Sequential Persuasion*

Fei Li† Peter Norman ‡

December 29, 2019

Abstract

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. Finally, we provide a simple condition that guarantees that full revelation is the unique equilibrium outcome regardless of the ordering of senders.

Keywords: Bayesian Persuasion, Communication, Competition in Persuasion, Multiple Senders, Sequential Persuasion.

JEL Classification Codes: D82, D83

*We thank the Editor and 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 participants at the 2017 CETC conference in Vancouver, the 2017 International Conference on Game Theory, 2017 Midwest Economic Theory Conference in Dallas, and the 2018 NSF/CEME Decentralization Conference in Chicago for helpful comments. The usual disclaimer applies.

†University of North Carolina at Chapel Hill. lifei@email.unc.edu

‡*Corresponding Author*. University of North Carolina at Chapel Hill. normanp@email.unc.edu

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 a number of 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 often consultations are sequential. For 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 (1968) 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 for the equilibrium, but instead of actually refining the information on the path, their preferences restrict what the first sender does through incentive compatibility constraints. This works on and 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

[1] Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll. (Harvard Corp.), Civil Action No. 14-cv-14176-ADB, 2019 U.S. Dist. LEXIS 170309 (D. Mass. Sep. 30, 2019).

by a finite set of vertices. We demonstrate that it is without loss of generality for every sender to only provide information that generates beliefs on these vertices.

Focusing on one-step strategies with support on a finite set of vertices, we then use backwards 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 in order to demonstrate that for a set of preferences of full measure, there is a unique equilibrium distribution over states and outcomes.

The equilibrium characterization is also convenient for asking questions about the effects of changes to the extensive form, because equilibrium distributions are recursively defined as *stable (vertex) beliefs*. Concretely, for the last sender, a stable belief is a probability of the state of the world that the last sender has no incentive to further refine, and we know that it is without loss of generality to consider only vertices of the polytopes defining optimal actions for the decision maker. The penultimate sender understands that the last sender (without loss of generality) will refine whatever information provided into a distribution of such stable beliefs, so he may as well only consider distributions over the set of vertex beliefs for the last sender. Hence, the set of stable beliefs in the truncated game starting with the penultimate sender is the set of stable beliefs in the one-sender game for which there is no mean-preserving spread onto stable beliefs in the one-sender game that the penultimate sender prefers. The set of stable beliefs for the full game is constructed recursively from this idea, and it is weakly smaller for each step of the backwards induction process.

By studying these stable beliefs, we find that adding a sender *who moves first* cannot reduce the informativeness in equilibrium. In contrast, strategic considerations may reduce information disclosure if another sender is added later in the game. Next, we consider whether multiple rebuttals are useful 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 to move 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 for the equilibrium outcomes.

We also compare sequential and simultaneous persuasion. We find that sequential persuasion can never generate a more informative equilibrium than simultaneous persuasion. This result holds for any equilibrium in the simultaneous move game, so this is not subject to the criticism of Hu and Sobel (2019). Finally, we provide a simple and easily interpretable sufficient condition for when full revelation is the unique equilibrium, which is invariant of the order of moves.

Literature. Our paper relates to a large body of work on information disclosure, but is most directly connected to the growing 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. See Gentzkow and Kamenica (2017a,b), Boleslavsky and Cotton (2016), Au and Kawai (2017, 2020), and Hwang, Kim, and Boleslavsky (2019). However, none of these papers deal 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. The equilibrium outcome is unique, which makes comparisons more straightforward and convincing, and it allows us to discuss some applications of our model. Second, our model clarifies how senders' experiments are combined. This enables us to transparently compare information revelation of the game where senders move sequentially with that of the game where senders move simultaneously.

A growing body of work embeds persuasion into dynamic models (see Ely, Frankel, and Kamenica (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, but they consider a very different information structure. 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, Lee, and Suen (2016, 2017), Bhattacharya and Mukherjee (2013) and Milgrom and Roberts (1986). In a recent paper, Hu and Sobel (2019) also compare simultaneous information disclosure 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 impose restrictions in the information transmission such as asymmetric information, limited information process ability, restricted forms of signals, etc. Instead, the Bayesian persuasion framework we adopt eliminates all of these frictions and focuses on the sole effect of strategic interaction among senders on information provision. It thus serves as a theoretical benchmark for identifying sources of communication inefficiency.

Organization. The remainder of this paper is organized as follows. In Section 2, we describe the model. Section 3 characterizes the set of equilibria. We show 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. Omitted proofs are collected in Appendix A.

2 Model

Players. Consider an environment with senders $i = 1, ..., n$ and a decision maker d. Player $i = 1, ..., n, d$ has a utility function $u_i : A \times \Omega \to \mathbb{R}$ where A is a finite set of actions, and Ω is the 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

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

Experiments. Players are uninformed about the state of the world, but a sender may provide information to the decision maker by creating an *experiment*. For reasons discussed in Section 2.1 we adopt the *partition representation* of experiments. This models the state-dependent noise explicitly, which makes it clear how to combine multiple experiments. Specifically, an experiment assumes a state-contingent partition of $[0, 1]$, i.e., for each state ω, $\{\pi(s|\omega)\}_{s \in S}$ are disjoint sets such that $\cup_{s \in S} \pi(s|\omega) = [0, 1]$ where S indexes the sets in the partitions. Given experiment π, one can relabel partitions as *signals* and assign probabilities to signals according to the measure of the corresponding partition in each state. In doing so, one obtains the state-contingent distribution over signals $p_\pi : \Omega \to \Delta(S)$, and the probability of signal $s \in S$ being realized conditional on state ω is

$$p_\pi(s|\omega) = \lambda(\pi(s|\omega)), \tag{2}$$

where $\lambda(\cdot)$ denotes the Lebesgue measure and $\sum_s p_\pi(s|\omega) = 1$ for each $\omega \in \Omega$. An experiment can be *informative* since it partitions $[0, 1]$ differently in different states, making the probability

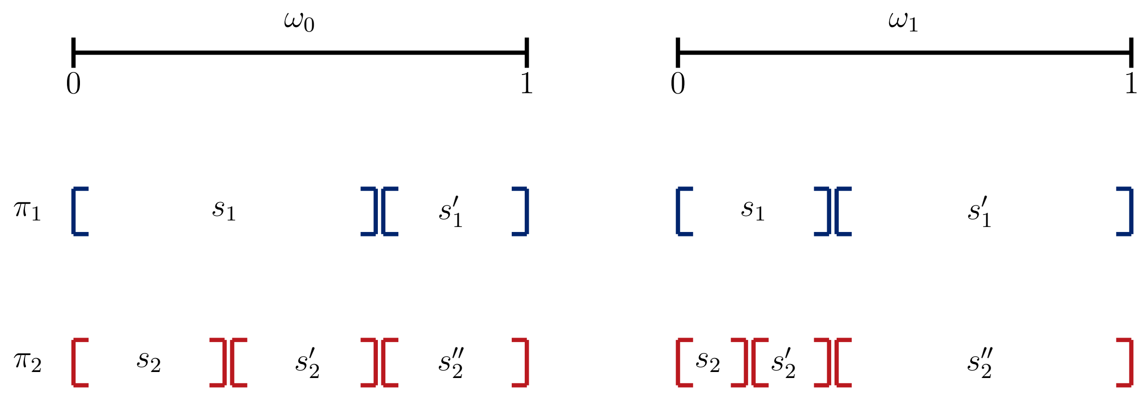

Figure 1: There are two states: ω_0 and ω_1 and two senders $i = 1, 2$. Sender 1's signal space contains two signals: s_1 and s_1'. Sender 2's experiment has three possible signals $\{s_2, s_2', s_2''\}$. It is evident that $\pi_1 \vee \pi_2 = \pi_2$.

realization of each signal state-dependent. We let Π denote the set of experiments and let $\succeq$ be the partial order in which $\pi \succeq \pi'$ if π is a finer partition than π'.[2] The space $(\Pi, \succeq)$ is a lattice and for any two experiments π, π', the join $\pi \vee \pi'$ is also an experiment finer than both π and π' , so combining two experiments by taking intersections generates a more informative new experiment. Figure 1 illustrates the simplicity of combining two experiments to get a new experiment using the partition representation.

Extensive Form. Senders move sequentially with sender $1, ..., n$ posting experiments $\pi_1 ...., \pi_n$, and each sender observing previous senders' experiments. Then nature draws ω. Finally, the decision maker observes $(\pi_1, ..., \pi_n)$ and a joint realization $s = (s_1, ..., s_n)$ according to the corresponding state-contingent probability $p_{\vee_i \pi_i}(s|\omega) = \lambda(\vee_i \pi_i(s|\omega))$ for $i = 1, ..., n$ and takes an action $a \in A$.

Strategies and Equilibrium. A pure *strategy* for sender i is a map $\sigma_i : \Pi^{i-1} \to \Pi$ where Π^0 is the trivial null history. That is, given a history $\{\pi_1, ..., \pi_{i-1}\}$, sender i chooses π_i that results in a finer experiment $\vee_{k=1}^{i} \pi_k$. A history for the decision maker is a vector $(\pi_1, .., \pi_n, s_1, ..., s_n)$. Let $\mathcal{H}_d$ be the set of all histories for the decision maker and $\sigma_d : \mathcal{H}_d \to A$ denote his 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.

[2]For any $\pi \succeq \pi'$, let p_π and $p_{\pi'}$ denote the corresponding state-contingent distributions over signals they generate respectively. It can be shown that there is a garbling that transforms p_π to $p_{\pi'}$ which is less informative in the sense of Blackwell (1953). See Green and Stokey (1978) for a proof.

2.1 Discussion

We make two unconventional assumptions. First, we adopt the partition representation of signal as in Green and Stokey (1978) and Gentzkow and Kamenica (2017b). Second, we assume nature moves after all senders. Before proceeding further, we briefly discuss these modelling choices.

On How to Represent Experiments. As shown by Green and Stokey (1978), it is without loss to interchangeably use the state-dependent distribution representation (S, p) and the partition representation π to model information disclosure in non-strategic settings. We decided to use the partition representation in our multi-sender setting as it is significantly more convenient for a number of reasons.

First, it transparently displays how a sender *adds information* to another sender's experiment. This is illustrated in Figure 1. Suppose that π_1 is the experiment chosen by sender 1. Given each signal, sender 2's experiment π_2 allows him to disclose additional information by further partitioning $\pi_1(s|\omega)$ for every $s \in S_1$ and every $\omega \in \Omega$. For instance, the interval $\pi_1(s_1|\omega)$ is further partitioned into $\pi_2(s_2|\omega)$ and $\pi_2(s_2'|\omega)$ for every ω. As long as $\lambda(\pi_2(s_2|\omega))/\lambda(\pi_1(s_1|\omega))$ varies over ω, observing signal s_2 (or s_2') helps to infer the underlying state conditional on the information generated by s_1.

In a single-sender persuasion model along the lines of Kamenica and Gentzkow (2011), it is equivalent to study the optimal information design problem and a reduced-form problem in which the choice variable is a distribution of beliefs. In a multi-sender model, things are less straightforward, and the representation of experiments, whether mixed strategies are allowed, and the extensive form all matter for how the decision maker aggregates information (see Gentzkow and Kamenica (2017b) and Li and Norman (2018)). One purpose of the current paper is to compare information revelation in multi-sender persuasion games where senders move sequentially with the simultaneous move benchmark (Gentzkow and Kamenica 2017b), that adopts the partition representation. To do this, one must (i) explicitly specify the experiments rather than jumping to the reduced-form belief-splitting game directly, and (ii) make the comparison fair in the sense that senders have access to similar information provision technologies in the two games. In our sequential model with the partition representation, the combination of multiple senders' experiments allows a sender to choose the distribution over signals conditional on the state and the signals of *previous* senders' experiments. This captures the idea that subsequent senders can observe and react upon the signals of the experiments chosen by previous senders. In contrast, when senders move simultaneously (Gentzkow and Kamenica 2017a), allowing senders to make refinements conditional on the realizations from other senders signals does not have this simple interpretation.

Signal Observability. We assume that nature moves after senders choose their experiments. Besides making subgame (rather than Bayesian) perfection applicable to the game, this specification also eases the presentation of the proofs. Because nature moves after all of the senders, a sender's partition choice must specify his information disclosure for each possible realized signal profile by precedent senders. Thus, we can consider histories while avoiding to keep track of signal realizations, which simplifies analyzing the impact on the continuation play from any particular experiment because partitions are much more easier to track than systems of conditional probability distributions.

In the sequential model of ours, a sender acts *as if* he observes and responds to the signal realizations of previous senders' experiments. Consider the experiments in Figure 1 once more. Sender 2's experiment π_2 allows him to disclose differential information conditional on the signal realization of π_1. As we pointed out, when the signal is s_1, π_2 adds additional information by further partitioning the corresponding intervals under both states into s_2 and s_2' differently. On the contrary, conditional on s_1' being realized, signal s_2'' has no information added to infer the state because $\pi_2(s_2''|\omega) = \pi_1(s_1'|\omega)$ for both states. Hence, the current specification is *strategically equivalent* to a setting in which a sender observes the signal realizations of experiments chosen by previous senders.

3 Equilibrium Characterization

In this section, we first prove a result similar to the revelation principle that simplifies the analysis considerably. Without loss of generality we may focus on equilibria where only the first sender discloses non-trivial information on the equilibrium path. Preferences of other senders are captured as incentive compatibility constraints in such equilibria. Then we construct an equilibrium and show that, generically, the game has a unique equilibrium outcome.

3.1 Simplifying the Problem

We begin with the observation that it is without loss of generality to consider certain simple strategies by the senders. 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 may exist many equilibrium information structures, but because players care only about distributions over $\Omega \times A$, everyone is indifferent across all outcome-equivalent equilibria. These equilibria can be Blackwell comparable, but since players do not care directly about informativeness, we consider them equivalent.

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 $\vee_{i=1}^{n}\sigma_i'(h_i') = \sigma_1'$.*

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 $h_i^* = \{\pi_1^*, ..., \pi_{i-1}^*\}$ be the equilibrium path history when it is sender i's turn to move. This equilibrium generates a joint experiment $\pi^* = \vee_{i=1}^{n}\pi_i^*$. To construct a one-step equilibrium, we let sender 1 play π^* and assume that *on the equilibrium path* players $i = 2, ..., n$ 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 equilibrium 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.

Finally, 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 following arbitrary histories of play. The same logic as on the equilibrium path generalizes to any *continuation equilibrium* following an arbitrary history of play, but the notation gets heavy.

The observation in 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 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 will be introduced through the process.

The equilibrium is constructed backwardly. We begin with the decision maker's problem. As in standard persuasion models, what matters for the decision maker is his posterior belief about the state. Also, we observe that it is without loss to focus on the vertices in the belief space, making players' problems essentially finite. Then we express each sender's best response as a mean-preserving spread of interim beliefs induced by previous senders' experiments. Moreover, we recursively define stable beliefs for which senders have no incentive to provide further information given the continuation play; thus it is sufficient to pin down the set of stable beliefs to characterize a one-step equilibrium.

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

$$\mu(\omega|s) = \frac{p(s|\omega)\,\mu_0(\omega)}{\sum_{\omega' \in \Omega} p(s|\omega')\,\mu_0(\omega')}, \tag{3}$$

where have dropped the subscript of $p(s|\omega)$ defined in (2). 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 standard *Bayes-plausibility* constraint

$$\sum_{s \in S} \mu(\omega|s)\, p(s) = \mu_0(\omega). \tag{4}$$

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

$$H(a \succeq a') \equiv \{\mu \in \Delta(\Omega) | \sum_{\omega \in \Omega} \mu(\omega)[u_d(a,\omega) - u_d(a',\omega)] \geq 0\}, \tag{5}$$

defines the set of 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

$$M(a) = \cap_{a' \in A} H(a \succeq a'), \tag{6}$$

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

Interim Beliefs. A history $h_i = \{\pi_j\}_{j=1}^{i-1}$ induces a joint experiment $\pi^{i-1} = \vee_{j=1}^{i-1} \pi_j$. For each signal s of π^{i-1}, the corresponding belief $\mu(\omega|s)$ is given by (3). This is the decision maker's posterior belief if senders $i+1, ..., n$ do not add any information in the continuation game and s

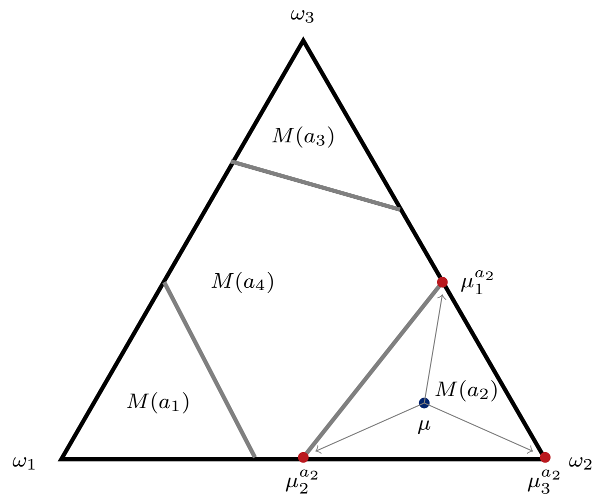

Figure 2: $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and $A = \{a_1, a_2, a_3, a_4\}$.

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 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 refinement 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 than 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 is on the form

$$\begin{aligned} V_n(\mu) &= \max_{\tau \in \Delta(\Delta(\Omega))} \sum_{\mu'} v_n(\sigma_d(\mu'), \mu')\tau(\mu') \quad (7) \\ \text{s.t.} \sum_{\mu'} \mu'\tau(\mu') &= \mu, \end{aligned}$$

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

In general, the support of τ_n could be infinite. The following observation makes it without loss to focus on τ_n with finite support. By construction, $M(a)$ is a finite convex polytope for each $a \in A$. Such a convex polytope has a finite set of $J(a)$ vertices $\left\{\mu_j^a\right\}_{j=1}^{J(a)}$ and these vertices span $M(a)$ so that every $\mu \in M(a)$ can be represented as a convex combination of the vectors $\left\{\mu_j^a\right\}_{j=1}^{J(a)}$.[3] Denote

$$X = \cup_{a\in A} \left\{\mu_j^a\right\}_{j=1}^{J(a)}, \tag{8}$$

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

Lemma 1. *Program (7) has a solution $\tau_n \in \Delta(X)$.*

The idea of Lemma 1 is that each $M(a)$ is spanned 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 τ_n to $\Delta(X)$ generates a utility at least as great as (7). But $\Delta(X)$ is a subset of the feasible set in (7), so the two problems must have the same value. Figure 2 provides an illustration. Suppose that μ is induced with probability τ by a solution to (7), then there is another solution in which $\{\mu_j^{a_2}\}_{j=1,2,3}$ is induced with probability τ_1, τ_2, τ_3 such that $\sum_j \mu_j^{a_2}\tau_j = \mu$ and $\sum_j \tau_j = \tau$.

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.,

$$X_n \equiv \{\mu \in X : v_n(\sigma_d(\mu), \mu) = V_n(\mu)\}. \tag{9}$$

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

$$X_i \equiv \{\mu \in X_{i+1} : v_i(\sigma_d(\mu), \mu) = \tilde{V}_i(\mu)\}, \tag{10}$$

[3] See Grünbaum, Klee, and Ziegler (1967).

with an auxiliary problem

$$\begin{aligned}\widetilde{V}_i(\mu) &= \max_{\tau\in\Delta(X_{i+1})} \sum v_i(\sigma_d(\mu'),\mu')\tau(\mu'|\mu) \\ \text{s.t.} & \sum_{\mu'\in X_{i+1}} \mu'\tau(\mu'|\mu) = \mu.\end{aligned} \tag{11}$$

Notice that (i) the solution to the auxiliary program (11) exists, (ii) $X_i \subseteq X_{i+1}$, and (iii) $X_1 \neq \emptyset$. In the auxiliary problem (11), sender i is restricted to use experiments that only induce vertex beliefs in X_{i+1}, and he believes that senders $i+1,...,n$ will not add any information. Because $X_i \subseteq X_j, \forall j > i$, sender i's belief is indeed justified.

Definition 3. *A belief is* ***stable*** *if $\mu \in X_i$ which is recursively defined by (9) and (10) for $i = 1,...,n$.*

By construction, no sender has the 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 non-trivial signal; if $\mu \notin X_1$, only sender 1 posts an informative experiment and the other senders remain silent. Off the equilibrium path, if one of sender i's interim beliefs $\mu_{i-1} \notin X_i$, he posts an experiment that "splits" the beliefs only in X_i and the subsequent senders do not add further information. The rest of the construction is relegated to the appendix. It is evident that no player has an incentive to deviate; thus we establish the existence of a one-step equilibrium. Formally,

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 $i = 1, 2..., n$, for every experiment profile $\pi_1, ..., \pi_{i-1}$ and possible signal profile $(s_1, ..., s_{i-1})$ 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,

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 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, which we show is without loss of generality. The basic idea is the same as for Lemma 1, but the proof is notationally more cumbersome, as we need to replicate the " incentives" corresponding to an arbitrary equilibrium with continuation play that is on the vertices only.

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.*

By Lemma 2, it is without loss to consider subgame perfect equilibria in which, after any history of play, continuation strategies are one step strategies inducing beliefs with support on vertices in the finite set X. In the same spirit, we note that when checking for subgame perfection it is without loss to consider one step deviations onto the set of vertices X. Hence, after any history, the continuation equilibrium outcome is determined by the one-step equilibrium of the corresponding subgame, which is immune to deviations that induce vertex beliefs. If there are multiple continuation equilibria that are not outcome equivalent, there must be a belief such that some sender is indifferent between staying at this belief and some mean-preserving spread over X.

Notice that there are two cases in which a sender is indifferent to splitting a belief into 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 $A \times \Omega$ is unchanged. Any failure of essential uniqueness must therefore correspond to indifferences over mean-preserving spreads that induce distinct actions. However, this case requires non-generic 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 a failure of essential uniqueness can occur for only a measure zero set of preferences.

4 Applications

This section discusses some applications of the equilibrium characterization. The aim is to shed light on some issues relevant for the design of the 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 are mainly interested in 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. Clearly, 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 therefore use set-wise 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)**.** *π is **essentially less informative** than π' if the finest experiment that is outcome equivalent to π is less informative than the finest experiment that is outcome equivalent to π' in the Blackwell order.*

Unlike some partial orderings considered in the literature on information design (e.g. individual sufficiency in Bergemann and Morris (2016)), the essential Blackwell ordering depends on optimal actions for the decision maker, so it is not an order applying for all preferences. Notice that the finest experiment among a class of outcome-equivalent experiments must put probability one on being on the vertices defining the optimality areas for the decision maker, allowing us to use stable beliefs to analyze equilibria. In the rest of this section, we adopt the essential Blackwell order to study how changes in the extensive form matters for outcomes.

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, Gentzkow and Kamenica (2017a,b) provide sufficient conditions under which additional senders do not reduce the amount of information revealed in simultaneous move Bayesian persuasion games.

The support for competition being beneficial becomes weaker in our sequential framework. Using a simple numerical example, Li and Norman (2018) show that adding senders to sequential Bayesian persuasion games may decrease the information revealed.

To start, the implications of adding additional senders depend on the *ordering* of the players as shown in the following simple example. Let the state space be $\{\omega_0, \omega_1\}$, the available actions be $\{a_0, a_1, a_2\}$, and assume that $M(a_i) = [i/3, (i+1)/3]$ for $i = 0, 1, 2$. Assume that the senders have state independent preferences with sender 1 having the strict preference ordering $a_1 \succ_1 a_2 \succ_1 a_0$ and sender 2 ranking the actions in accordance with $a_2 \succ_2 a_0 \succ_2 a_1$. Let the prior probability that the state is ω_0 be some $\mu_0 > 1/3$. In the single-sender persuasion game with the players being sender 1 and the decision maker, sender 1 may without loss optimally choose a mean preserving spread with support on $\{1/3, 2/3\}$ if $\mu_0 \in [1/3, 2/3]$. Notice that this is a simple example of outcome irrelevant multiplicity of Blackwell rankable information systems, as any mean-preserving spread in $[1/3, 2/3]$ is an equilibrium. When $\mu_0 > 2/3$ the only equilibrium is the unique mean-preserving spread with support $\{2/3, 1\}$.

Next, we add sender 2 to the game. First, suppose that sender 2 moves *before* sender 1. We illustrate the equilibrium analysis in panels (a) and (c) of Figure 3. Any interim belief μ such that $0 < \mu < 1/3$ will be split by sender 1 onto $\{0, 1/3\}$ which implies that actions a_0 and a_1 are both chosen with some strictly positive probabilities. Any belief on $[1/3, 2/3]$ ultimately leads to action a_1 for sure and a belief on $(2/3, 1)$ implies that action a_1 is taken with some probability and a_2 with some probability. It follows that the unique best response by sender 2 is to play the unique mean-preserving spread onto $\{0, 1\}$, so the state is fully revealed. As the payoffs are generic and do not satisfy the condition in Proposition 10 this illustrates that we cannot obtain a simple necessary and sufficient condition for full revelation.

The case where sender 1 moves before sender 2 is illustrated in panels (b) and (d) of Figure 3. We note that sender 2 will split any $0 < \mu < 1/3$ onto $\{0, 2/3\}$, implying that a best response for 1 is to split the prior to $\{0, 2/3\}$ if $\mu_0 < 2/3$ and to $\{2/3, 1\}$ if $\mu_0 \geq 2/3$. Hence, we see

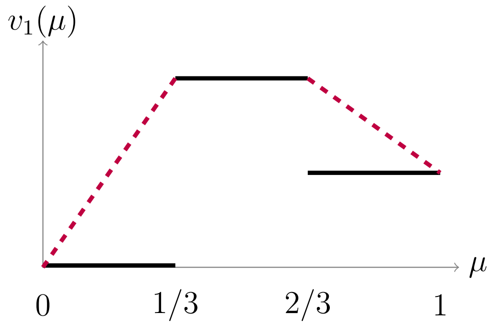

(a) Sender 1's payoff when sender 2 moves first

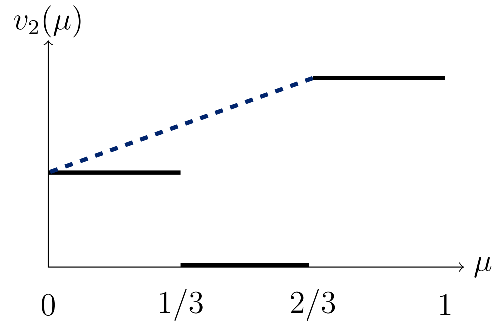

(b) Sender 2's payoff when sender 1 moves first

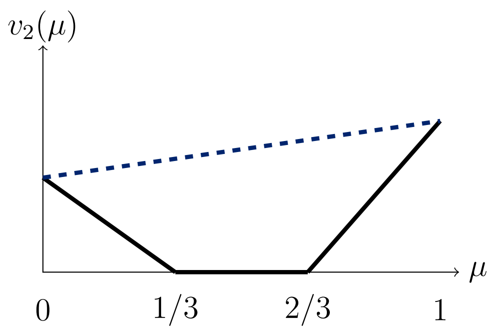

(c) Sender 2's payoff when he moves first

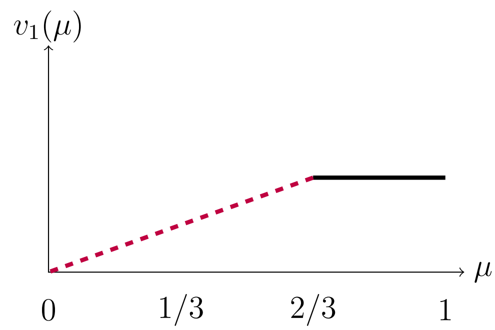

(d) Sender 1's payoff when he moves first

Figure 3: The effect of the ordering of moves. The solid line represents the sender's continuation payoff given subsequent players' best responses, while the dashed line represents his "concavified" continuation payoff.

that the informativeness and equilibrium payoffs depend on the order of moves, and also that the equilibrium in the model with sender 1 moving before sender 2 is strictly less informative in the Blackwell order than when sender 1 is the only sender, provided that $\mu_0 \in [2/3, 1)$.

While special, the example above delivers a general negative message. It is hopeless to search for general results unless we make restrictions on when a new sender moves. However, in the example, 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 persuasion 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 persuasion game with n senders and 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 $1, ..., n$ and the set of stable beliefs is defined backwardly, we have that

$$X_1^{n+1} \subseteq X_1^n. \tag{12}$$

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 and $X_1^{n+1}(\mu_0)$ be the support of the equilibrium in the game with $n+1$ senders. Since we focus on the finest experiment, these beliefs must be vertices and stable, i.e., $X_1^j(\mu_0) \subseteq X_1^j$ for $j = n, n+1$.

For the sake of 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 to (12).

□

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 reason 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 reduces the set of stable beliefs. As a result, such a change cannot make the persuasion outcome essentially less informative. In the special case where there are only two states, the incompleteness of essential Blackwell order is no longer substantial, so we have 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. When there are only two states, the support of the finest equilibrium contains at most two stable beliefs for generic preferences; otherwise, we can construct multiple mean-preserving spreads of the prior onto these stable beliefs, implying multiple equilibrium outcomes, a contradiction. Let $X_1^j(\mu_0) = \{\mu_L^j, \mu_H^j\}$ be the support of the finest equilibrium in the game with j senders where $j = n, n+1$. Without loss, assume $\mu_L^j \leq \mu_0 \leq \mu_H^j$ for every j. Then by Proposition 4, for each $\mu \in X_1^{n+1}$, either $\mu \in X_1^n$, or μ is not in the convex hull of $X_1^n(\mu_0)$. Therefore, $\mu_L^{n+1} \leq \mu_L^n \leq \mu_H^n \leq \mu_H^{n+1}$. The mean-preserving spread that splits μ_0 onto $X_1^{n+1}(\mu_0)$ can be obtained in two steps. First, splits μ_0 onto $\{\mu_L^n, \mu_H^n\}$. Second, further split μ_L^n, μ_H^n onto $X_1^{n+1}(\mu_0)$ accordingly. Therefore, the result follows. □

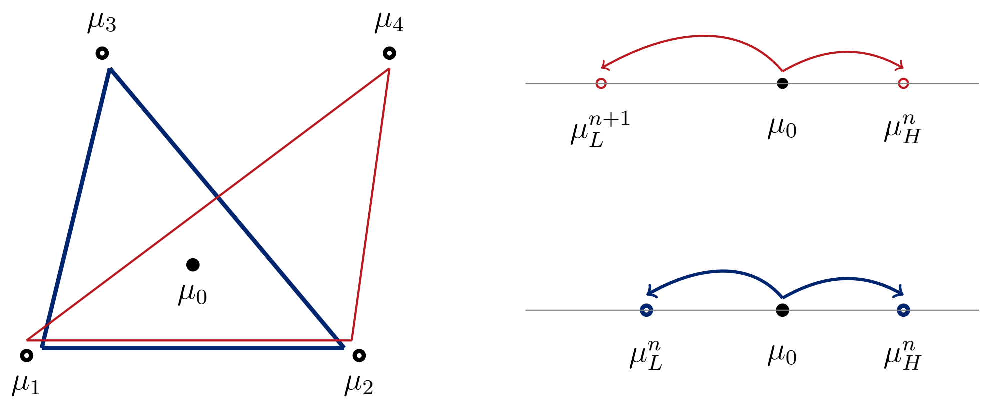

Figure 4: The left panel illustrates a case with $|\Omega| = 3$ while the right panel corresponds to a case with $|\Omega| = 2$.

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 $X_1^n(\mu_0) = \{\mu_1, \mu_2, \mu_3\}$ in the original n-sender game. When a new sender is added to speak before other senders, the support of the finest equilibrium becomes $X_1^{n+1}(\mu_0) = \{\mu_1, \mu_2, \mu_4\}$. Proposition 4 leaves the possibility that two equilibria are incomparable in the sense of Blackwell. On 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.

We would like to make a final remark on the effect of adding senders. The example in Li and Norman (2018) implies that, for some preference profile of senders, adding a new sender who moves after some other senders strictly reduces information revelation. Combining this observation and an implication of Proposition 4, we summarize the relation between the position of the newly added sender and the amount of information revelation as follows.

Corollary 1. *Fix a game with senders $1, 2, ..., n$, denoted by Γ. Add a new sender j^* and construct another game Γ^* with senders $1, 2, ..., n$ and j^* such that sender i moves after sender $i-1$ for every $i = 2, ..., n$. For generic preferences, there is no equilibrium of Γ^* essentially less informative than that of Γ if and only if the newly added sender j^* moves before every sender $i = 1, 2, ..., n$.*

Corollary 1 implies that, for an arbitrary preference profile, the information revelation will not be reduced if and only if the newly added sender speaks first. When the state space is binary, the result becomes even stronger: regardless of the preference profile of players, adding a new sender who speaks first will either increase or has no effect on information revelation; thus it will never lower the payoff of the decision maker.

4.3 Multiple Moves by the Same Sender

Our second application regards the communication protocol for a given set of senders. Up to this point we have allowed each player to move only once, which is without loss for results having to do with the characterization, existence, and uniqueness of equilibria because we can always add multiple players with identical preferences. In contrast, it is not obvious whether it is useful to allow multiple counterarguments, and this section addresses this question.

Senders who speak at late stages can make counterarguments to early movers' arguments, i.e. disclosing information conditional on the signals sent by previous senders. Is there any value in letting senders respond to counterarguments from other senders? If so, what is the source of the value? The literature has offered some discussion on this issue when senders have restrictions on their ability to provide information in each round (Glazer and Rubinstein 2001).

The Bayesian persuasion framework we adopt offers a frictionless alternative, which helps to identify the conditions needed to rationalize multi-rounds of rebuttals and counter-rebuttals. In our model, the preferences of senders are common knowledge, and a sender can provide as much information as he wants in a single round of disclosure. Hence, the only constraint on communication is strategic considerations. Our results imply that such a strategic friction is *per se* insufficient to justify multiple rounds of argument and counter-argument, except that an opportunity to offer a rebuttal may be useful for the first sender that moves.

Formally, we let $i \in \{1, ..., n\}$ denote the set of senders and we let the stage when senders move be denoted by $t = 1, ..., T$ 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 sequential persuasion game where each sender only moves once. For example, consider a sequential persuasion game with three senders $i = 1, 2, 3$ and five stages. Exactly one sender moves at each stage, and the order of moves is $1 \to 2 \to 3 \to 3 \to 2$. 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 $1 \to 3 \to 2$. The intuition is very simple. Consider the incentive of a sender who can speak at stage 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 in a Bayesian persuasion model 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 $1, ..., n$ 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 $1 \to 2 \to 3$, so that each player moves only once. Change the game to $2 \to 1 \to 2 \to 3$, 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 nevertheless generate different equilibrium outcomes because the first mover in the game can choose a Bayes plausible distribution of stable beliefs.

If the prior belief is stable, this choice doesn't matter, as any first mover is happy to not provide any information. If there are only two states it is also irrelevant. 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 disadvantegeous, 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 in Appendix B.2.

4.4 Simultaneous vs Sequential Persuasion

Now we fix the set of senders and the order of consultation. When the decision maker receives disclosures from senders sequentially, he 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. We conclude that information revealed in the simultaneous game cannot be essentially less informative than in the sequential game.

Consider a simultaneous-move persuasion game. Suppose that $\tau \in \Delta(\Delta(\Omega))$ is an equilibrium distribution of beliefs. 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 can use the same reasoning as in the basic sequential setup and 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 and the sequential games 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 between the models is 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.[4]

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.*

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

[4] A 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).

support of the equilibrium of 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.

□

Note that the comparison in Proposition 8 holds for any equilibrium in the simultaneous-move game including the ones in surviving iterated elimination of weakly dominated strategies, so it is not subject to the criticism in Hu and Sobel (2019).

Just like in the case of adding senders, a weakness of the result is the incompleteness of Blackwell's ordering, which implies that experiments may be non-comparable. However, we can again obtain a sharp characterization for the case with two states. Note, however, that the qualifier "pure strategy" is important. Li and Norman (2018) construct a numerical example where a mixed strategy equilibrium in the simultaneous game with two states is strictly less informative than the unique equilibrium in the sequential game.

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 there exist non-Blackwell comparable distributions also 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.

The last sender can always deviate to an equilibrium of the simultaneous-move game when the equilibrium experiments are comparable (as they are in the case of two states). Therefore 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 On Fully-Revealing Equilibria

A short-cut 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 non-degenerate 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 non-degenerate $\mu \in X$, there exists a sender i such that*

$$v_i(\sigma_d(\mu), \mu) < \sum_{\omega \in \Omega} u_i(\sigma_d(\delta_\omega, \omega))\mu(\omega), \tag{13}$$

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 (13) 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 way so that the corresponding sequential persuasion game does not have non-degenerate stable beliefs. To do so, it must be the case that every particular non-degenerate vertex belief is "disliked" by at least one sender.

It is worth mentioning that the full revelation sufficient condition (13) can be applied regardless of the extensive form of the game. As discussed in Sobel (2010), 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 can show that subgame perfect equilibria exist and generate a unique joint distribution over states and outcomes for generic preferences. The fact that a finite set of stable beliefs characterizes the equilibrium makes it convenient to identify the unique equilibrium outcome and to apply the model to study the optimal structure of consultations. In particular, (1) 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, (2) 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 (3) sequential persuasion cannot generate a more informative equilibrium than simultaneous persuasion, and is less informative in the case of two states.

A Appendix: Omitted Proofs

A.1 Proofs: One-Step Equilibrium

Proof of Proposition 1. To proceed, we extend the definition of one-step equilibrium to off-the-path of play:

Definition 5. *Consider strategy σ' and let h_i be an arbitrary history when sender $i \in \{1, ..., n-1\}$ 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.*[5] *We say that $\sigma'|_{h_i}$ is* ***one step*** *if $\vee_{j=i}^{n} \sigma'_j(h'_j|h_i) = \sigma'_i(h_i)$.*

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

[5]That is, $h'_i|h_i = h_i$, $h'_{i+1}|h_i = (h_i, \sigma'_i(h_i))$, $h'_{i+2}|h_i = \left(h_i, \sigma'_i(h_i), \sigma'_{i+1}(h_i, \sigma'_i(h_i))\right)$ and so on.

following h_i. Let

$$\pi^*|_{h_i} = \left(\vee_{i=i}^{i-1}\pi_i\right) \vee \left(\vee_{i=i}^{n}\pi_i^*|_{h_i}\right) \tag{A.1}$$

be the joint experiment generated by the continuation equilibrium path. Replace the continuation equilibrium strategies following h_i by $(\sigma_i', ..., \sigma_n', \sigma_d')$ where on the continuation outcome path

$$\begin{aligned} \sigma_i'(h_i) &= \pi^*|_{h_i} \\ \sigma_j'(h_i, \pi^*|_{h_i}, ..., \pi^*|_{h_i}) &= \pi^*|_{h_i} \text{ for } j \in \{i+1, .., n\} \\ \sigma_d'(h_i, \pi^*|_{h_i}, ..., \pi^*|_{h_i}, s) &= \sigma_d\left(h_i, \left(\pi_i^*|_{h_i}, ..., \pi_n^*|_{h_i}\right), s\right), \end{aligned} \tag{A.2}$$

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

$$\begin{aligned} \sigma_k'(h_i, \pi^*|_{h_i}, ..., \pi^*|_{h_i}, \pi_j, ..., \pi_k) &= \sigma_k^*\left(h_i, \pi_i^*|_{h_i}, ..., \pi_j^*|_{h_i}, \pi_j, ..., \pi_k\right) \\ \sigma_d'(h_i, \pi^*|_{h_i}, ..., \pi^*|_{h_i}, \pi_j, ..., \pi_n) &= \sigma_d^*\left(h_i, \pi_i^*|_{h_i}, ..., \pi_j^*|_{h_i}, \pi_j, ..., \pi_n\right), \end{aligned} \tag{A.3}$$

and for any other history; let

$$\begin{aligned} \sigma_j'(h_i, \pi_i, .., \pi_{j-1}) &= \sigma_j^*(h_i, \pi_i, .., \pi_{j-1}) \text{ for } j \in \{i+2, .., n\} \\ \sigma_d'(h_i, \pi_i, .., \pi_{j-1}, s) &= \sigma_d^*(h_i, \pi_i, .., \pi_{j-1}, s) \end{aligned} \tag{A.4}$$

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.2) that the implied distribution over $\Omega \times A$ is identical if each $j \geq i$ plays in accordance with the original equilibrium σ^*. Also, the strategies in (A.4) 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.4) is trivially subgame perfect. Finally, (A.3) 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 $\left(h_i, \pi_i^*|_{h_i}, ..., \pi_{j-1}^*|_{h_i}\right)$ in the original equilibrium, so $j \in \{i+1, ..., n\}$ 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.2), (A.3) and (A.4) following any history i and h_i we obtain a subgame perfect strategy profile which is one step after every history h with the same equilibrium outcome. □

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 $s' \in \pi^*|_{h_i}$ that induces a decision maker posterior belief $\mu' \notin X$ with positive probability. Let $a' = \sigma_d(h_i, \pi^*|_{h_i}, ..., \pi^*|_{h_i})$ be the equilibrium action induced by s'. Furthermore, let $M(a')$ be the belief polytope where a' is optimal and $X(a') = \{\mu_j(a')\}_{j=1}^m$ 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 $\{s_1, ..., s_m\}$, where each s_j generates posterior $\mu_j(a')$ and has unconditional probability $p(s')\lambda_j$, and everything else in π' is like the original equilibrium.[6] We also assume that the decision maker follows a strategy in which

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

where σ^*_d is the strategy of the decision maker in the original equilibrium. Since each $\mu_j(a') \in M(a')$ this must be 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 $j = \{i, .., n\}$ play

$$\sigma'_j(h_j) = \begin{cases} \pi' & \text{if } h_j = (h_i, \pi', ..., \pi') \\ \sigma^*_i(h_i, \pi^*, .., \pi^*, \pi_k, ..., \pi_{j-1}) & \text{if } h_j = (h_i, \pi', .., \pi', \pi_k, ..., \pi_{j-1}) \\ \sigma^*_i(h_j) & \text{if } h_j = (h_i, \pi_i, ..., \pi_{j-1}) \text{ is such that } \pi_i \neq \pi' \end{cases}, \quad \text{(A.5)}$$

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

$$v_i(a,\mu) = \sum_{j=1}^m \lambda_j v_n(a, \mu_j^a) = \sum_{j=1}^m \lambda_j v_n(\sigma_d(\pi', .., \pi', s_j), \mu_j^a), \quad \text{(A.6)}$$

while nothing is changed for signal realizations that are kept like 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

[6] It is possible that $\lambda_j = 0$ for some j. Instead of eliminating these beliefs we may simply generate a probability zero signal in order not to treat this case separately.

off-path continuation play in the original equilibrium, so there is no profitable deviation off the path. Repeating the same argument for each history h_i, every continuation experiment $\pi^*|_{h_i}$ and every realization $s' \in \pi^*|h_i$ with corresponding belief $\mu' \notin X$ completes the proof.

□

Proof of Corollary 2. It is immediate (see Lemma 1 below for details) that n has no profitable deviation following any history unless there is a profitable deviation onto X, so consider sender $i < n$ and history h_i with one-step continuation experiment $\pi^*|_{h_i}$. Consider a deviation π_i with some induced belief $\mu' \notin X$ given all subsequent senders play sequentially rational, which is profitable for i. By Proposition 2 (applied to truncated game with i moving first and using one of the beliefs in the interim belief distribution as the prior) there is a continuation one-step equilibrium after (h_i, π_i) that generates the same distribution over $\Omega \times A$. Hence, there is a deviation over vertex beliefs and a one-step continuation equilibrium that is profitable.

□

A.2 Proofs: Equilibrium Construction

Proof of Lemma 1. For each program on form (7), we consider a restricted *finite* linear program

$$\begin{aligned} \widetilde{V}_n(\mu) &= \max_{\tau \in \Delta(X)} \sum_{\mu' \in X} v_n(\sigma_d(\mu'), \mu')\tau(\mu') \qquad \text{(A.7)} \\ \text{s.t.} \sum_{\mu'} \mu' \tau(\mu') &= \mu, \end{aligned}$$

where X is defined in (8). Hence, (A.7) is well defined as it is a finite dimensional bounded linear program.

Pick any feasible solution τ to program (7). For each $a \in A$ write $\tau^a(\mu')$ for μ' such that $\sigma_d(\mu') = a$ and $\tau = \left\{\{\tau^a(\mu')\}_{\mu' \in \widehat{W}(a)}\right\}_{a \in A}$ where $\widehat{M}(a) = \{\mu \in \Omega | \sigma_d(\mu) = a\}$ is the "decision area" of action a defined by $\sigma_d(\cdot)$. Obviously, $\widehat{M}(a) \subset M(a), \forall a$.

For each $\mu' \in M(a)$ there exists $\lambda' \in \Delta\left(\{\mu_j^a\}_{j=1}^{J(\alpha)}\right)$ such that $\mu' = \sum_{j=1}^{J(a)} \lambda'_j \mu_j^a$. For every $a \in A$ and μ_j^a spanning $M(a)$ let $\widehat{\tau}\left(\mu_j^a\right) = \sum_{\mu' \in \widehat{M}(a)} \tau(\mu') \lambda'_j$ so that

$$\sum_{j=1}^{J(a)} \widehat{\tau}\left(\mu_j^a\right) = \sum_{\mu' \in \widehat{M}(a)} \tau(\mu') \sum_{j=1}^{J(a)} \lambda'_j = \sum_{\mu' \in \widehat{M}(a)} \tau(\mu'). \qquad \text{(A.8)}$$

Since it is possible that $v_n\left(a, \mu_j^a\right) < v_n\left(a', \mu_j^a\right)$ 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.7) satisfies

$$\begin{aligned}\widetilde{V}_n(\mu) &\geq \sum_{a\in A}\sum_{j=1}^{J(a)} v_n(a,\mu_j^a)\widehat{\tau}\left(\mu_j^a\right) = \sum_{a\in A}\sum_{j=1}^{J(a)}\sum_{\omega\in\Omega} u_n(a,\omega)\mu_j^a(\omega)\,\widehat{\tau}\left(\mu_j^a\right) \qquad \text{(A.9)}\\ &= \sum_{a\in A}\sum_{\omega\in\Omega} u_n(a,\omega)\sum_{j=1}^{J(a)}\left[\mu_j^a(\omega)\,\lambda_j'\right]\left[\sum_{\mu'\in\widehat{M}(a)}\tau(\mu')\right] = \sum_{a\in A}\sum_{\omega\in\Omega} u_n(a,\omega)\mu'\left[\sum_{\mu'\in\widehat{M}(a)}\tau(\mu')\right]\\ &= \sum_{\mu'} v_n(\sigma_d(\mu'),\mu')\tau(\mu').\end{aligned}$$

This holds for any feasible solution to (7). Hence, $\widetilde{V}_n(\mu) \geq V_n(\mu)$. Moreover, any optimal solution to (A.7) is a feasible solution to (7), so $\widetilde{V}_n(\mu) \leq V_n(\mu)$. This establishes that solutions to (7) exist and that $\widetilde{V}_n(\mu) = V_n(\mu)$ and that every $\widetilde{\tau}_n \in \Delta(X)$ that solves (A.7) also solves (7). Finally, if τ_n solves (7) and μ' is such that $\tau_n(\mu') > 0$, 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.9). □

Proof of Proposition 2. In what follows, we constructs a subgame perfect equilibrium where sender i's equilibrium strategy coincides with the solution to program (11). 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 only.

Fix a pair (σ_d, τ_n) that solves the decision maker and sender n such that

- σ_d breaks the tie of the decision maker's problem in favor of sender n, and
- sender n induces vertex beliefs only, i.e., $\tau_n \in \Delta(X_n)$, which is without loss by Lemma 1.

Then, sender $n-1$'s problem can be formulated as follows

$$\begin{aligned} V_{n-1}(\mu) \quad &= \quad \max_{\tau}\left[\sum_{\mu'\in\Delta(\Omega)}\left(\sum_{\mu''\in\Delta(X_n)} v_{n-1}(\sigma_d(\mu''),\mu'')\tau_n(\mu''|\mu')\right)\tau(\mu'|\mu)\right] \qquad \text{(A.10)}\\ &\text{s.t.} \quad \sum_{\mu'\in\Delta(\Omega)}\mu'\tau(\mu'|\mu) = \mu. \end{aligned}$$

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

In general, there may exist multiple $\tau_n \in \Delta(X_n)$ solving program (7), and whether program (A.10) is well-defined depends on the selection of τ_n. Whenever there is multiplicity, we select $\tau_n(\mu|\mu) = 1$ for $\mu \in X_n$. That is, the last sender adds no information whenever he has no strict incentive to do so. As will be clear soon, such a selection ensures sender $n-1$'s problem is well-defined. As in the proof of Lemma 1, our strategy is to bound the value function of program (A.10) and derive a feasible mean-preserving spread in $\Delta(X_n)$ that achieves the upper bound.

Fix a feasible strategy τ_{n-1} in program (A.10). For each μ, construct a *compound mean-preserving spread* $\tilde{\tau}_{n-1} : \Delta(\Omega) \to \Delta(X_n)$ using τ_{n-1} and τ_n such that

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

Since $\tilde{\tau}_{n-1}$ is feasible in program (11) for $i = n-1$, we have

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

for every μ. On the other hand, in program (A.10), it is feasible to choose any mean-preserving spread $\tau \in \Delta(X_n)$. By the selection rule imposing on τ_n, sender n will not add information when $\mu \in X_n$, so

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

for every μ. Notice that this inequality crucially relies on the selection rule of τ_n. If sender n further provides an informative signal when at some interim belief $\mu \in X_i$, some feasible mean-preserving spreads in program (11) are no longer feasible in program (A.10).

Consequently, $V_{n-1}(\cdot) = \tilde{V}_{n-1}(\cdot)$. Since $\tilde{V}_{n-1}(\mu)$ is well-defined, we establish the optimal mean-preserving spread of sender $n-1$, denoted by τ_{n-1}. The support of τ_{n-1} is X_{n-1}. Again, whenever there exist multiple τ_{n-1}, we select ones that sender $n-1$ adds no information at every $\mu \in X_{n-1}$, ensuring that the best responses of senders $1, ..., n-2$ well-define. Other senders' optimal mean-preserving spread can be constructed by induction. □

A.3 Proof: 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.3.1 Preliminaries

The following Corollary is more or less a direct consequence of Proposition 2.

Corollary 2. *Fix an equilibrium strategy profile and a history. For any deviation of sender i, there is a payoff-equivalent one-step continuation play where senders $i+1, ..., n$ add no information and the resulting posterior beliefs are vertices.*

Proof. See Appendix A.1. □

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, such as which sender revealed state ω, can be used to construct non-Markov mixed-strategies for the decision maker.[7] 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.

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 generate the same posterior belief $\mu \in X$. Then the decision maker's equilibrium choice must be identical.*

Proof. Consider some 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

$$\sigma_d^*(h^*, s^*) = a \neq a' = \sigma_d^{**}(h^{**}, s^{**}). \tag{A.11}$$

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

$$v_d(a, \omega) = v_d(a', \omega). \tag{A.12}$$

A decision maker's payoff function may be viewed as element in $|\Omega \times A|$ dimensional Euclidean space and the payoff functions that satisfy (A.12) defines a $|\Omega \times A| - 1$ dimensional subspace. As there are a finite number of triples $(a, a', \omega) \in A^2 \times \Omega$, the set of utility functions in which (A.12) holds for some triple (a, a', ω) is of Lebesgue measure zero. Next, consider the case with (A.11) 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 one, implying that

$$\sum_{\omega \in \Omega} v_n(a, \omega) \mu_j(a) = \sum_{\omega \in \Omega} v_n(a', \omega) \mu_j^a, \tag{A.13}$$

[7] In a numerical example, we construct a non-Markov equilibrium where the decision maker's tie-breaking rule determines by payoff-irrelevant endogenous choice senders. See Appendix B.1 for detail.

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 $(a, a^{'}, \mu_j^a)$ to consider and for each triple (A.13) 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.

□

By Lemma 3, for generic preferences, it is without loss to restrict to a Markov strategy on form $\sigma_d : \Delta(\Omega) \to A$ for the decision maker. It is then useful to define $\widehat{v}_i : \Delta(\Omega) \to \mathbb{R}$, where

$$\hat{v}_i(\mu) \equiv v_i(\sigma_d(\mu), \mu), \tag{A.14}$$

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 $\sigma_d : \Delta(\Omega) \to A$. Then, for any sender $i \in \{1, ..., n\}$ 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:*

1. *$\sigma_d(\mu') = \sigma_d(\mu)$ for every $(\mu, \mu') \in Y$,*
2. *there exists $(\mu, \mu') \in Y$ such that $\sigma_d(\mu) \neq \sigma_d(\mu')$. In this case*

$$\hat{v}_i(\mu) > \sum_{\mu' \in Y} \hat{v}_i(\mu') \tau(\mu'), \tag{A.15}$$

for a set of sender i' 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 exists $\mu \in X_i$ and $Y \subset X_i$ and $\tau \in \Delta(Y)$ such that $\mu = \sum_{\mu' \in Y} \mu' \tau(\mu')$ and that (A.15) is violated for sender i. Denote by $\{\mu_1, ..., \mu_{m+1}\} = Y$ and $\tau = (\tau_1, ..., \tau_{m+1})$ and write the failure of (A.15) as

$$\hat{v}_i(\mu) = \sum_{\boldsymbol{j}=1}^{m+1} \hat{v}_i(\mu_{\boldsymbol{j}}) \tau_{\boldsymbol{j}}. \tag{A.16}$$

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 $\{\mu_2 - \mu_1, ..., \mu_{m+1} - \mu_1\}$ are linearly dependent. Then there are scalars $(\alpha_2, ..., \alpha_{m+1}) \neq (0, ..., 0)$ such that $\sum_{\boldsymbol{j}=2}^{m+1} \alpha_{\boldsymbol{j}} (\mu_{\boldsymbol{j}} - \mu_1) = 0$. So

$$\left(-\sum_{j=2}^{m+1} \alpha_j\right) \mu_1 + \sum_{j=2}^{m+1} \alpha_j \mu_j = \sum_{\boldsymbol{j}=1}^{m+1} \alpha_j \mu_j = 0, \tag{A.17}$$

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

$$\mu = \sum_{j=1}^{m+1} \mu_j \tau_j = \sum_{j=1}^{m+1} \mu_j \tau_j - \beta \sum_{j=1}^{m+1} \alpha_j \mu_j = \sum_{j=1}^{m+1} (\tau_j - \beta \alpha_j) \mu_j. \tag{A.18}$$

Let $I^+ = (j \in \{1, ..., m+1\} \,|\tau_j > 0)$ and let j^* be chosen so that $0 < \frac{\tau_{j^*}}{\alpha_j^*} \leq \frac{\tau_j}{a_j}$ 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

$$\tau_j^* = \tau_j - \frac{\tau_{j^*}}{\alpha_j^*} \alpha_j. \tag{A.19}$$

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 $\{\mu_1, ..., \mu_{m+1}\}$ and still find a convex combination that generates μ. By induction, there exists an affinely independent set of vectors $\{\widehat{\mu}_1, ..., \widehat{\mu}_k\} \subseteq Y$ such that μ is in its convex hull, implying that there exists a unique solution $\widehat{\tau}$ such that $\mu = \sum_{j=1}^{k} \widehat{\mu}_j \widehat{\tau}_j$. [8] If $\sigma_d(\widehat{\mu}_j) = \sigma_d(\widehat{\mu}_{j'})$ for every pair of beliefs in $\{\widehat{\mu}_1, ..., \widehat{\mu}_k\}$, then μ and $\widehat{\tau}$ are outcome equivalent. If $\sigma_d(\widehat{\mu}_j) \neq \sigma_d(\widehat{\mu}_{j'})$ for some beliefs in $\{\widehat{\mu}_1, ..., \widehat{\mu}_k\}$

$$\hat{v}_i(\mu) = \sum_{j=1}^{k} \hat{v}_i(\widehat{\mu}_j) \widehat{\tau}_j, \tag{A.20}$$

then $\widehat{v}_i : \Delta(\Omega) \to R$ belongs to a Lebesgue measure zero set of utility functions.[9] 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 indifferences 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 on form

$$\mu = \sum_{l=1}^{L} \beta_l \sum_{j=1}^{k(j)} \widehat{\mu}_j(l) \tau_{\boldsymbol{j}}(l), \tag{A.21}$$

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

[9]By repeating the steps in (A.30), (A.31) and (A.32) below, the measure zero condition in belief space implies measure zero in terms of maps $u_i : A \times \Omega . \to \mathbb{R}$.

where $\beta_l \geq 0$ for each l, $\sum_{l=1}^{L} \beta_j = 1$ and every set $\{\widehat{\mu}_1(j), ..., \widehat{\mu}_k(j)\}$ is affinely independent. Hence, if (A.15) holds for every affinely independent subset of X_i it holds for all subsets of X_i. The result follows. □

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 non-trivial 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 to τ generates identical joint distribution over actions and states.

In the second case of Lemma 4, X_i, the set of beliefs of the truncated game with senders $i, i+1, ..., n$, 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.15) doesn't hold, implying that

$$\hat{v}_i(\mu) = \sum_{\mu' \in Y} \hat{v}_i(\mu') \tau(\mu'), \tag{A.22}$$

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.22) can only hold for a non-generic set of functions $\widehat{v}_i : \Delta(\Omega) \to \mathbb{R}$. [10] If, instead, Y is an affinely dependent set, then there must be an affinely independent subset of Y such that (A.22) holds for some mean-preserving spread with support on the affinely independent subset. For each affinely independent subset of Y, this requires non-generic preferences, and since there is a finite number of senders and affinely independent subsets, the result follows by induction.

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

Lemma 5. *Fix any $i \in \{1, ..., n\}$. Then*

$$\sum_{\mu' \in Y} \widehat{v}_i(\mu') \tau(\mu') \neq \sum_{\mu' \in \widetilde{Y}} \widehat{v}_i(\mu') \widetilde{\tau}(\mu'), \tag{A.23}$$

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

[10] This also implies that a non-generic set of utility functions $u_i : A \times \Omega \to \mathbb{R}$ can satisfy the equality

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

$$\mu_0 = \sum_{\mu \in X(\mu_0)} \mu \tau(\mu), \tag{A.24}$$

$$\widetilde{\mu}_0 = \sum_{\mu \in X(\mu_0)} \mu \lambda(\mu). \tag{A.25}$$

Hence, for any β

$$\mu_0 = \sum_{\mu \in X(\mu_0)} \mu \left(\tau(\mu) - \beta \lambda(\mu)\right) + \beta \widetilde{\mu}_0, \tag{A.26}$$

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

$$\widetilde{\mu}_0 = \sum_{\mu \in X(\widetilde{\mu}_0)} \mu \widetilde{\tau}(\mu). \tag{A.27}$$

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

$$\{\tau(\mu) - \beta \lambda(\mu) + \beta \widetilde{\tau}(\mu)\}_{\mu \in X(\mu_0) \cup X(\widetilde{\mu}_0)}, \tag{A.28}$$

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

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

Hence,

$$\sum_{\mu \in X(\widetilde{\mu}_0)} \hat{v}_1(\mu) \widetilde{\tau}(\mu) < \sum_{\mu \in X(\mu_0)} \hat{v}_1(\mu) \lambda(\mu), \tag{A.29}$$

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

A.3.2 Proof of Proposition 3

Lemma 2 and corollary 2 imply that for sender $i = 2, ..., n$ we only need 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 $\widetilde{Y} \subset X_i$ such that

$$\sum_{\mu' \in Y} \widehat{v}_i(\mu') \tau(\mu') = \sum_{\mu' \in \widetilde{Y}} \widehat{v}_i(\mu') \widetilde{\tau}(\mu'). \tag{A.30}$$

where τ is the unique mean-preserving spread of μ onto Y and $\widetilde{\tau}$ is the unique mean-preserving of μ onto $\widetilde{Y}$. 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.30) can be rewritten as

$$\sum_{\mu' \in Y} \sum_{\omega \in \Omega} [u_i(\sigma_d(\mu'), \omega) \mu'(\omega)] \tau(\mu') = \sum_{\mu' \in \widetilde{Y}} \sum_{\omega \in \Omega} [u_i(\sigma_d(\mu'), \omega) \mu'(\omega)] \widetilde{\tau}(\mu'). \tag{A.31}$$

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 $\widetilde{Y}(a)$ we may rewrite (A.31) further as

$$\sum_{a \in A} \left\{ \sum_{\omega \in \Omega} u_i(a, \omega) \left[\sum_{\mu' \in X(\mu, a)} \mu'(\omega) \tau(\mu') - \sum_{\mu' \in \widetilde{X}(\mu, a)} \mu'(\omega) \widetilde{\tau}(\mu') \right] \right\} = 0. \tag{A.32}$$

Since τ and $\widetilde{\tau}$ are uniquely defined, this defines a lower dimensional subspace of $|A \times \Omega|$ −dimensional Euclidean space, so the set of sender i payoff functions such that (A.30) holds is measure zero. Since X_i is finite, there is a finite set of pairs of affinely independent sets spanning μ and we only consider μ from the finite set $X \cup \{\mu_0\}$. The result follows.

A.4 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 (reference to recursive definition of stable beliefs). 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 at $X_i^t = X_i^{t'}$, 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 vertex 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' doesn't affect the set of stable vertex 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 $X_i^t \subseteq X_i^{t'}$. 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. □

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 the argument can be repeated for any μ not in $\Delta(X)$.

□

Proof of Proposition 9. By the proof of Proposition 5, the support of the finest equilibrium in a sequential game contains at most two stable beliefs, i.e., $\{\mu_L, \mu_H\}$ such that $\mu_L \leq \mu_0 \leq \mu_H$. Take a belief μ induced by an equilibrium of the simultaneous game. By Proposition 8, either $\mu \leq \mu_L$ or $\mu \geq \mu_H$. The rest of the proof follows exactly the same argument in the proof of Proposition 5, so it is omitted. □

B Examples

B.1 Non-Markov Equilibrium

In this section, we consider an example which 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

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

Also suppose that two senders have state-independent preferences

$$u_1(a,\omega) = \begin{cases} 3 & \text{if } a \in \{a_1, a_4\} \\ 1 & \text{if } a = a_3 \\ 0 & \text{if } a \in \{a_2, a_5\} \end{cases}, \text{ and, } u_2(a,\omega) = \begin{cases} 3 & \text{if } a \in \{a_2, a_5\} \\ 1 & \text{if } a = a_3 \\ 0 & \text{if } a \in \{a_1, a_4\} \end{cases}. \tag{B.2}$$

Consider a Markov equilibrium. Allowing for mixed strategies let $\sigma_1(0)$ be the probability for a_1 given belief $\mu = 0$ and $\sigma_4(1)$ 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 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 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 will just consider two of then, 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 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 and expected utility of 2/3 and player 2 and expected utility of $-1/3$. In contrast, when sender 2 moves first, μ_1 is played with probability 2/3 and e_3 with probability 1/3 resulting in the opposite expected utilities.

References

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

AU, P. H., AND K. KAWAI (2017): "Competitive disclosure of correlated information," *Economic Theory*, pp. 1–33.

——— (2020): "Competitive information disclosure by multiple senders," *Games and Economic Behavior*, 119, 56–78.

AUMANN, R., AND M. MASCHLER (1968): "Repeated games of incomplete information, the zero-sum extensive case," *Reports ST-143, Mathematica Inc., Princeton, NJ.*

BATTAGLINI, M. (2002): "Multiple referrals and multidimensional cheap talk," *Econometrica*, 70(4), 1379–1401.

BERGEMANN, D., AND S. MORRIS (2016): "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 11(2), 487–522.

BHATTACHARYA, S., AND A. MUKHERJEE (2013): "Strategic information revelation when experts compete to influence," *The RAND Journal of Economics*, 44(3), 522–544.

BLACKWELL, D. (1953): "Equivalent comparisons of experiments," *The annals of mathematical statistics*, 24(2), 265–272.

BOARD, S., AND J. LU (2018): "Competitive information disclosure in search markets," *Journal of Political Economy*, 126(5), 1965–2010.

BOLESLAVSKY, R., AND C. COTTON (2016): "Limited capacity in project selection: competition through evidence production," *Economic Theory*, pp. 1–37.

ELY, J., A. FRANKEL, AND E. KAMENICA (2015): "Suspense and surprise," *Journal of Political Economy*, 123(1), 215–260.

Ely, J. C. (2017): "Beeps," *The American Economic Review*, 107(1), 31–53.

Gentzkow, M., and E. Kamenica (2017a): "Bayesian persuasion with multiple senders and rich signal spaces," *Games and Economic Behavior*, 104, 411–429.

——— (2017b): "Competition in persuasion," *The Review of Economic Studies*, 84(1), 300–322.

Glazer, J., and A. Rubinstein (2001): "Debates and decisions: On a rationale of argumentation rules," *Games and Economic Behavior*, 36(2), 158–173.

Green, J., and N. Stokey (1978): "Two Representations of Information Structures and their Comparisons," *Institute for Mathematical Studies in the Social Sciences, Stanford University, technical report no. 271.*

Grünbaum, Brankoand Kaibel, V., V. Klee, and G. M. Ziegler (1967): *Convex polytopes*, vol. 1967. Springer.

Harris, C. (1985): "Existence and characterization of perfect equilibrium in games of perfect information," *Econometrica: Journal of the Econometric Society*, pp. 613–628.

Hu, P., and J. Sobel (2019): "Simultaneous Versus Sequential Disclosure," *Working Paper, University of California San Diego.*

Hwang, I., K. Kim, and R. Boleslavsky (2019): "Competitive Advertising and Pricing," *Unpublished Manuscript.*

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

Kartik, N., F. X. Lee, and W. Suen (2016): "A Theorem on Bayesian Updating and Applications to Signaling Games," *Unpublished Manuscript.*

——— (2017): "Investment in concealable information by biased experts," *The RAND Journal of Economics*, 48(1), 24–43.

Kawai, K. (2015): "Sequential cheap talks," *Games and Economic Behavior*, 90, 128–133.

Krishna, V., and J. Morgan (2001): "A model of expertise," *Quarterly Journal of Economics*, 116(2).

Li, F., and P. Norman (2018): "On Bayesian persuasion with multiple senders," *Economics Letters*, 170, 66–70.

Lipnowski, E., and L. Mathevet (2017): "Simplifying Bayesian Persuasion," *Unpublished Manuscript.*

——— (2018): "Disclosure to a psychological audience," *American Economic Journal: Microeconomics*, 10(4), 67–93.

Milgrom, P., and J. Roberts (1986): "Relying on the information of interested parties," *The RAND Journal of Economics*, pp. 18–32.

Rayo, L., and I. Segal (2010): "Optimal information disclosure," *Journal of Political Economy*, 118(5), 949–987.

Sobel, J. (2010): "Giving and receiving advice," in *Econometric Society 10th World Congress.*

Wu, W. (2018): "Sequential Bayesian Persuasion," *Unpublished Manuscript.*