

PROPOSITION 7. *Suppose that $\tau \in \Delta(\Delta(\Omega))$ is an equilibrium distribution of beliefs in a simultaneous-persuasion game. Then there exists an outcome-equivalent equilibrium in which $\tau' \in \Delta(X)$.*

Hence, the difference between the sequential model and the simultaneous model boils down to a comparison that can be done vertex belief by vertex belief. A vertex belief in the support of an equilibrium of the sequential model must be unimprovable with respect to Bayes plausible deviations over the set of stable beliefs, that is, vertex beliefs that no sender would like to further refine. In contrast, a belief in the support of an equilibrium in the simultaneous move game must be unimprovable with respect to any Bayes plausible deviation.

It thus follows that for both the simultaneous game and the sequential game, we need to make sure that there is no vertex belief such that an admissible mean-preserving spread is preferred to a sender. The difference is thus that we have to check stability against arbitrary mean-preserving spreads in the simultaneous model, whereas some mean-preserving spreads can be ruled out in the sequential model because they would be undone by future senders. The following proposition therefore follows.⁹

PROPOSITION 8. *For generic preferences, there exists no pure-strategy equilibrium in the simultaneous game that is essentially less informative than the equilibrium in the sequential game.*

⁹A similar comparison is made in the multi-sender cheap talk literature. The conditions under which a fully revealing equilibrium exists is weaker in a simultaneous move cheap talk model than a sequential move one. See [Ambrus and Takahashi \(2008\)](#), [Battaglini \(2002\)](#), [Kawai \(2015\)](#), and [Krishna and Morgan \(2001\)](#).

PROOF. Suppose that the simultaneous game has an equilibrium essentially less informative than the finest equilibrium in the sequential game. Then there exists an μ such that (i) it is in the support of the equilibrium of the simultaneous move game and (ii) it is in the interior of the convex hull of the beliefs in the support of the finest equilibrium in the sequential move game. Since preferences are generic, μ cannot be stable belief in the sequential move game. Hence, some sender in the simultaneous move game has a profitable deviation, a contradiction. \square

There are two important caveats to [Proposition 8](#). Arbitrarily correlated experiments must be allowed and it applies only to pure-strategy equilibria in the simultaneous game. For the sequential game, generic preferences rule out mixed strategies and arbitrarily correlated experiments are without loss of generality, but this is not so in the simultaneous game.

The needs for pure strategies and arbitrary correlation are related. Together these assumptions imply that a simultaneous move equilibrium must be immune to profitable deviations at any realized signal in support of the joint equilibrium experiment. When experiments are independent or mixed strategies are used, it is impossible to fine tune deviations in this way.

When arbitrary correlation is violated, [Li and Norman \(2018\)](#) show that adding a sender may result in a strict loss of information. From [Proposition 4](#), we know that we cannot lose information in the sequential setting by adding the second player at the top, so the example combined with [Proposition 4](#) generates an explicit example where the sequential game is more informative when signals are independent. Similarly, [Li and Norman \(2018\)](#) provide an example in which a strictly less informative mixed-strategy equilibrium emerges when a sender is added. Again combining with [Proposition 4](#), we obtain an example with an equilibrium in the sequential model being more informative than an equilibrium in the simultaneous model while still allowing for arbitrarily correlated signals.

In each counterexample above, a fully revealing equilibrium also exists in the simultaneous game. In a related setting, [Hu and Sobel \(2019\)](#) argue that when multiple equilibria exist, this is not the most plausible equilibrium, because agents use strategies that are eliminated by iterations on weak dominance. However, this problem does not apply to [Proposition 4](#) as it is for *any* equilibrium in the simultaneous model.

Just as in the case of adding senders, the incompleteness of Blackwell's ordering implies that experiments may be noncomparable. However, we can again obtain a sharp characterization for the case with two states.

PROPOSITION 9. *Suppose that $\Omega = \{\omega_0, \omega_1\}$ and that there is an essentially unique equilibrium in the sequential game. Then any pure-strategy equilibrium in the simultaneous move game is weakly essentially more informative.*

The proof is similar to that of [Proposition 5](#) and is relegated to the Appendix. While non-Blackwell comparable distributions also exist in the case of two states, it is immediate to see that if the result fails, there is some belief μ in the support of an equilibrium

with simultaneous moves that lies strictly between the smallest and the largest beliefs in the support of the equilibrium with sequential moves. But then at least one sender must have an incentive to split the beliefs onto the smallest and the largest sequential move beliefs. Otherwise there must be an indifference, which is ruled out in the generic case. Again, [Figure 4](#) illustrates how the two state case is different from the general case.

We can also compare payoffs between simultaneous and sequential games. An implication of [Proposition 9](#) is that the last sender prefers the sequential move game to the simultaneous move game. The same is true for the general model whenever equilibria can be ranked using the Blackwell order. Hence, the persuasion framework generates the opposite result compared to duopolistic quantity competition. An intuition for this is that the reason why the Stackelberg leader is better off and the follower is worse off than under Cournot competition is that there is commitment value to overproduction, which allows the leader to grab a larger share of the pie. In contrast, in the persuasion model, the follower can always refine whatever the leader does. It is for this reason that the follower is made better off than in the simultaneous move game. Whether senders moving earlier are made better or worse off than in the simultaneous game is ambiguous.

4.5 Fully-revealing equilibria

A shortcut to the optimal design of the consultation structure problem is to look for conditions under which full revelation is an equilibrium. Then the decision maker can select senders and organize the order of moves to satisfy the conditions and achieve the complete information payoff.

Thanks to the one-step vertex characterization of the equilibrium outcome, we can identify an easy-to-check sufficient condition for when the unique equilibrium is fully revealing. One can rule out non-fully-revealing equilibria as long as at each nondegenerate vertex belief, there exists at least one sender who prefers full revelation to the current belief being observed by the decision maker.

PROPOSITION 10. *All equilibria are fully revealing if for each nondegenerate $\mu \in X$, there exists a sender i such that*

$$v_i(\sigma_d(\mu), \mu) < \sum_{\omega \in \Omega} u_i(\sigma_d(\delta_\omega, \omega))\mu(\omega), \quad (9)$$

where δ_ω is the degenerate belief about state ω .

Given the characterization of equilibrium outcomes in terms of stable vertex beliefs, the proof is obvious, so it is omitted. It is easy to check condition (9) as it depends only on the decision maker's strategy and the current sender's payoff at a small number of vertices. Although persuasion is sequential, the one-step characterization makes it unnecessary to take the subsequent senders' actions into account, which explains why the condition is order invariant (it also applies to the simultaneous model, and the case of both sequential and simultaneous moves).

Proposition 10 suggests a simple method to achieve full revelation. The decision maker selects senders in a such a way that the corresponding sequential-persuasion game does not have nondegenerate stable beliefs. To do so, it must be the case that every particular nondegenerate vertex belief is “disliked” by at least one sender.

It is worth mentioning that condition (9) applies regardless of the extensive form of the game. As discussed in Sobel (2013), in most multi-sender strategic communication models, a fully-revealing equilibrium exists under very weak conditions. The key reason is that when others fully reveal the state, a sender has no way to further affect the outcome. However, this means that full revelation can be supported as an equilibrium outcome even if it is Pareto dominated in a simultaneous move game, making the prediction less convincing. Some natural questions are prompted by this. In a multi-sender Bayesian persuasion game where senders move simultaneously, when should we expect full revelation as an equilibrium outcome if senders are coordinating on a plausible equilibrium, and under what conditions is full revelation the unique equilibrium outcome? **Proposition 10** offers some insight into these questions.

5. CONCLUDING REMARKS

We consider a sequential Bayesian persuasion model with multiple senders. Because it is without loss of generality to focus on equilibria corresponding to a finite set of beliefs, we establish that subgame perfect equilibria exist and generate a unique joint distribution over states and outcomes for generic preferences. Having a finite set of stable beliefs characterizing the equilibrium makes it easy to identify the unique equilibrium outcome and to apply the model to study changes in the extensive form. In particular, (i) adding a sender who moves first cannot reduce informativeness in equilibrium, and will result in a more informative equilibrium in the case of two states, (ii) it is without loss to let each sender speak only once, with the exception that the first mover may benefit from having a second move, and (iii) sequential persuasion cannot generate a more informative equilibrium than simultaneous persuasion and is less informative in the case of two states.

APPENDIX A: OMITTED PROOFS

A.1 Proofs: One-step equilibrium and equilibrium construction

PROOF OF PROPOSITION 1. To proceed, we extend the definition of one-step equilibrium to off the path of play.

DEFINITION 5. Consider a strategy σ' and let h_i be an arbitrary history when sender $i \in \{1, \dots, n-1\}$ moves. Also, for $j \geq i$, let $h'_j|_h_i$ be the implied continuation outcome path induced if each player $j \geq i$ follows σ'_j after history h_i and let $\sigma'|_{h_i}$ denote the continuation strategy profile.¹⁰ We say that $\sigma'|_{h_i}$ is *one-step* if $\bigvee_{j=i}^n \sigma'_j(h'_j|_h_i) = \sigma'_i(h_i)$.

¹⁰That is, $h'_i|_h_i = h_i$, $h'_{i+1}|_h_i = (h_i, \sigma'_i(h_i))$, $h'_{i+2}|_h_i = (h_i, \sigma'_i(h_i), \sigma'_{i+1}(h_i, \sigma'_i(h_i)))$, and so on.

Now we are ready to proceed. Fix a subgame perfect equilibrium σ^* and let $h_i = (\pi_1, \dots, \pi_{i-1})$ be an arbitrary history when i moves. Let $(\pi_i^*|h_i, \dots, \pi_n^*|h_i)$ be the continuation equilibrium path following h_i . Let

$$\pi^*|h_i = \left(\bigvee_{i=i}^{i-1} \pi_i \right) \vee \left(\bigvee_{i=i}^n \pi_i^* \Big|_{h_i} \right)$$

be the joint experiment generated by the continuation equilibrium path. Replace the continuation equilibrium strategies following h_i by $(\sigma'_i, \dots, \sigma'_n, \sigma'_d)$ where, on the continuation outcome path,

$$\begin{aligned} \sigma'_i(h_i) &= \pi^*|h_i \\ \sigma'_j(h_i, \pi^*|h_i, \dots, \pi^*|h_i) &= \pi^*|h_i \quad \text{for } j \in \{i+1, \dots, n\} \\ \sigma'_d(h_i, \pi^*|h_i, \dots, \pi^*|h_i, s) &= \sigma_d(h_i, (\pi_i^*|h_i, \dots, \pi_n^*|h_i), s). \end{aligned} \tag{A.1}$$

For a history in which i plays $\pi^*|h_i$ but some $j \in \{i+1, \dots, n\}$ deviates, let

$$\begin{aligned} \sigma'_k(h_i, \pi^*|h_i, \dots, \pi^*|h_i, \pi_j, \dots, \pi_k) &= \sigma_k^*(h_i, \pi_i^*|h_i, \dots, \pi_{j-1}^*|h_i, \pi_j, \dots, \pi_k) \\ \sigma'_d(h_i, \pi^*|h_i, \dots, \pi^*|h_i, \pi_j, \dots, \pi_n) &= \sigma_d^*(h_i, \pi_i^*|h_i, \dots, \pi_{j-1}^*|h_i, \pi_j, \dots, \pi_n), \end{aligned} \tag{A.2}$$

and for any other history, let

$$\begin{aligned} \sigma'_j(h_i, \pi_i, \dots, \pi_{j-1}) &= \sigma_j^*(h_i, \pi_i, \dots, \pi_{j-1}) \quad \text{for } j \in \{i+2, \dots, n\} \\ \sigma'_d(h_i, \pi_i, \dots, \pi_n, s) &= \sigma_d^*(h_i, \pi_i, \dots, \pi_n, s). \end{aligned} \tag{A.3}$$

The decision maker plays an optimal response following any path of play after h_i , as after each continuation path, the response is selected as some response for an identical joint experiment. Moreover, if each $j \geq i$ plays in accordance with σ'_j , it follows from (A.1) that the implied distribution over $\Omega \times \mathcal{A}$ is identical if each $j \geq i$ plays in accordance with the original equilibrium σ^* . Also, the strategies in (A.3) imply that the continuation play after a deviation by i is the same under σ' as under σ^* , so i has no incentive to deviate. As σ^* is subgame perfect, the continuation play in (A.3) is trivially subgame perfect. Finally, (A.2) implies that if j is the first player after i to deviate from $\pi^*|h_i$, then continuation play replicates that after the same deviation from the σ^* equilibrium following history $(h_i, \pi_i^*|h_i, \dots, \pi_{j-1}^*|h_i)$ in the original equilibrium, so $j \in \{i+1, \dots, n\}$ have no incentives to deviate. Clearly, σ' is not one-step after any history, but i and h_i were arbitrary, so adjusting σ^* in accordance with (A.1), (A.2), and (A.3) following any history i and h_i , we obtain a subgame perfect strategy profile that is one-step after every history h with the same equilibrium outcome. \square

PROOF OF LEMMA 1. For each program on form (3), we consider a restricted *finite* linear program

$$\begin{aligned} \tilde{V}_n(\mu) &= \max_{\tau \in \Delta(X)} \sum_{\mu' \in X} v_n(\sigma_d(\mu'), \mu') \tau(\mu') \\ \text{s.t. } &\sum_{\mu'} \mu' \tau(\mu') = \mu, \end{aligned} \tag{A.4}$$

where X is defined in (4). Hence, (A.4) is well defined as it is a finite-dimensional bounded linear program.

Pick any feasible solution τ to program (3). For each $a \in A$, define $\widehat{M}(a) \subset M(a)$ as the beliefs under which the decision maker takes action a : $\widehat{M}(a) = \{\mu \in \Omega \mid \sigma_d(\mu) = a\}$. Since $\widehat{M}(a) \subset M(a)$, it follows that for each $\mu' \in \widehat{M}(a)$, there exists $\lambda' \in \Delta(\{\mu_j^a\}_{j=1}^{J(a)})$ such that $\mu' = \sum_{j=1}^{J(a)} \lambda'_j \mu_j^a$. Hence, all beliefs that generate action a under τ may be split onto the vertices of $M(a)$ and aggregated into

$$\sum_{j=1}^{J(a)} \widehat{\tau}(\mu_j^a) = \sum_{\mu' \in \widehat{M}(a)} \tau(\mu') \sum_{j=1}^{J(a)} \lambda'_j = \sum_{\mu' \in \widehat{M}(a)} \tau(\mu').$$

Since it is possible that $v_n(a, \mu_j^a) < v_n(a', \mu_j^a)$ for some $\mu_j^a \in M(a)$ (and $\mu_j^a \notin \widehat{M}(a)$, because breaking the tie in favor of a' may be better than a), it follows that the solution to (A.4) satisfies

$$\begin{aligned} \tilde{V}_n(\mu) &\geq \sum_{a \in A} \sum_{j=1}^{J(a)} v_n(a, \mu_j^a) \widehat{\tau}(\mu_j^a) = \sum_{a \in A} \sum_{j=1}^{J(a)} \sum_{\omega \in \Omega} u_n(a, \omega) \mu_j^a(\omega) \widehat{\tau}(\mu_j^a) \\ &= \sum_{a \in A} \sum_{\omega \in \Omega} u_n(a, \omega) \sum_{j=1}^{J(a)} [\mu_j^a(\omega) \lambda'_j] \left[\sum_{\mu' \in \widehat{M}(a)} \tau(\mu') \right] \\ &= \sum_{a \in A} \sum_{\omega \in \Omega} u_n(a, \omega) \mu' \left[\sum_{\mu' \in \widehat{M}(a)} \tau(\mu') \right] \\ &= \sum_{\mu'} v_n(\sigma_d(\mu'), \mu') \tau(\mu'). \end{aligned} \tag{A.5}$$

This holds for any feasible solution to (3). Hence, $\tilde{V}_n(\mu) \geq V_n(\mu)$. Moreover, any optimal solution to (A.4) is a feasible solution to (3), so $\tilde{V}_n(\mu) \leq V_n(\mu)$. This establishes that solutions to (3) exist, and that $\tilde{V}_n(\mu) = V_n(\mu)$ and that every $\tau \in \Delta(X)$ that solves (A.4) also solves (3). Finally, if τ solves (3) and μ' is such that $\tau(\mu') > 0$, there can be no $\mu_k^a \in M(a)$ such that $v_n(a, \mu_k^a) < v_n(a', \mu_k^a)$ and $\lambda'_k > 0$ for the weight on vector μ_k^a in the convex combination such that $\mu' = \sum_{j=1}^{J(a)} \lambda'_j \mu_j^a$. This is seen from noting that this would generate a strict inequality in the first inequality of (A.5). \square

PROOF OF LEMMA 2. Proposition 1 implies that for every subgame perfect equilibrium, there is an outcome-equivalent equilibrium in which strategies are one-step for every

history, so we assume that σ^* is such a strategy profile. Suppose that there is a sender i and history h_i with associated continuation experiment $\pi^*|_{h_i}$ such that there exists some realization s' of experiment $\pi^*|_{h_i}$ that induces a decision maker posterior belief $\mu' \notin X$ with positive probability. Let $a' = \sigma_d(h_i, \pi^*|_{h_i}, \dots, \pi^*|_{h_i})$ be the equilibrium action induced by s' . Furthermore, let $M(a')$ be the belief polytope where a' is optimal and let $X(a') = \{\mu_j^a\}_{j=1}^m$ be the set of vertices of $M(a')$. Since $M(a')$ is the convex hull spanned by $X(a')$, there exists $\lambda \in \Delta(X(a'))$ such that $\mu' = \sum_{j=1}^m \lambda_j \mu_j^a$. Consider an alternative one-step strategy with $\pi^*|_{h_i}$ replaced by some π' in which the realization s' is replaced by the set $\{s_1, \dots, s_m\}$, where each s_j generates posterior μ_j^a and has unconditional probability $p(s')\lambda_j$, and everything else in π' is like the original equilibrium.¹¹ We also assume that the decision maker follows a strategy in which

$$\sigma'_d(h, s) = \begin{cases} a' & \text{if } h = (h_i, \pi', \dots, \pi') \text{ and } s \in \{s_1, \dots, s_m\} \\ \sigma_d^*(\pi^*|_{h_i}, \dots, \pi^*|_{h_i}, s) & \text{if } h = (h_i, \pi', \dots, \pi'), \text{ and } s \neq s' \\ & \text{is a realization of } \pi^* \\ \sigma_d^*(\pi^*|_{h_i}, \dots, \pi^*|_{h_i}, \pi_j, \dots, \pi_n, s) & \text{if } h = (h_i, \pi', \dots, \pi', \pi_j, \dots, \pi_n) \text{ where} \\ & j \geq i \text{ is the first player playing } \pi_j \neq \pi' \\ \sigma_d^*(h, s) & \text{for any other } h, \end{cases}$$

where σ_d^* is the strategy of the decision maker in the original equilibrium. For each $\mu_j^a \in M(a')$, σ'_d is a best response if σ_d^* is a best response. Also assume that all senders with $j < i$ follow the original equilibrium strategy σ_i^* and that sender $j = \{i, \dots, n\}$ plays

$$\sigma'_j(h_j) = \begin{cases} \pi' & \text{if } h_j = (h_i, \pi', \dots, \pi') \\ \sigma_i^*(h_i, \pi^*, \dots, \pi^*, \pi_k, \dots, \pi_{j-1}) & \text{if } h_j = (h_i, \pi', \dots, \pi', \pi_k, \dots, \pi_{j-1}) \\ \sigma_i^*(h_j) & \text{if } h_j = (h_i, \pi_i, \dots, \pi_{j-1}) \text{ is such that } \pi_i \neq \pi', \end{cases}$$

and leaves everything as in the original equilibrium if h_i is not played by $\{1, \dots, i - 1\}$. The continuation outcome path following h_i is then (π', \dots, π') and

$$v_i(a, \mu) = \sum_{j=1}^m \lambda_j v_n(a, \mu_j^a) = \sum_{j=1}^m \lambda_j v_n(\sigma_d(\pi', \dots, \pi', s_j), \mu_j^a),$$

while nothing is changed for signal realizations that are kept as in π^* , so the distribution over states and outcomes is the same as in the original equilibrium if no player deviates after h_i . Moreover, if $j \geq i$ is the first sender deviating from playing π' to π_j , the path of play replicates what happens if j is the first sender to deviate from π^* to π_j in the original continuation equilibrium. Hence, there is no profitable deviation on the path. Finally, off-path play replicates off-path continuation play in the original equilibrium, so there is no profitable deviation off the path. Repeating the same argument for each

¹¹It is possible that $\lambda_j = 0$ for some j . Instead of eliminating these beliefs, we may simply generate a probability zero signal so as not to treat this case separately.

history h_i , every continuation experiment $\pi^*|_{h_i}$ and every realization s' of $\pi^*|_{h_i}$ with corresponding belief $\mu' \notin X$ completes the proof. \square

PROOF OF PROPOSITION 2. In what follows, we construct a subgame perfect equilibrium where sender i 's equilibrium strategy coincides with the solution to program (7). That is, every sender i adds no information as long as $\mu \in X_i$ and posts an experiment that induces beliefs on X_i after any history.

Fix a pair (σ_d, τ_n) such that the following statements hold:

- The variable σ_d is optimal for the decision maker and breaks the ties in favor of sender n .
- We have $\tau_n : \Delta(\Omega) \rightarrow \Delta(X_n)$, so that only vertex beliefs are induced following any history, which is without loss by Lemma 1. Additionally, τ_n leaves any belief in X_n unchanged, so that $\tau_n(\mu|\mu) = 1 \forall \mu \in \Delta(X_n)$.

Sender $n - 1$'s problem can then be formulated as

$$\begin{aligned}
 V_{n-1}(\mu) = \max_{\tau} & \left[\sum_{\mu' \in \Delta(\Omega)} \left(\sum_{\mu'' \in \Delta(X_n)} v_{n-1}(\sigma_d(\mu''), \mu'') \tau_n(\mu''|\mu') \right) \tau(\mu'|\mu) \right] \\
 \text{s.t.} & \sum_{\mu' \in \Delta(\Omega)} \mu' \tau(\mu'|\mu) = \mu.
 \end{aligned}
 \tag{A.6}$$

That is, sender $n - 1$ chooses a mean-preserving spread that splits an interim belief μ into some updated interim beliefs τ , and for each induced interim belief μ' in τ , sender n further splits it into $\Delta(X_n)$ according to the selected τ_n .

Fix an arbitrary interim belief μ and a feasible strategy τ for program (A.6). Additionally, let τ_n be an any best response by player n that induces vertex beliefs only following any history and also satisfies $\tau_n(\mu|\mu) = 1$ for any $\mu \in X_n$. Together, τ_n and τ induce a compound mean-preserving spread $\tau_{n-1} : \Delta(\Omega) \rightarrow \Delta(X_n)$ from τ and τ_n defined as

$$\tau_{n-1}(\mu''|\mu) = \sum_{\mu' \in \Delta(\Omega)} \tau_n(\mu''|\mu') \tau(\mu'|\mu).$$

Since sender n always splits beliefs onto vertices, every compound mean-preserving spread τ_{n-1} has support on vertex beliefs only. Hence, every feasible solution to program (A.6) is feasible also in the restricted program (7) for $i = n - 1$, so

$$\tilde{V}_{n-1}(\mu) \geq V_{n-1}(\mu)$$

for every μ . In program (A.6), it is feasible to choose any mean-preserving spread $\tau \in \Delta(X_n)$. Since sender n does not add information when $\mu \in X_n$,

$$\tilde{V}_{n-1}(\mu) \leq V_{n-1}(\mu),$$

holds for every μ . Notice that this inequality crucially relies on our restriction on behavior on X_n . If sender n adds information at some interim belief $\mu \in X_n$, some feasible mean-preserving spreads in program (7) may no longer be feasible in program (A.6).

Consequently, $V_{n-1}(\cdot) = \tilde{V}_{n-1}(\cdot)$.¹² Since $\tilde{V}_{n-1}(\mu)$ is well defined, an optimal mean-preserving spread τ_{n-1} exists for $n - 1$ and has support on X_{n-1} . Whenever there exist multiple τ_{n-1} , we select ones such that sender $n - 1$ adds no information at every $\mu \in X_{n-1}$, ensuring that the best response of sender $n - 2$ is well defined. By induction, continuation strategies exist such that best responses for senders $1, \dots, n - 3$ are also defined. \square

A.2 Proofs: Outcome uniqueness

The proof of Proposition 3 has two parts. First, we state and prove a few intermediate results. Then we use these intermediate results to prove the uniqueness of equilibrium outcome.

A.2.1 Preliminaries The following corollary is more or less a direct consequence of Proposition 1.

COROLLARY 1. *Fix an equilibrium σ^* and a history h_i . For any deviation σ'_i by sender i , there exists a one-step continuation strategy profile σ^\dagger of senders $i + 1, \dots, n$ after history h_i such that the following statements hold:*

- (i) *Strategy profiles $(\sigma_i^\dagger, \dots, \sigma_n^\dagger, \sigma_d^*)$ and $(\sigma'_i, \sigma_{i+1}^*, \dots, \sigma_d^*)$ are outcome equivalent.*
- (ii) *Strategy profile $(\sigma_{i+1}^\dagger, \dots, \sigma_n^\dagger, \sigma_d^*)$ is a subgame perfect equilibrium of the continuation game after history $(h_i, \sigma_i^\dagger(h_i))$.*
- (iii) *The resulting posterior beliefs are vertices.*

PROOF. If $i = n$, then there is nothing to prove, so assume that $i < n$. A history h_i and a deviation strategy σ'_i by sender i yields a history $h'_{i+1} = (h_i, \sigma'_i(h_i))$. By Lemma 2, there exists an outcome-equivalent subgame perfect equilibrium in which senders $i + 1, \dots, n$ play one-step strategies σ^{**} with implied posterior beliefs with support on X after h'_{i+1} . Then by the same logic of the proof of Proposition 2, there is a one-step strategy profile $(\sigma_i^\dagger, \dots, \sigma_n^\dagger, \sigma_d^*)$ that yields the same outcome as $(\sigma'_i, \sigma_{i+1}^{**}, \dots, \sigma_n^{**}, \sigma_d^*)$. \square

We begin by ruling out non-Markov strategies for the decision maker. There are two pathological cases to address. First, it may be that there is some state $\omega \in \Omega$ in which the decision maker is indifferent between two actions. In that case, payoff-irrelevant aspects of the history can be used to construct non-Markov mixed strategies for the decision maker.¹³ The second case is that there is some interior vertex associated with some decision area $M(a)$ where both sender n and the decision maker are indifferent. Both these cases are rare in the sense that the associated payoff functions are measure zero subsets of all conceivable payoff functions.

¹²The argument does not use continuity of the objective function in (A.6). Best responses by n making the maximand for $n - 1$ discontinuous are admissible, as every choice by $n - 1$ ultimately results in a finite set.

¹³In a numerical example, we construct a non-Markov equilibrium where the decision maker's tie-breaking rule is determined by payoff-irrelevant endogenous choice senders. See Appendix B.1 for details.

LEMMA 3. *Pick any utility functions for the decision maker and sender n that belong to a set of full Lebesgue measure. Take any pair of histories h^*, h^{**} that generates the same posterior belief $\mu \in X$. Then the decision maker's equilibrium choice must be identical.*

PROOF. Consider an action a that is taken in equilibrium and some vertex $\mu_j^a \in M(a) \cap X$. Assume that there exist equilibria σ^* and σ^{**} , and histories h^*, h^{**} that generate joint experiments π^*, π^{**} with realizations $s^* \in \pi^*$ and $s^{**} \in \pi^{**}$ such that $\mu(s^*) = \mu(s^{**}) = \mu_j^a$, but that

$$\sigma_d^*(h^*, s^*) = a \neq a' = \sigma_d^{**}(h^{**}, s^{**}). \tag{A.7}$$

Suppose first that μ_j^a is a degenerate belief, i.e., a vertex of the simplex $\Delta(\Omega)$. Then there is some ω such that

$$v_d(a, \omega) = v_d(a', \omega). \tag{A.8}$$

A decision maker's payoff function may be viewed as an element in $|\Omega \times A|$ -dimensional Euclidean space and the payoff functions that satisfy (A.8) defines a $|\Omega \times A| - 1$ -dimensional subspace. As there is a finite number of triples $(a, a', \omega) \in A^2 \times \Omega$, the set of bounded Bernoulli payoff functions in which (A.8) holds for some triple (a, a', ω) is of Lebesgue measure zero. Next consider the case with (A.7) holding at some μ_j^a that is not a vertex of the simplex $\Delta(\Omega)$. Then sender n can deviate in a way so that either a or a' is chosen with probability arbitrarily close to 1, implying that

$$\sum_{\omega \in \Omega} v_n(a, \omega) \mu_j^a = \sum_{\omega \in \Omega} v_n(a', \omega) \mu_j^a, \tag{A.9}$$

which again defines a $|\Omega \times A| - 1$ -dimensional subspace of an $|\Omega \times A|$ -dimensional Euclidean space given any a, a' , and μ_j^a . There is a finite set of triples (a, a', μ_j^a) to consider and for each triple, (A.9) is satisfied for a set of payoff functions of Lebesgue measure zero, implying that the set of payoff functions for sender n that allows for multiple tie-breaking rules at an interior vertex is of measure zero. \square

By Lemma 3, for generic preferences, the decision maker must follow a Markov strategy of the form $\sigma_d : \Delta(\Omega) \rightarrow A$. It is then useful to define $\hat{v}_i : \Delta(\Omega) \rightarrow \mathbb{R}$, where

$$\hat{v}_i(\mu) \equiv v_i(\sigma_d(\mu), \mu),$$

which is the implied payoff function directly over decision maker beliefs for each sender i .

Next we show that for full measure of stable beliefs, no sender has a weak incentive to add information. To state this result, recall that X_i is the set of stable vertex beliefs in the truncated game starting with sender i .

LEMMA 4. *Suppose that the decision maker plays a Markov strategy $\sigma_d : \Delta(\Omega) \rightarrow A$. Then, for any sender $i \in \{1, \dots, n\}$ and for any $\mu \in X_i, Y \subseteq X_i$, and τ such that $\sum_{\mu' \in X_i} \mu' \tau(\mu') = \mu$, exactly one of the following two cases holds:*

(i) We have $\sigma_d(\mu') = \sigma_d(\mu)$ for every $(\mu, \mu') \in Y$.

(ii) There exists $(\mu, \mu') \in Y$ such that $\sigma_d(\mu) \neq \sigma_d(\mu')$. In this case,

$$\hat{v}_i(\mu) > \sum_{\mu' \in Y} \hat{v}_i(\mu') \tau(\mu') \tag{A.10}$$

for a set of sender i Bernoulli utility functions over $A \times \Omega$ with full Lebesgue measure.

PROOF. If $\sigma_d(\mu') = \sigma_d(\mu)$ for each $\mu \in X_i$ and every i , there is nothing to prove. Suppose instead that there exist $\mu \in X_i$, $Y \subset X_i$, and $\tau \in \Delta(Y)$ such that $\mu = \sum_{\mu' \in Y} \mu' \tau(\mu')$ and that (A.10) is violated for sender i . Denote $\{\mu_1, \dots, \mu_{m+1}\} = Y$ and $\tau = (\tau_1, \dots, \tau_{m+1})$, and write the failure of (A.10) as

$$\hat{v}_i(\mu) = \sum_{j=1}^{m+1} \hat{v}_i(\mu_j) \tau_j.$$

If Y is an affinely independent set, there is a unique mean-preserving spread of μ onto Y . In this case, the next step in which we find an affinely independent set that spans μ can be skipped. The case that requires more work is when Y is an affinely dependent set of vectors. This is true if and only if $\{\mu_2 - \mu_1, \dots, \mu_{m+1} - \mu_1\}$ are linearly dependent. Then there are scalars $(\alpha_2, \dots, \alpha_{m+1}) \neq (0, \dots, 0)$ such that $\sum_{j=2}^{m+1} \alpha_j(\mu_j - \mu_1) = 0$. So

$$\left(-\sum_{j=2}^{m+1} \alpha_j\right) \mu_1 + \sum_{j=2}^{m+1} \alpha_j \mu_j = \sum_{j=1}^{m+1} \alpha_j \mu_j = 0$$

by defining $\alpha_1 = -\sum_{j=2}^{m+1} \alpha_j$, which also implies that $\sum_{j=1}^{m+1} \alpha_j = 0$. For every β , we have

$$\mu = \sum_{j=1}^{m+1} \mu_j \tau_j = \sum_{j=1}^{m+1} \mu_j \tau_j - \beta \sum_{j=1}^{m+1} \alpha_j \mu_j = \sum_{j=1}^{m+1} (\tau_j - \beta \alpha_j) \mu_j.$$

Let $I^+ = \{j \in \{1, \dots, m+1\} \mid \tau_j > 0\}$ and let j^* be chosen so that $0 < \frac{\tau_j}{\alpha_j^*} \leq \frac{\tau_j}{\alpha_j}$ for all j such that $\alpha_j > 0$. Such j^* exists as there is at least one j such that $\alpha_j > 0$. Let $\beta^* = \frac{\tau_{j^*}}{\alpha_{j^*}^*}$ and

$$\tau_j^* = \tau_j - \frac{\tau_{j^*}}{\alpha_{j^*}^*} \alpha_j.$$

It follows that $\tau_j^* \geq 0$ for all j , that $\sum_{j=1}^{m+1} \tau_j^* = 1$ and $\tau_{j^*}^* = 0$. Hence, we can remove μ_{j^*} from $\{\mu_1, \dots, \mu_{m+1}\}$ and still find a convex combination that generates μ . By induction, there exists an affinely independent set of vectors $\{\hat{\mu}_1, \dots, \hat{\mu}_k\} \subseteq Y$ such that μ is in its convex hull, implying that there exists a unique solution $\hat{\tau}$ such that $\mu = \sum_{j=1}^k \hat{\mu}_j \hat{\tau}_j$.¹⁴

¹⁴If $\hat{\tau} \neq \hat{\tau}$ are distinct mean-preserving spreads of μ onto $\{\hat{\mu}_1, \dots, \hat{\mu}_k\}$, then $0 = \sum_{i=1}^k \hat{\mu}_i (\hat{\tau}_i - \hat{\tau}_i)$ or $0 = \sum_{i=2}^k (\hat{\mu}_i - \hat{\mu}_1) (\hat{\tau}_k - \hat{\tau}_k)$, which implies $\{\hat{\mu}_1, \dots, \hat{\mu}_k\}$ is affinely dependent, as $\hat{\tau}_i - \hat{\tau}_i \neq 0$ for at least one $i \in \{2, \dots, k\}$.

If $\sigma_d(\widehat{\mu}_j) = \sigma_d(\widehat{\mu}_{j'})$ for every pair of beliefs in $\{\widehat{\mu}_1, \dots, \widehat{\mu}_k\}$, then μ and $\widehat{\tau}$ are outcome equivalent. If $\sigma_d(\widehat{\mu}_j) \neq \sigma_d(\widehat{\mu}_{j'})$ for some beliefs in $\{\widehat{\mu}_1, \dots, \widehat{\mu}_k\}$,

$$\widehat{v}_i(\mu) = \sum_{j=1}^k \widehat{v}_i(\widehat{\mu}_j) \widehat{\tau}_j,$$

then $\widehat{v}_i : \Delta(\Omega) \rightarrow R$ belongs to a Lebesgue measure zero set of utility functions.¹⁵ We conclude that for every affinely independent subset of X_i , there is a Lebesgue measure zero of utility functions for i that can generate indifference that are not outcome equivalent. There is a finite number of affinely independent subsets and every mean-preserving spread of μ with support on X_i can be written in the form

$$\mu = \sum_{l=1}^L \beta_l \sum_{j=1}^{k(j)} \widehat{\mu}_j(l) \tau_j(l),$$

where $\beta_l \geq 0$ for each l , $\sum_{l=1}^L \beta_l = 1$ and every set $\{\widehat{\mu}_1(j), \dots, \widehat{\mu}_k(j)\}$ is affinely independent. Hence, if (A.10) holds for every affinely independent subset of X_i , it holds for all subsets of X_i . The result follows. \square

The first case of Lemma 4 simply points out that it is possible that the decision maker action is constant on a subset of stable beliefs. This is relevant because it is possible that there may exist a nontrivial mean-preserving spread $\tau \in \Delta(X_i)$ of $\mu \in X_i$ and if $\sigma_d(\mu') = \sigma_d(\mu)$ for each μ' in the support of τ , the sender is indifferent. However, this multiplicity is not essential because staying on μ or splitting beliefs in accordance with τ generates identical joint distribution over actions and states.

In the second case of Lemma 4, X_i , the set of stable beliefs of a sequential game played by senders $i, i + 1, \dots, n$, contains beliefs that result in at least two distinct actions according to σ_d . Suppose that $\tau \in \Delta(Y)$ is a vector such that (A.10) does not hold, implying that

$$\widehat{v}_i(\mu) = \sum_{\mu' \in Y} \widehat{v}_i(\mu') \tau(\mu'), \tag{A.11}$$

as otherwise μ could not be a stable belief. If Y is an affinely independent set of vectors, there is a unique mean-preserving spread of μ onto Y and it should be clear that (A.11) can only hold for a nongeneric set of functions $\widehat{v}_i : \Delta(\Omega) \rightarrow \mathbb{R}$.¹⁶ If, instead, Y is an affinely dependent set, then there must be an affinely independent subset of Y such that (A.11) holds for some mean-preserving spread with support on the affinely independent subset. For each affinely independent subset of Y , this requires nongeneric preferences, and since there is a finite number of senders and affinely independent subsets, the result follows by induction.

¹⁵By repeating the steps in (A.12), (A.13), and (A.14) below, the measure zero condition in belief space implies measure zero in terms of maps $u_i : A \times \Omega \rightarrow \mathbb{R}$.

¹⁶This also implies that a nongeneric set of Bernoulli utility functions $u_i : A \times \Omega \rightarrow \mathbb{R}$ can satisfy the equality.

In a similar spirit, we establish that indifferences over distinct distributions over stable continuation beliefs are rare.

LEMMA 5. Fix any $i \in \{1, \dots, n\}$. Then

$$\sum_{\mu' \in Y} \hat{v}_i(\mu') \tau(\mu') \neq \sum_{\mu' \in \tilde{Y}} \hat{v}_i(\mu') \tilde{\tau}(\mu')$$

for every $\mu \in X \cup \{\mu_0\}$ and every distinct pair $(\tau, Y), (\tilde{\tau}, \tilde{Y})$ with $Y \subseteq X_i$ and $\tilde{Y} \subseteq X_i$ being affinely independent sets, and τ ($\tilde{\tau}$) being a the unique mean-preserving spread of μ onto Y (\tilde{Y}) holds for a set of sender i Bernoulli utility functions over $A \times \Omega$ with full Lebesgue measure.

PROOF. Let $X(\mu_0)$ be the support for the unique equilibrium given prior μ_0 and let τ be the associated equilibrium distribution. We note that τ and λ are unique vectors so that

$$\begin{aligned} \mu_0 &= \sum_{\mu \in X(\mu_0)} \mu \tau(\mu) \\ \tilde{\mu}_0 &= \sum_{\mu \in X(\mu_0)} \mu \lambda(\mu). \end{aligned}$$

Hence, for any β ,

$$\mu_0 = \sum_{\mu \in X(\mu_0)} \mu (\tau(\mu) - \beta \lambda(\mu)) + \beta \tilde{\mu}_0,$$

and all coefficients are positive if β is small enough. Also, we assume that $\tilde{\tau}$ has support on $X(\tilde{\mu}_0) \neq X(\mu_0)$ so that

$$\tilde{\mu}_0 = \sum_{\mu \in X(\tilde{\mu}_0)} \mu \tilde{\tau}(\mu).$$

This implies that when the prior is μ_0 , it is feasible to split beliefs over $X(\mu_0) \cup X(\tilde{\mu}_0)$ in accordance with

$$\{\tau(\mu) - \beta \lambda(\mu) + \beta \tilde{\tau}(\mu)\}_{\mu \in X(\mu_0) \cup X(\tilde{\mu}_0)},$$

provided that β is small enough. But since τ is the generically unique equilibrium given μ_0 , this is suboptimal, so

$$\begin{aligned} \sum_{\mu \in X(\mu_0)} \hat{v}_1(\mu) \tau(\mu) &> \sum_{\mu \in X(\mu_0) \cup X(\tilde{\mu}_0)} \hat{v}_1(\mu) [\tau(\mu) - \beta \lambda(\mu) + \beta \tilde{\tau}(\mu)] \\ &= \sum_{\mu \in X(\mu_0)} \hat{v}_1(\mu) \tau(\mu) + \beta \left[\sum_{\mu \in X(\tilde{\mu}_0)} \hat{v}_1(\mu) \tilde{\tau}(\mu) - \sum_{\mu \in X(\mu_0)} \hat{v}_1(\mu) \lambda(\mu) \right]. \end{aligned}$$

Hence,

$$\sum_{\mu \in X(\tilde{\mu}_0)} \hat{v}_1(\mu) \tilde{\tau}(\mu) < \sum_{\mu \in X(\mu_0)} \hat{v}_1(\mu) \lambda(\mu),$$

which contradicts that $\tilde{\tau}$ is better than λ for prior belief $\tilde{\mu}_0$. □

A.2.2 Proof of Proposition 3 Lemma 2 and Corollary 1 imply that for sender $i = 2, \dots, n$, we need only to consider responses at X onto $\Delta(X_i)$. Lemma 4 implies that, generically, each sender has a strict incentive not to refine any $\mu \in X_i$. By linearity, an optimal mean-preserving spread with support on an affinely independent set must exist, so Lemma 5 implies that for generic preferences, each deviation onto $\Delta(X)$ generates an essentially unique response for generic preferences and, since every deviation is equivalent to a deviation onto $\Delta(X)$, we conclude that the off-equilibrium path is generically unique. Finally, Lemma 5 applied to sender 1 also implies that sender 1 generically has a unique optimal mean-preserving spread of the prior onto the set of stable beliefs.

Assume that there exist two distinct affinely independent sets of vectors $Y \subseteq X_i$ and $\tilde{Y} \subset X_i$ such that

$$\sum_{\mu' \in Y} \widehat{v}_i(\mu') \tau(\mu') = \sum_{\mu' \in \tilde{Y}} \widehat{v}_i(\mu') \tilde{\tau}(\mu'), \tag{A.12}$$

where τ is the unique mean-preserving spread of μ onto Y and $\tilde{\tau}$ is the unique mean-preserving of μ onto \tilde{Y} . Also assume there are at least two distinct actions chosen by the decision maker. In terms of the primitive preferences over $A \times \Omega$, (A.12) can be rewritten as

$$\sum_{\mu' \in Y} \sum_{\omega \in \Omega} [u_i(\sigma_d(\mu'), \omega) \mu'(\omega)] \tau(\mu') = \sum_{\mu' \in \tilde{Y}} \sum_{\omega \in \Omega} [u_i(\sigma_d(\mu'), \omega) \mu'(\omega)] \tilde{\tau}(\mu'). \tag{A.13}$$

Notice that if for each $a \in A$, we let $Y(a) = \{\mu' \in Y \text{ s.t. } \sigma_d(\mu') = a\}$ and symmetrically for $\tilde{Y}(a)$, we may rewrite (A.13) further as

$$\sum_{a \in A} \left\{ \sum_{\omega \in \Omega} u_i(a, \omega) \left[\sum_{\mu' \in X(\mu, a)} \mu'(\omega) \tau(\mu') - \sum_{\mu' \in \tilde{X}(\mu, a)} \mu'(\omega) \tilde{\tau}(\mu') \right] \right\} = 0. \tag{A.14}$$

Since τ and $\tilde{\tau}$ are unique, this defines a lower dimensional subspace of $|A \times \Omega|$ -dimensional Euclidean space, so the set of sender i payoff functions such that (A.12) holds is measure zero. Since X_i is finite, there is a finite set of pairs of affinely independent sets spanning μ and we consider only μ from the finite set $X \cup \{\mu_0\}$. The result follows.

A.3 Proofs: Applications

PROOF OF PROPOSITION 6. Since the stage and the player identity no longer coincide, let X_i^t denote the stable beliefs in the truncated game starting with player i moving at stage t . Suppose that t is the final move of player i and that i also moves at t' , with $t' < t$. If t' and t are consecutive stages, it is immediate that $X_i^t = X_i^{t'}$, so assume that there exists a player j moving in between t' and t . Without loss of generality, let j move at time $t' + 1$ and let $X_j^{t'+1} \subseteq X_i^t$ be the set of stable beliefs in the truncated game starting

with player j at time $t' + 1$. We claim that $X_i^{t'} = X_j^{t'+1}$, that is, that player i moving at t' does not affect the set of stable beliefs in the truncated game starting at the next stage, so the move by i at t' is redundant. For contradiction, assume that the move by i at t' refines the set of stable beliefs, so that there exists $\mu \in X_j^{t'+1}$ such that $\mu \notin X_i^{t'}$. But if $\mu \in X_j^{t'+1}$, then $\mu \in X_i^{t'}$, which implies that i has no incentive to create a mean-preserving spread of μ with support in $X_i^{t'} \subseteq X_j^{t'+1}$. Since any mean-preserving spread that is feasible at time t' is feasible also at t , this contradicts $X_i^{t'}$ being the set of stable beliefs in the truncated game starting a time t . Since $t' < t$ and i were arbitrary, the proposition follows. \square

PROOF OF PROPOSITION 7. Consider some μ in the support of τ that is not in $\Delta(X)$. Assume that $\sigma_d(\mu) = a$ is the action taken by the decision maker following μ and let $M(a)$ be the set of beliefs for which a is optimal. Replace μ with any mean-preserving spread τ' of onto beliefs in $M(a)$, suppose that $\sigma_d(\mu') = a$ for each μ' in the support of τ' , and let the probability of any other belief in τ be unchanged. Clearly, this belief distribution is outcome equivalent with τ . To see that it must also be an equilibrium, assume that it is not. Then there exists some player i and belief μ' in the support of τ' and a mean-preserving spread τ'' of μ' such that i strictly prefers τ'' to μ' . But then i strictly prefers the compound mean-preserving spread constructed by first splitting μ into τ and then further splitting μ' into τ'' . Since this compound mean-preserving spread is a feasible deviation for i given belief μ , this contradicts μ being in the support of an equilibrium distribution. Since τ' is any mean-preserving spread with support in $M(a)$, we may choose one with support on the vertices of $M(a)$, which is always possible. The proof is completed by noting that the argument can be repeated for any μ not in $\Delta(X)$. \square

PROOF OF PROPOSITION 9. Fix the prior μ_0 and begin by noting that for the result to fail some information must be provided in the sequential model. Hence, without loss there must be a pair $\mu_L, \mu_H \in X$ such that $\mu_L < \mu_0 < \mu_H$, where μ_L and μ_H are in the support of the equilibrium in the sequential model. Suppose that there is some μ with $\mu_L < \mu < \mu_H$ that is in the support of an equilibrium in the simultaneous move model. As in the proof of Proposition 8, there are two cases. First, suppose first that the action is the same at μ_L and μ_H . Then putting positive probability on μ or the unique mean-preserving spread onto $\{\mu_L, \mu_H\}$ has no effect on the distribution over actions and states, so putting positive probability on μ does not affect the essential informativeness. Second, suppose that μ_L and μ_H generate distinct actions. Then for μ to be part of an equilibrium in the simultaneous game, all senders must weakly prefer μ to the unique mean-preserving spread to $\{\mu_L, \mu_H\}$. But then μ must be an equilibrium (not necessarily on a vertex) in the sequential game, which since μ and the mean-preserving spread to $\{\mu_L, \mu_H\}$ generate different distribution over states and action contradicts essential uniqueness. Hence, an equilibrium in the simultaneous game is either more or equally informative as the finest equilibrium of the sequential game. \square

APPENDIX B: EXAMPLES

B.1 *Non-Markov equilibrium*

In this section, we consider an example that has a non-Markov equilibrium that is qualitatively different from the Markov equilibrium. Suppose that $\Omega = \{\omega_0, \omega_1\}$ and the optimal choice correspondence for the decision maker is

$$\sigma(\mu) = \begin{cases} \{a_1, a_2\} & \text{if } \mu \leq 1/10 \\ a_3 & \text{if } 0.1 \leq \mu \leq 9/10 \\ \{a_4, a_5\} & \text{if } \mu \geq 9/10. \end{cases}$$

Also suppose that two senders have state-independent preferences

$$u_1(a, \omega) = \begin{cases} 3 & \text{if } a \in \{a_1, a_4\} \\ 1 & \text{if } a = a_3 \\ 0 & \text{if } a \in \{a_2, a_5\}, \end{cases} \quad u_2(a, \omega) = \begin{cases} 3 & \text{if } a \in \{a_2, a_5\} \\ 1 & \text{if } a = a_3 \\ 0 & \text{if } a \in \{a_1, a_4\}. \end{cases}$$

Consider a Markov equilibrium. Allowing for mixed strategies, let $\sigma_1(0)$ be the probability for a_1 given belief $\mu = 0$ and let $\sigma_4(1)$ be the probability of a_4 given belief $\mu = 1$. Suppose that the decision maker has full information. Then the payoffs of sender 1 and sender 2 are $3[\sigma_1(0) + \sigma_4(1)]/2$ and $3[2 - \sigma_1(0) - \sigma_4(1)]/2$, respectively, so the payoff is greater than or equal to $3/2$ for at least one sender. Hence, beliefs in $[1/10, 9/10]$ can be ruled out in any Markov equilibrium. In contrast, if the decision maker always breaks the tie against the sender who first splits the belief into $[0, 1/10]$ or $[9/10, 1]$, each sender may as well not provide any information, and qualitatively different equilibria with action a_3 can be supported by such non-Markov strategies.

B.2 *First-mover advantage*

To illustrate the first-mover advantage, assume that there are three states, i.e., $\Omega = \{\omega_1, \omega_2, \omega_3\}$, and that the prior is $(1/3, 1/3, 1/3)$. For simplicity, take the set of stable beliefs as a primitive. We assume that the stable vertex beliefs are $e_1 = (1, 0, 0)$, $e_2 = (0, 1, 0)$, $e_3 = (0, 0, 1)$, $\mu_1 = (1/2, 1/2, 0)$, and $\mu_2 = (0, 1/2, 1/2)$. There can be an arbitrary number of senders, but we consider just two of them, labeled 1 and 2. Let their expected utilities evaluated at the stable beliefs be

$$\begin{aligned} (\widehat{v}_1(e_1), \widehat{v}_1(e_2), \widehat{v}_1(e_3), \widehat{v}_1(\mu_1), \widehat{v}_1(\mu_2)) &= (0, -1, -1, 0, 1) \\ (\widehat{v}_2(e_1), \widehat{v}_2(e_2), \widehat{v}_2(e_3), \widehat{v}_2(\mu_1), \widehat{v}_2(\mu_2)) &= (-1, -1, 0, 1, 0). \end{aligned}$$

While e_1, e_2 , and e_3 are trivially stable, we need to check the stability of μ_1 and μ_2 . We have that μ_1 is stable because

$$\begin{aligned} \widehat{v}_1(\mu_1) = 0 &> \frac{1}{2}\widehat{v}_1(e_1) + \frac{1}{2}\widehat{v}_1(e_2) = -\frac{1}{2} \\ \widehat{v}_2(\mu_1) = 1 &> \frac{1}{2}\widehat{v}_2(e_1) + \frac{1}{2}\widehat{v}_2(e_2) = -1, \end{aligned}$$

and μ_2 is stable by a symmetric computation. It follows that in the game in which sender 1 moves first, the equilibrium will be that sender 1 puts probability $1/3$ on e_1 and $2/3$ on μ_2 , giving player 1 an expected utility of $2/3$ and player 2 an expected utility of $-1/3$. In contrast, when sender 2 moves first, μ_1 is played with probability $2/3$ and e_3 is played with probability $1/3$, resulting in the opposite expected utilities.

REFERENCES

- Ambrus, Attila and Satoru Takahashi (2008), “Multi-sender cheap talk with restricted state spaces.” *Theoretical Economics*, 3, 1–27. [642, 657]
- Au, Pak Hung and Keiichi Kawai (2019), “Competitive disclosure of correlated information.” *Economic Theory*. [641]
- Au, Pak Hung and Keiichi Kawai (2020), “Competitive information disclosure by multiple senders.” *Games and Economic Behavior*, 119, 56–78. [641]
- Aumann, Robert J. and Michael B. Maschler (1995), *Repeated Games of Incomplete Information, the Zero-Sum Extensive Case*. MIT. [640]
- Battaglini, Marco (2002), “Multiple referrals and multidimensional cheap talk.” *Econometrica*, 70, 1379–1401. [642, 657]
- Bergemann, Dirk and Stephen Morris (2016), “Bayes correlated equilibrium and the comparison of information structures in games.” *Theoretical Economics*, 11, 487–522. [640, 651]
- Bhattacharya, Sourav and Arijit Mukherjee (2013), “Strategic information revelation when experts compete to influence.” *RAND Journal of Economics*, 44, 522–544. [642]
- Blackwell, David (1953), “Equivalent comparisons of experiments.” *Annals of Mathematical Statistics*, 24, 265–272. [643]
- Board, Simon and Jay Lu (2018), “Competitive information disclosure in search markets.” *Journal of Political Economy*, 126, 1965–2010. [641]
- Boleslavsky, Raphael and Christopher Cotton (2015), “Grading standards and education quality.” *American Economic Journal: Microeconomics*, 7, 248–279. [641]
- Boleslavsky, Raphael and Christopher Cotton (2018), “Limited capacity in project selection: Competition through evidence production.” *Economic Theory*, 65, 385–421. [641]
- Chen, Ying and Hülya Eraslan (2017), “Dynamic agenda setting.” *American Economic Journal: Microeconomics*, 9, 1–32. [656]
- Ely, Jeffrey, Alexander Frankel, and Emir Kamenica (2015), “Suspense and surprise.” *Journal of Political Economy*, 123, 215–260. [641]
- Ely, Jeffrey C. (2017), “Beeps.” *American Economic Review*, 107, 31–53. [641]
- Gentzkow, Matthew and Emir Kamenica (2017a), “Bayesian persuasion with multiple senders and rich signal spaces.” *Games and Economic Behavior*, 104, 411–429. [641, 647, 652, 657]

- Gentzkow, Matthew and Emir Kamenica (2017b), “Competition in persuasion.” *Review of Economic Studies*, 84, 300–322. [641, 642, 652]
- Glazer, Jacob and Ariel Rubinstein (2001), “Debates and decisions: On a rationale of argumentation rules.” *Games and Economic Behavior*, 36, 158–173. [642]
- Green, Jerry and Nancy Stokey (1978), “Two representations of information structures and their comparisons.” Institute for Mathematical Studies in the Social Sciences, Stanford University, Technical report no. 271. [642, 643, 647]
- Grünbaum, Branko, Volker Kaibel, Victor Klee, and Günter M. Ziegler (1967), *Convex Polytopes*. John Wiley and Sons, New York, New York. [647]
- Harris, Christopher (1985), “Existence and characterization of perfect equilibrium in games of perfect information.” *Econometrica*, 53, 613–628. [641]
- Hu, Peicong and Joel Sobel (2019), “Simultaneous versus sequential disclosure.” Unpublished Paper, University of California, San Diego. [642, 658]
- Hwang, Ilwoo, Kyungmin Kim, and Raphael Boleslavsky (2019), “Competitive advertising and pricing.” [641]
- Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian persuasion.” *American Economic Review*, 101, 2590–2615. [640, 641]
- Kartik, Navin, Frances Xu Lee, and Wing Suen (2017), “Investment in concealable information by biased experts.” *RAND Journal of Economics*, 48, 24–43. [642]
- Kartik, Navin, Frances Xu Lee, and Wing Suen (2019), “A theorem on Bayesian updating and applications to signaling games.” Unpublished Paper, Department of Economics, Columbia University. [642]
- Kawai, Keiichi (2015), “Sequential cheap talks.” *Games and Economic Behavior*, 90, 128–133. [642, 657]
- Krishna, Vijay and John Morgan (2001), “A model of expertise.” *Quarterly Journal of Economics*, 116, 747–775. [642, 657]
- Li, Fei and Peter Norman (2018), “On Bayesian persuasion with multiple senders.” *Economics Letters*, 170, 66–70. [641, 654, 658]
- Lipnowski, Elliot and Laurent Mathevet (2017), “Simplifying Bayesian persuasion.” Unpublished Paper, Columbia University. [642]
- Lipnowski, Elliot and Laurent Mathevet (2018), “Disclosure to a psychological audience.” *American Economic Journal: Microeconomics*, 10, 67–93. [642]
- McKelvey, Richard (1976), “Intransitivities in multidimensional voting models and some implications of agenda control.” *Journal of Economic Theory*, 12, 472–473. [656]
- Milgrom, Paul R. and John Roberts (1986), “Relying on the information of interested parties.” *Rand Journal of Economics*, 17, 18–32. [642]

Rayo, Luis and Ilya Segal (2010), “Optimal information disclosure.” *Journal of Political Economy*, 118, 949–987. [641]

Romer, Thomas and Howard Rosenthal (1978), “Political resource allocation, controlled agendas, and the status quo.” *Public Choice*, 33, 27–43. [656]

Sobel, Joel (2013), “Giving and receiving advice.” In *Advances in Economics and Econometrics: Tenth World Congress*, volume 1, 305–341, Cambridge University Press, New York, New York. [660]

Wu, Wenhao (2018), “Sequential Bayesian persuasion.” Unpublished Paper, University of Arizona. [641]

Co-editor Simon Board handled this manuscript.

Manuscript received 15 October, 2018; final version accepted 9 July, 2020; available online 20 July, 2020.