Private Bayesian Persuasion with Sequential Games

Andrea Celli,[1] Stefano Coniglio,[2] Nicola Gatti[1]
[1]Politecnico di Milano, Piazza Leonardo da Vinci 32, Milan, Italy
[2]University of Southampton, University Road SO17 1BJ, Southampton, United Kingdom
{andrea.celli, nicola.gatti}@polimi.it, s.coniglio@soton.ac.uk

Abstract

We study an information-structure design problem (a.k.a. a persuasion problem) with a *single* sender and *multiple* receivers with actions of *a priori* unknown types, independently drawn from action-specific marginal probability distributions. As in the standard *Bayesian persuasion* model, the sender has access to additional information regarding the action types, which she can exploit when committing to a (noisy) signaling scheme through which she sends a private signal to each receiver. The novelty of our model is in considering the much more expressive case in which the receivers interact in a *sequential game with imperfect information*, with utilities depending on the game outcome and the realized action types. After formalizing the notions of *ex ante* and *ex interim* persuasiveness (which differ by the time at which the receivers commit to following the sender's signaling scheme), we investigate the *continuous optimization* problem of computing a signaling scheme which maximizes the sender's expected revenue. We show that computing an optimal *ex ante* persuasive signaling scheme is NP-hard when there are three or more receivers. Instead, in contrast with previous hardness results for *ex interim* persuasion, we show that, for games with two receivers, an optimal *ex ante* persuasive signaling scheme can be computed in polynomial time thanks to the novel algorithm we propose, based on the ellipsoid method.

Bayesian persuasion, introduced by Kamenica and Gentzkow (2011), revolves around influencing the behavior of self-interested agents through the provision of payoff-relevant information. Differently from traditional *mechanism design*, where the designer influences the outcome by providing tangible incentives, in Bayesian persuasion the designer influences the outcome by deciding *who gets to know what* (Bergemann and Morris 2016b). Real-world applications are ubiquitous. For instance, this framework has been recently applied to security problems (Rabinovich et al. 2015; Xu et al. 2015; 2016), financial-sector stress testing (Goldstein and Leitner 2018), voting (Alonso and Câmara 2016; Castiglioni, Celli, and Gatti 2019), and online advertisement (Badanidiyuru, Bhawalkar, and Xu 2018; Emek et al. 2012).

The classical Bayesian persuasion framework comprises a *single sender* and a *single receiver*. The sender, who has access to some private information, designs a signaling scheme in order to persuade the receiver to select a favorable action. The model assumes that the sender commits to the selected signaling scheme. This hypothesis is realistic in many settings where reputation and credibility are a key factor for the long-term utility of the sender (Rayo and Segal 2010), as well as whenever an automated signaling scheme either has to abide by a contractual service agreement or it is enforced by a trusted authority (Dughmi 2017).

The extension to the case with multiple receivers is of major interest, see, *e.g.*, its applications to private-value auctions studied by Kaplan and Zamir (2000). In this setting, a number of works assume a *public signal* model in which all the receivers observe the same information (Dughmi, Immorlica, and Roth 2014; Alonso and Câmara 2016; Dughmi 2018). A more general setting is the *private signal* case, in which the sender may tailor receiver-specific signals. Persuasion with private signals has been explored only in very specific settings, such as two-agents two-action games (Taneva 2019), unanimity elections (Bardhi and Guo 2018), voting with binary action spaces and binary states of Nature (Wang 2013), and games with binary action spaces and no inter-agent externalities (Babichenko and Barman 2016; Arieli and Babichenko 2016; Dughmi and Xu 2017). As pointed out by Dughmi (2017), the problem of computing private signaling schemes in general multi-receiver settings still lacks a general algorithmic framework.

In this work, we make a further step in this direction and consider a general model with the following key features: i) it admits an arbitrary number of receivers' actions and states of Nature; ii) it allows inter-agent externalities; [1] iii) it models sequential interactions among receivers. The last point constitutes the major difference from classical Bayesian persuasion models, which typically assume that the receivers take their actions simultaneously (Dughmi 2017; Kamenica 2018) as, to the best of our knowledge, we address here the multi-receiver case with sequential interactions among receivers for the first time in the literature. As most of the real-world economic interactions take place sequentially, this allows for a greater modeling flexibility which could be exploited in the context of, *e.g.*, sequential auctions (Leme,

[1]When there are inter-agent externalities, the utility of a receiver is determined by the state of Nature, her action, and (crucially) the actions selected by all the other receivers.

Syrgkanis, and Tardos 2012). In this paper, we show how to address sequential, multi-receiver settings algorithmically via the notion of *ex ante* persuasive signaling scheme, where the receivers commit to following the sender's recommendations having observed only the signaling scheme. This is motivated by the fact that the classical notion of persuasiveness (*ex interim* persuasiveness), which allows the receivers to deviate after observing the sender's signal, renders most of the associated design problems (with the exception of very narrow settings) computationally intractable even when the interaction is simultaneous (Dughmi and Xu 2016), ultimately making its adoption impractical in real-world applications where the receivers act sequentially. In parallel with our work, Xu (2019) introduced a notion of *ex ante* persuasion similar to ours, but studies it in a more restrictive setting: public signaling, simultaneous moves, binary actions, and no inter-agent externalities.

Ex ante persuasive signaling schemes may be employed every time the environment allows for a credible receivers' commitment before the recommendations are revealed. As argued by Kamenica and Gentzkow (2011), this is *not* unrealistic. On a general level, the receivers will uphold their *ex ante* commitment every time they reason with a long-term horizon where a reputation for credibility positively affects their utility (Rayo and Segal 2010). In some cases, they could also be forced to stick to their *ex ante* commitment by contractual agreements or penalties. Many real-world problems involve *ex ante* commitments. This happens, for example, when the signaling schemes are implemented as software (*e.g.*, recommender systems) and receivers can decide whether to adopt it or not. This is the case in sequential auctions in online advertising, where a (trusted) third party service (*e.g.*, programmatic advertising platforms) could allow bidders for coordinated behaviors during the sequential auction, leading to better outcomes in terms of bidders' payoffs, and to more efficient allocations of the ads.

Original contributions. We investigate private persuasion games with multiple receivers interacting in a sequential game, and study the continuous optimization problem of computing a private signaling scheme which maximizes the sender's expected utility. We focus on the framework with independent action types, similarly to what is previously done by Dughmi and Xu (2016). We introduce the notion of *ex ante* persuasive signaling scheme, and formalize its differences from *ex interim* persuasive schemes. Then, we show that *ex ante* persuasiveness can provide the sender with a utility that can be arbitrarily larger than that provided by *ex interim* persuasiveness. Motivated by the hardness results for the *ex interim* setting with simultaneous moves provided by Dughmi and Xu (2016), we study the problem of computing optimal *ex ante* signaling schemes. First, we prove a result of independent interest that plays a crucial role in the following proofs. More precisely, we show that, given a multi-player game and a behavioral strategy of a perfect-recall player, it is possible to find, in polynomial time, a realization-equivalent mixed strategy (defined on the normal form) with a polynomially-sized support. We show that an optimal *ex ante* signaling scheme may be computed in polynomial time in settings with two receivers and independent action types, which makes *ex ante* persuasive signaling schemes a persuasion tool which is applicable in practice. Moreover, we show that this result cannot be extended to settings with more than two receivers, as the problem of computing an optimal *ex ante* signaling scheme becomes NP-hard.

Bayesian Persuasion with Sequential Games

In our model, we assume a *sender* denoted by S and a set of *receivers* $\mathcal{R} = \{1, \dots, n\}$. Each receiver $i \in \mathcal{R}$ is faced with the problem of selecting actions from a set A_i with *a priori* uncertain payoffs. We adopt the perspective of the sender, whose goal is persuading the receivers to take actions which are favorable for her. The fundamental feature of our model is that receivers confront themselves in a *sequential* decision problem, which we describe as an *extensive-form game* (EFG) with imperfect information and perfect recall.

Payoffs are a function of the actions taken by the receivers and of an unknown *state of nature* θ, drawn from a set of potential realizations Θ. We follow the standard framework of Dughmi and Xu (2016), where each action a has a set of possible types Θ_a and in which a state of nature θ is a vector specifying the realized type of each action of the receivers, *i.e.*, $\theta \in \Theta = \times_{i \in \mathcal{R}} \times_{a \in A_i} \Theta_a$.[2] Furthermore, as also done by Dughmi and Xu (2016), we assume action types which are drawn independently from action-specific marginal probability distributions denoted by $\tilde{\pi}_a \in \text{int}(\Delta^{|\Theta_a|})$, where $\tilde{\pi}_a(t)$ is the probability of a having type $t \in \Theta_a$.[3] These marginal probability distributions form a common prior over the states of nature which we assume to be known explicitly to both sender and receivers. This common knowledge can be equivalently represented by the distribution $\mu_0 \in \Delta^{|\Theta|}$, where $\mu_0(\theta) = \prod_{i \in \mathcal{R}} \prod_{a \in A_i} \tilde{\pi}_a(\theta_a)$.

In the following, we provide some background on EFGs, describe the two models of persuasion (*ex interim* and *ex ante*), and highlight their differences with some examples.

Background on EFGs

An EFG—here denoted by Γ—is composed of a set H of nodes, each of which is identified by the ordered sequence of actions leading to it from the root node. The set of terminal nodes of the game is denoted by $Z \subseteq H$. The game is played by the receivers $\mathcal{R}$. A_i is the set of actions available to each receiver $i \in \mathcal{R}$. Let $A = \{A_i\}_{i \in \mathcal{R}}$. For each nonterminal node $h \in H \setminus Z$, we denote by, respectively, $P(h)$ and $A(h)$ the unique receiver acting at h and the set of actions available at that node. Imperfect information is represented via *information sets* (or *infosets*), which group together decision nodes which are indistinguishable for a certain receiver. For each receiver i, we denote her set of infosets by $\mathcal{I}_i$. $\mathcal{I}_i$ defines a partition of $\{h \in H \mid P(h) = i\}$. Each $I \in \mathcal{I}_i$ is such that $A(h) = A(h') \; \forall h, h' \in I$. With a little notation overload, we let $A(I)$ be the set of actions available at each

[2] Standard (*i.e.*, non Bayesian) EFGs can be represented by assigning to each Θ_a a singleton. Note that this model also encompasses Bayesian games *á la* Harsanyi (1967).

[3] $\text{int}(X)$ is the interior of set X, and $\Delta^{|X|}$ is the set of all probability distributions on X.

decision node in I. Receiver i has *perfect recall* if she has perfect memory of her past actions and observations.

We denote a *behavioral strategy* of receiver i by π_i. It corresponds to a vector defining a probability distribution over $A(I)$, $\forall I \in \mathcal{I}_i$. Given π_i, let $\pi_{i,I}$ be the (sub)vector representing the probability distribution at $I \in \mathcal{I}_i$. Letting, for each receiver i, $\Sigma_i = \times_{I \in \mathcal{I}_i} A(I)$, a *plan* is a vector $\sigma_i \in \Sigma_i$ which specifies an action for each of the receiver's infosets.[4] We denote by $\sigma_i(I)$ the action selected at infoset $I \in \mathcal{I}_i$. Letting $\Sigma = \times_{i \in \mathcal{P}} \Sigma_i$, we denote by $\sigma \in \Sigma$ the tuple which specifies the plan chosen by each receiver. Finally, a *mixed strategy* x_i is a probability distribution over Σ_i. We let $\mathcal{X}_i$ be the mixed strategy space of receiver i, and $\mathcal{X}$ be the set of joint probability distributions over Σ.

The *sequence form* (Koller, Megiddo, and von Stengel 1996; von Stengel 1996) of a game is a compact representation applicable to games with perfect recall. It decomposes strategies into sequences of actions and their realization probabilities. A sequence q_i for receiver i associated with a node h is a tuple specifying receiver i's actions on the path from the root to h. We denote the set of all sequences for receiver i by Q_i. A sequence is said terminal if, together with some sequences of the other receivers, leads to a terminal node. We let $q_\emptyset$ be the fictitious sequence leading to the root node and qa the extended sequence obtained by appending action a to q. A *sequence-form strategy* (usually called *realization plan*) for a receiver i is a function $r_i : Q_i \to [0,1]$ such that $r_i(q_\emptyset) = 1$ and, for each $I \in \mathcal{I}_i$ and sequence q leading to I, $-r_i(q) + \sum_{a \in A(I)} r_i(qa) = 0$ holds. We denote by $Q(I)$ the set of sequences originating in I. For each $q \in Q_i$, we denote by $I_\downarrow(q) \subseteq \mathcal{I}_i$ the set of infosets reachable by i after selecting q without making other intermediate moves, whereas $I_\uparrow(q) \in \mathcal{I}_i$ denotes the unique infoset where the last action of q was taken.[5] We call two strategies of receiver i *realization equivalent* if, for any fixed strategy of the other receivers, they induce the same distribution over the terminal nodes in Z.

Ex interim Persuasiveness

Let $u_S : \Sigma \times \Theta \to \mathbb{R}$ and $u_i : \Sigma \times \Theta \to \mathbb{R}$ be the payoff functions of the sender and receiver $i \in \mathcal{R}$. We assume that the sender is allowed to tailor signals to individual receivers through private communications. Let Ω_i be the set of signals available to receiver i, and let $\Omega = \times_{i \in \mathcal{R}} \Omega_i$. We assume that the sender has access to private information and her goal is designing a *signaling scheme* $\varphi : \Theta \to \Delta^{|\Omega|}$ to persuade the receivers to select actions which are favorable for her. We denote by φ_θ the probability distribution over Ω having observed θ. In the classical Bayesian persuasion framework by Kamenica and Gentzkow (2011), the receivers decide their behavior after observing the sender's signal and updating their posterior over Θ accordingly. The sender-receivers interaction goes as follows.

- The sender chooses φ and publicly discloses it.
- Nature draws a state $\theta \sim \mu_0$, observed by the sender.
- The sender draws a tuple $\omega \sim \varphi_\theta$ and privately sends signal ω_i to each receiver $i \in \mathcal{R}$.
- Each receiver i updates her posterior distribution knowing φ and having observed ω_i. Then, each of the receivers selects a plan $\sigma_i \in \Sigma_i$. Together, their joint choices form the tuple $\sigma = (\sigma_1, \dots, \sigma_n)$.
- Sender and receivers get, respectively, payoffs $u_S(\theta, \sigma)$ and $u_i(\theta, \sigma)$, for all $i \in \mathcal{R}$.

In this setting, a result similar to the *revelation principle* (see, *e.g.*, (Myerson 1979)) holds. Specifically, an *optimal signaling scheme* (*i.e.*, a signaling scheme maximizing the sender's expected utility) can always be obtained by restricting the set of signals Ω to the set of plans Σ; see Proposition 1 by Kamenica and Gentzkow (2011). In the following, we assume $\Omega = \Sigma$ (*i.e.*, the sender recommends a plan to follow to each receiver). The receivers have an incentive to follow the sender's recommendation $\hat{\sigma}_i$ if the recommended plan is preferred to any other action, conditional on the knowledge of $\hat{\sigma}_i$. We call this condition *ex interim persuasiveness*, which is precisely the kind of constraint characterizing a *Bayes Correlated Equilibrium* (BCE) (Bergemann and Morris 2013; 2016a). We remark that, according to the definition of BCE, the signaling scheme must necessarily be defined on plans and cannot be compactly represented by using sequences or actions.

Definition 1 (*Ex interim* persuasiveness)**.** *A signaling scheme $\varphi : \Theta \to \Delta^{|\Sigma|}$ is* ex interim persuasive *if the following holds for all $i \in \mathcal{R}$ and $\sigma_i, \sigma'_i \in \Sigma_i$:*

$$\sum_{\substack{\theta \in \Theta, \\ \sigma_{-i} \in \Sigma_{-i}}} \mu_0(\theta) \varphi_\theta(\sigma_i, \sigma_{-i}) \Big(u_i(\theta, (\sigma_i, \sigma_{-i})) - u_i(\theta, (\sigma'_i, \sigma_{-i})) \Big) \geq 0.$$

Definition 2. *A signaling scheme $\varphi : \Theta \to \Delta^{|\Sigma|}$ is a BCE if it is* ex interim *persuasive.*

Ex ante Persuasiveness

We introduce the setting in which the receivers have to decide whether to follow the sender's recommendations before actually observing them, basing their decision only on the knowledge of μ_0 and φ.[6] The interaction between sender and receivers goes as follows.

- The sender computes φ, and publicly discloses it.
- The receivers decide whether to adhere to the recommendations drawn according to φ or not.
- Nature draws a state $\theta \sim \mu_0$, observed by the sender.
- If $i \in \mathcal{R}$ decided to opt-in to the signaling scheme:
  - the sender draws $\hat{\sigma}_i \sim \varphi_\theta$ and privately communicates it to receiver i;
  - receiver i acts according to the recommended $\hat{\sigma}_i$.

[4] A plan is an action of the corresponding normal-form game, whose size is exponentially larger than the extensive form.

[5] When the context requires disambiguation between different games, we write $I^\Gamma_\downarrow(q)$ to denote the result for EFG Γ.

[6] As discussed in the introduction of the paper, the receivers' commitment to follow a certain signaling scheme is not an unrealistic assumption for the same reason why it is realistic to assume the sender's commitment power.

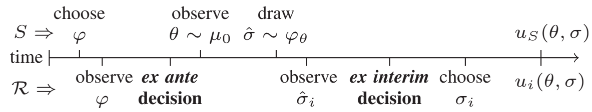

Figure 1: Interaction between sender and receivers in the *ex ante* and *ex interim* settings.

	In	Out	P
E	$(-1, 1)$	$(1, 0)$	$(0, 1/2)$
H	$(-1, -1)$	$(1, 0)$	$(0, 0)$

Figure 2: A game where *ex ante* persuasion guarantees the sender a higher expected utility with respect to *ex interim* persuasion.

- Sender and receivers get, respectively, payoffs $u_S(\theta, \sigma)$ and $u_i(\theta, \sigma)$, $\forall i \in \mathcal{R}$, where $\sigma_i = \hat{\sigma}_i$ if i adhered to the signaling scheme.

In this setting, the receivers adhere to the signaling scheme (*i.e.*, $\sigma_i = \hat{\sigma}_i$) if it is *ex ante* persuasive:

Definition 3 (*Ex ante* persuasiveness)**.** *The signaling scheme* $\varphi : \Theta \to \Delta^{|\Sigma|}$ *is* ex ante *persuasive if, for all* $i \in \mathcal{R}$ *and* $\sigma_i \in \Sigma_i$*, the following holds:*

$$\sum_{\substack{\theta \in \Theta, \sigma_i' \in \Sigma_i \\ \sigma_{-i} \in \Sigma_{-i}}} \mu_0(\theta) \varphi_\theta(\sigma_i', \sigma_{-i}) \Big(u_i(\theta, (\sigma_i', \sigma_{-i})) - u_i(\theta, (\sigma_i, \sigma_{-i})) \Big) \geq 0.$$

Such constraints define *Bayes Coarse Correlated Equilibria* (BCCE), *i.e.*, the generalization of coarse correlated equilibria to incomplete-information games, see Forges (1993), Cai and Papadimitriou (2014), Hartline, Syrgkanis, and Tardos (2015), and Caragiannis et al. (2015), Celli, Coniglio, and Gatti (2019b).[7]

Definition 4. *A signaling scheme* $\varphi : \Theta \to \Delta^{|\Sigma|}$ *is a BCCE if it is* ex ante *persuasive.*

Comparison Between Notions of Persuasiveness

Figure 1 summarizes the interaction flow between sender and receivers in the two aforementioned settings. The key difference is the time at which the receivers decide whether to adhere to the signaling scheme or not.

We also propose the following illustrative example (in the basic single-receiver setting) to further illustrate the main differences between the two notions of persuasiveness.

Example 1. *The incumbent of an industry wants to persuade a potential new entrant to the market. The market can be either easy (E), with probability 0.3, or hard (H), with the remaining probability. The incumbent knows the state of the market. The entrant has three possible actions available: entering the market (In), staying out of the market (Out), or proposing a partnership to the incumbent (P). Figure 2 depicts the utility matrix for the game (the first values are the incunbent's payoffs).*

[7] The set of (non Bayesian) coarse correlated equilibria is characterized by the constraints of Definition 3, with $|\Theta_a| = 1 \; \forall a \in A$.

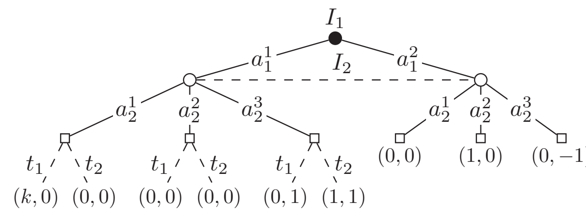

Figure 3: A game with two receivers in which action a_1^1 has two possible types t_1 and t_2. Terminal nodes report receivers' utilities.

The incumbent wants the entrant to stay out of the market, values its entrance negatively, and is indifferent towards a partnership. The entrant values entering the new market positively only when it has favorable conditions. A partnership in a hard market gives the entrant 0 (rather than a negative score) as no fixed costs have to be sustained. In this setting, forcing the entrant (contractually) to commit to following the incumbent's recommendations ex ante *is strictly better (in terms of expected utility) for the incumbent.*

An optimal ex ante *signaling scheme (e.g.,* $\varphi_E(In) = \varphi_E(Out) = \frac{1}{2}$, $\varphi_H(Out) = 1$*) guarantees the sender an expected utility of* 0.7*. An optimal* ex interim *signaling scheme (e.g.,* $\varphi_E(P) = 1$, $\varphi_H(Out) = \frac{11}{14}$, $\varphi_H(P) = \frac{3}{14}$*,) guarantees a sender's expected utility of 0.55. Therefore,* ex ante *persuasion provides a 27% increase in utility for the incumbent w.r.t. ex interim persuasion.*

We remark that the set of *ex ante* persuasive signaling schemes strictly includes the set of *ex interim* signaling schemes. In particular, an optimal *ex ante* persuasive signaling scheme may lead to an expected utility for the sender that is arbitrarily larger than the one she would obtain with an optimal *ex interim* scheme. This is shown by means of the following example.

Example 2. *Consider the game in Figure 3, with two receivers with one information set each (I_1 for receiver 1 and I_2 for receiver 2), and parametric in $k \gg 1$. Action $a_1^1 \in A_1$ is such that $\Theta_{a_1^1} = \{t_1, t_2\}$ and $\tilde{\pi}_{a_1^1}(t_1) = \tilde{\pi}_{a_1^1}(t_2) = 1/2$. The figure only reports the receivers' utilities, as we assume $u_S(\theta, \sigma) = u_1(\theta, \sigma) + u_2(\theta, \sigma)$, $\forall (\theta, \sigma)$. The signaling scheme with $\varphi'_{t_1}(a_1^1, a_2^1) = 1/2$, $\varphi'_{t_1}(a_1^2, a_2^2) = 1/2$, and $\varphi'_{t_2}(a_1^1, a_2^3) = 1$ is* ex ante persuasive*, but it is not* ex interim *persuasive. The optimal* ex interim *persuasive signaling scheme is such that $\varphi''_{t_1}(a_1^2, a_2^2) = 1$, and $\varphi''_{t_2}(a_1^1, a_2^3) = 1$. Signaling scheme φ' guarantees the sender an expected utility of $(k + 5)/4$, while signaling scheme φ'' guarantees $3/2$. Therefore, for increasing values of k, an optimal* ex ante *signaling scheme provides an arbitrarily larger utility than what can be obtained by* ex interim *persuasion.*

Positive Result

In the independent-actions setting, Dughmi and Xu (2016) show that computing an optimal *ex interim* signaling scheme is #P-hard even with a single receiver. Motivated by this negative result, we study the problem of computing an optimal (for the sender) *ex ante* persuasive signaling scheme.

We denote this problem by OPT-EA. It amounts to computing a Coarse Correlated Equilibrium (CCE) for the game of complete information obtained by treating Nature as a player with a *trivial* (*i.e.*, constant everywhere) payoff function and subject to having marginal strategies constrained to be μ_0.[8]

In contrast with the known hardness results for the *ex interim* setting, we show that OPT-EA with $|\mathcal{R}| = 2$ can be solved in polynomial time (see Theorem 5 below). To prove our main theorem, we first show how to build, in polynomial time, a *small* (*i.e.*, with a support of size upper bounded by a polynomial) mixed strategy which is realization-equivalent to a given behavioral strategy. Omitted proofs are presented in extended version of the paper (see Celli, Coniglio, and Gatti (2019a)).

Small–supported Mixed Strategies

Given a behavioral strategy profile π_i^* for a generic perfect-recall player i, we show (see Theorem 4 below) that it is always possible to compute in polynomial time some $x_i^* \in \mathcal{X}_i$ such that (i) it is realization-equivalent to π_i^* and (ii) it has a support of polynomial size.[9]

For each $\sigma_i \in \Sigma_i$, let $\xi(\sigma_i) := \{q \in Q_i | \exists I \in \mathcal{I}_i, \sigma_i(I) = q\}$ (*i.e.*, the set of sequences selected with probability 1 in a realization plan equivalent to σ_i). Analogously, $\forall \sigma = (\sigma_1, \sigma_2) \in \Sigma$ we denote by $\xi(\sigma)$ the set of tuples (q_1, q_2) such that $q_1 \in \xi(\sigma_1)$ and $q_2 \in \xi(\sigma_2)$. In the remainder of the section, we drop the dependency on i when it is not strictly necessary. We denote by M an $|Q_i| \times |\Sigma_i|$ matrix where $M(q, \sigma) = 1$ iff $q \in \xi(\sigma)$ and $M(q, \sigma) = 0$ otherwise. We denote by M_q the row of M specifying the plans containing q in their support. Let r^* be the $|Q_i|$-dimensional vector representing the realization plan of player i which is realization-equivalent to π^*. In order to compute x^*, it is enough to find an optimal solution to the LP $\max_{x \in \mathbb{R}_{\geq 0}^{|\Sigma_i|}} \{\mathbb{1}^\top x \quad \text{s.t.} \quad Mx \leq r^*\}$, which we denote by (A), which has a polynomial number of constraints and an exponential number of variables.

By relying on the assumption of perfect recall and proceeding by contradiction, we establish the following lemma:

Lemma 1. *An optimal solution x^* to (A) satisfies $Mx^* = r^*$.*

Proof. Consider a behavioral strategy π^* whose realization-equivalent realization plan is denoted by r^*. Since player i has perfect recall, there always exists at least a mixed strategy $\hat{x} \in \mathcal{X}_i$ realization equivalent to π^* (Maschler, Solan, and Zamir 2013, Th. 6.11). Therefore, the optimal value of (A) is 1 (since $\mathbb{1}^\top \hat{x} = 1$). Given $\hat{x} \in \mathcal{X}_i$, a distribution assigning to each sequence $q \in Q_i$ value $\sum_{\sigma \in \Sigma_i : q \in \xi(\sigma)} \hat{x}_\sigma$ is a valid realization plan. Therefore, if $x \in \Delta^{|\Sigma_i|}$ then Mx is a well defined realization plan for i. Now, by contradiction, assume that x^* is an optimal solution to (A) and that there exists $q' \in Q_i$ such that $M_{q'} x^* < r^*(q')$. Optimality implies $\mathbb{1}^\top x^* = 1$ and, therefore, $x^* \in \Delta^{|\Sigma_i|}$. Let $Mx^* = r$. We have $r(q') < r^*(q')$. Since the sequence-form constraints hold, there must exist some $q'' \in Q_i$ such that $r(q'') > r^*(q'')$. This leads to a contradiction since x^* would not be a feasible solution. □

We now characterize an optimal solution to (A) by two properties which are proven by relying on Lemma 1 and on the fact that, as LP (A) contains a polynomial number of constraints, it admits an optimal basic solution with only a polynomial number of nonzero variables:

Theorem 2. *These two properties hold:*

- *(i) An optimal solution x^* to (A) is a normal-form strategy ($x^* \in \mathcal{X}_i$) realization equivalent to r^*,*
- *(ii) there exists an optimal solution x^* with* $\text{supp}(x^*)$ *of polynomial size.*

Let $\mathcal{D}$ be the dual of problem (A). By showing that an optimal plan corresponding to a violated dual constraint can be found in polynomial time by backward induction, we can establish the following:

Lemma 3. *$\mathcal{D}$ admits a polynomial-time separation oracle.*

Next, by relying on Lemma 3 and on the *ellipsoid method* for solving LPs we prove a result which is the basis for our main theorem, Theorem 5 (whose statement and proof are given in full in the next subsection):

Theorem 4. *Given an EFG, a perfect-recall player i, and a behavioral strategy profile π^* for i (with the realization-equivalent realization plan r^*), a solution to LP (A) can be found in polynomial time.*

Finally, we show that we can efficiently compute a solution with support size of at most $|Q_i|$ by applying the ellipsoid method for at most a polynomial number of iterations:

Corollary 1. *A* basic feasible solution *to (A) can be computed in polynomial time.*

Optimal *Ex Ante* Persuasive Schemes

Computing an *ex ante* persuasive signaling scheme is equivalent to computing a CCE for an EFG of complete information where Nature is treated as a player with constant utility and marginal strategies constrained to be equal to μ_0. We focus on the setting where $|\mathcal{R}| = 2$ and show that OPT-EA can be solved in polynomial time. We reason over an auxiliary game where each action of the receivers is followed by one of Nature's nodes, determining its type. Marginal probabilities $\tilde{\pi}$ determining action types are treated as behavioral strategies of the Nature player, which we denote by N. Formally:

Definition 5. *Given an EFG Γ describing the interaction between receivers and a set of marginal distributions $\{\tilde{\pi}_a \in \text{int}(\Delta^{|\Theta_a|})\}_{a \in A}$, the auxiliary game $\hat{\Gamma}$ is an EFG such that:*

- *It has a set of players $\mathcal{R} \cup \{\mathsf{N}\}$.*
- *For each receiver $i \in \mathcal{R}$, her utility function is the same as in Γ, i.e., $\forall (\theta, \sigma) \in \Theta \times \Sigma$, $u_i(\theta, \sigma) = \hat{u}_i(\theta, \sigma)$. Nature has $\hat{u}_{\mathsf{N}}(\cdot) = k \in \mathbb{R}$ constant everywhere.*

[8] Notice that finding an optimal CCE with two players and Nature is hard in the worst-case, while our problem is a variation.

[9] As customary, we define the support of a mixed strategy $x_i \in \mathcal{X}_i$ as $\text{supp}(x_i) := \{\sigma_i \in \Sigma_i | x_i(\sigma_i) > 0\}$.

- *The receivers have the same information structures as in Γ, i.e., $\forall i \in \mathcal{R}$, $\mathcal{I}_i = \hat{\mathcal{I}}_i$, and $\forall q \in Q_i$, $I_\downarrow^\Gamma(q) = I_\downarrow^{\hat\Gamma}(q)$.*
- *$\forall i \in \mathcal{R}$, each $a \in A_i$ is immediately followed by a singleton infoset $I \in \mathcal{I}_\mathsf{N}$ such that $A(I) = \Theta_a$.*
- *$\forall I \in \mathcal{I}_\mathsf{N}$, with I following $a \in A$, N selects actions (types) at I according to the marginal distribution $\tilde\pi_a$.*

The first step is devising an LP to compute a BCCE with a polynomial number of constraints and an exponential number of variables. We do so by providing an LP to find an optimal CCE over $\hat\Gamma$. First, notice that θ is a plan of player N in $\hat\Gamma$. A distribution in $\Delta^{|\Theta|}$ is a mixed strategy of N. Denote by μ^* the mixed strategy realization equivalent to $\tilde\pi$ computed (in poly-time) as in the proof of Theorem 4. Let $\Theta^* := \text{supp}(\mu^*)$. Due to Corollary 1, the set Θ^* has polynomial size. Then, we write the problem as a function of $\gamma \in \Delta^{|\Sigma\times\Theta^*|}$ (i.e., we look for a correlated distribution in $\hat\Gamma$, encompassing the Nature player). Let v_i be the $|\mathcal{I}_i|$-dimensional vector of variables of the dual of the best-response problem for receiver i in sequence form. Moreover, we employ sparse $(|\mathcal{R}|+1)$-dimensional matrices describing the utility function of sender and receivers for the profiles (θ, q_1, q_2) leading to terminal nodes of $\hat\Gamma$. We denote them by $U_i \in \mathbb{R}^{|\Theta^*|\times|Q_1|\times|Q_2|}$, with $i \in \mathcal{R}\cup\{S\}$.[10] In the following, we let $q = (q_1, q_2)$. The problem of computing a CCE over $\hat\Gamma$ reads:

$$\max_{\substack{\gamma\geq 0,\\ v_1, v_2}} \sum_{\substack{\theta\in\Theta^*\\ \sigma\in\Sigma}} \gamma(\theta,\sigma) \sum_{q\in\xi(\sigma)} U_S(\theta, q) \tag{1a}$$

$$\text{s.t.} \quad \sum_{\substack{\theta\in\Theta^*\\ \sigma\in\Sigma}} \gamma(\theta,\sigma) \sum_{q\in\xi(\sigma)} U_i(\theta, q) \geq \sum_{\substack{I'\in\mathcal{I}_i:\\ I'\in I_\downarrow(q_\emptyset)}} v_i(I') \quad \forall i \in \mathcal{R} \tag{1b}$$

$$v_1(I_\uparrow(q_1)) - \sum_{I'\in I_\downarrow(q_1)} v_1(I') + - \sum_{\substack{\theta\in\Theta^*\\ \sigma\in\Sigma}} \gamma(\theta,\sigma) \sum_{q_2\in\xi(\sigma_2)} U_1(\theta, q_1, q_2) \geq 0 \quad \forall q_1 \in Q_1 \tag{1c}$$

$$v_2(I_\uparrow(q_2)) - \sum_{I'\in I_\downarrow(q_2)} v_2(I') + - \sum_{\substack{\theta\in\Theta^*\\ \sigma\in\Sigma}} \gamma(\theta,\sigma) \sum_{q_1\in\xi(\sigma_1)} U_2(\theta, q_1, q_2) \geq 0 \quad \forall q_2 \in Q_2 \tag{1d}$$

$$\sum_{\sigma\in\Sigma} \gamma(\theta,\sigma) = \mu^*(\theta) \qquad \forall\theta\in\Theta^*. \tag{1e}$$

We make the following observations on the above LP, which we denote by (B):

- The left term of constr. (1b) is the expected utility of i at the equilibrium. Incentive constraints (1c) and (1d) are compactly encoded by exploiting the sequence form. Intuitively, we decompose the best-response problem locally at each infoset. The constraints impose that the utility at the equilibrium be no smaller than the value achieved when playing the plan obtained by letting the receiver best respond in each infoset.
- Constraint (1e) forces Nature's marginal distribution to be equal to the prior μ^*.
- Once a solution γ^* to (B) has been computed, an optimal solution to OPT-EA is the signaling scheme which, having observed θ, recommends σ with probability $\gamma^*(\theta,\sigma)/\mu^*(\theta)$.

The following key positive result holds:

Theorem 5. *OPT-EA can be solved in polynomial time when $|\mathcal{R}| \leq 2$.*

Proof. Let $\mathcal{D}_B$ be the dual of (B). Let α_1, α_2 be the dual variables of constraints (1b), $\beta_1 \in \mathbb{R}^{|Q_1|}$ and $\beta_2 \in \mathbb{R}^{|Q_2|}$ the dual variables of (1c) and (1d), and $\delta \in \mathbb{R}^{|\Theta^*|}$ the dual variables of (1e). We show that, given $(\bar\alpha_1, \bar\alpha_2, \bar\beta_1, \bar\beta_2, \bar\delta)$, the problem of finding either a hyperplane separating the solution from the feasible set of $\mathcal{D}_B$ or proving that no such hyperplane exists can be solved in polynomial time. Along the lines of Theorem 4, this implies that (B) is solvable in polynomial time by the ellipsoid method. As the number of dual constraints corresponding to variables v_i is linear, all these constraints can be checked efficiently for violation. Besides those, the dual problem $\mathcal{D}_B$ features the following constraint for each $(\theta,\sigma) \in \Theta^*\times\Sigma$:

$$\sum_{i\in\mathcal{R}} \sum_{q\in\xi(\sigma)} U_i(\theta,q)\bar\alpha_i + \frac{\bar\delta(\theta)}{\mu^*(\theta)} - \sum_{q\in\xi(\sigma)} U_S(\theta,q) + - \sum_{q\in Q_1\times\xi(\sigma_2)} U_1(\theta,q)\bar\beta_1(q_1) + - \sum_{q\in\xi(\sigma_1)\times Q_2} U_2(\theta,q)\bar\beta_2(q_2) \geq 0.$$

Given $(\bar\alpha_1, \bar\alpha_2, \bar\beta_1, \bar\beta_2, \bar\delta)$, the *separation problem* of finding a maximally violated inequality of $\mathcal{D}_B$ reads:

$$\min_{\theta,\sigma} \left\{ \sum_{q\in\xi(\sigma)} \left[\sum_{i\in\mathcal{R}} U_i(\theta,q) - U_S(\theta,q) \right] + \frac{\bar\delta(\theta)}{\mu^*(\theta)} + - \sum_{q\in Q_1\times\xi(\sigma_2)} U_1(\theta,q)\bar\beta_1(q_1) - \sum_{q\in\xi(\sigma_1)\times Q_2} U_2(\theta,q)\bar\beta_2(q_2) \right\}.$$

A pair (θ,σ) yielding a violated inequality exists iff the separation problem admits an optimal solution of value < 0. If such a (θ,σ) exists, it can be determined in polynomial time by enumerating all the (polynomially many) $(\theta, z) \in \Theta^*\times\hat Z$, where $\hat Z$ is the outcome set of $\hat\Gamma$. For each pair (θ, z), we look for a $\sigma\in\Sigma$ which, together with some actions of N, minimizes the objective function of the separation problem and could lead to z. The procedure halts as soon as a plan σ such that (θ,σ) yielding a violated inequality is found; if it terminates without finding any, $\mathcal{D}_B$ has been solved. First, by fixing a pair (θ, z) the first two terms of the objective function are completely determined. The remaining terms can be minimized independently for each receiver. Let us consider the problem of finding $\sigma_2 \in \Sigma_2$ (the other one is solved analogously). It reads:

$$\max_{\sigma_2\in\Sigma_2} \left\{ \sum_{q_1\in Q_1} \sum_{q_2\in\xi(\sigma_2)} U_1(\theta, q_1, q_2)\bar\beta_1(q_1) \right\},$$

[10] U_i employs both the sequence form (for receivers), and plans of N. However, polynomiality of Θ^* implies polynomiality of U_i.

subject to the constraint that σ_2 be an admissible plan for the given z (*i.e.*, given the solution plan, it has to be possible to reach z together with some actions of the other players). This problem can be solved in poly-time as shown in Algorithm 1, where $\mathcal{I}_i^z$ and Q_i^z are, respectively, the set of infosets and sequences of i encountered on the path from the root to z. Once Q^* has been determined by visiting each $I \in \mathcal{I}_2$, the corresponding optimal σ_2 can be built directly. As in Corollary 1, an optimal solution to B has polynomial support size. Then, it is used to determine an optimal solution to OPT-EA in poly-time. □

Algorithm 1 `Separation plan search for` (θ, z)

function F(I,Q^*) ▷ $I \in \mathcal{I}_2$ is the current infoset
  $\hat{Q} \leftarrow \emptyset$, $w(q_2) \leftarrow -\infty \quad \forall q_2 \in Q_2$
  if $I \in \mathcal{I}_2^z$ **then**
    $\hat{Q} \leftarrow \{q_2 \in Q_2 | q_2 \in Q(I)$ and $q_2 \in Q_2^z\}$
  else
    $\hat{Q} \leftarrow Q(I)$
  for $q_2 \in \hat{Q}$ **do**

$$w(q_2) \leftarrow \sum_{q_1 \in Q_1} U_1(\theta, q_1, q_2)\bar{\beta}_1(q_1) + \\ + \sum_{I' \in I_\downarrow(q_2)} \mathrm{F}(I', Q^*)$$

  $q_2^* = \arg\max_{q_2 \in Q_2} w(q_2)$
  $Q^* \leftarrow Q^* \cup \{q_2^*\}$
  return $w(q_2^*)$

Negative Result

We conclude by showing that the previous approach cannot be extended to settings where $|\mathcal{R}| > 2$ and that, in particular, the border between easy and hard cases coincides with $|\mathcal{R}| = 2$. Indeed, the fact that computing an optimal CCE for a three-player EFG is NP-hard (von Stengel and Forges 2008, Th. 1.3) directly implies the following:

Theorem 6. *OPT-EA is NP-hard when* $|\mathcal{R}| > 2$.

Proof. Let $|\mathcal{R}| = 3$ and, $\forall a \in A$, $|\Theta_a| = 1$. Then, the problem amounts to computing an optimal CCE for a three player EFG, which is NP-hard since the reduction described in Th. 1.3 by von Stengel and Forges (2008) applies. □

For completeness, we also provide the following result, showing that the separation problem of OPT-EA is NP-hard when $|R| > 2$.

Theorem 7. *Computing an optimal solution to the separation problem of OPT-EA is NP-hard when* $|\mathcal{R}| > 2$.

Proof. Consider the case in which $\mathcal{R} = 3$ and $\mathcal{D}_B$ is adapted accordingly. Let $q = (q_1, q_2, q_3)$. Given the dual variables $(\bar{\alpha}_1, \bar{\alpha}_2, \bar{\alpha}_3, \bar{\beta}_1, \bar{\beta}_2, \bar{\beta}_3, \bar{\delta})$, the separation problem reads:

$$\min_{\theta,\sigma} \Bigg\{ \sum_{q \in \xi(\sigma)} \Bigg[\sum_{i \in \mathcal{R}} U_i(\theta, q) - U_S(\theta, q) \Bigg] + \frac{\bar{\delta}(\theta)}{\mu^*(\theta)} + \\ - \sum_{q \in Q_1 \times \xi(\sigma_2) \times \xi(\sigma_3)} U_1(\theta, q)\bar{\beta}_1(q_1) \\ - \sum_{q \in \xi(\sigma_1) \times Q_2 \times \xi(\sigma_3)} U_2(\theta, q)\bar{\beta}_2(q_2) + \\ - \sum_{q \in \xi(\sigma_1) \times \xi(\sigma_2) \times Q_3} U_3(\theta, q)\bar{\beta}_3(q_3) \Bigg\}.$$

Consider a setting with the following features: $\forall \theta \in \Theta^*$, $\bar{\delta}(\theta) = 0$ (a valid assumption since $\delta \in \mathbb{R}^{|\Theta^*|}$); $\forall (\theta, q) \in \Theta^* \times (\times_{i \in \mathcal{R}} Q_i)$, $U_S(\theta, q) = U_1(\theta, q)$; $\forall (\theta, q) \in \Theta^* \times (\times_{i \in \mathcal{R}} Q_i)$, $U_2(\theta, q) = U_3(\theta, q) = 0$. Then, finding a maximally violated inequality corresponds to solving: $\arg\max_{\theta, \sigma_2, \sigma_3} \{\sum_{q \in Q_1 \times \xi(\sigma_2) \times \xi(\sigma_3)} U_1(\theta, q)\bar{\beta}_1(q_1)\}$. Let $U_1' \in \mathbb{R}^{|\Theta^*| \times |Q_2| \times |Q_3|}$ be such that, for each (θ, q_2, q_3), $U_1'(\theta, q_2, q_3) = \sum_{q_1 \in Q_1} U_1(\theta, q_1, q_2, q_3)\bar{\beta}(q_1)$. If Θ^* is a singleton, the problem becomes $\arg\max_{\sigma_2, \sigma_3} \{\sum_{(q_2, q_3) \in \xi(\sigma_2) \times \xi(\sigma_3)} U_1'(q_1, q_2)\}$. This is a joint best-response problem between receivers 2 and 3, which is known to be NP-hard (von Stengel and Forges 2008). This concludes our proof. □

The last result is worth some further remarks. Since the separation problem of $\mathcal{D}_B$ is NP-hard, this implies, as a consequence of the equivalence between optimization and separation (see Theorem 3.1 of (Grötschel, Lovász, and Schrijver 1981)) that $\mathcal{D}_B$ is NP-hard for at least one linear objective function. Theorem 7 shows that one such objective function is precisely the one obtained from the RHS of B.

Discussion

We have studied persuasion in the multi-receiver setting with private signals, introducing, for the first time, a model encompassing receivers with sequential interactions, as well as the notion of *ex ante* persuasiveness. In contrast with previous complexity results on computing optimal CCEs and optimal *ex interim* persuasive schemes, we show that with $|\mathcal{R}| \leq 2$ an optimal *ex ante* scheme can be computed in polynomial time with the ellipsoid method by relying on a polynomial-time separation oracle. This result is rather surprising since a very similar problem (the computation of a CCE with two players and Nature) is known to be hard. We also show that $|\mathcal{R}| = 2$ constitutes the border between easy and hard cases as, even for $|\mathcal{R}| = 3$, the problem is NP-hard.

In the future, we are interested both in combining other forms of correlations with Bayesian persuasion, *e.g.*, extensive-form correlations, and in investigating forms of perfection in sequential information-design problems. We also plan to assess the scalability of the method we proposed for solving OPT-EA with simplex-based column generation algorithms which are, in practice, more efficient than the ellipsoid method, also employing techniques for solving the separation oracle along the lines of (Amaldi, Coniglio, and Gualandi 2010; 2014; Coniglio and Tieves 2015) to achieve a faster convergence.

Acknowledgments

This work is partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets".

References

Alonso, R., and Câmara, O. 2016. Persuading voters. *AM ECON REV* 106(11):3590–3605.

Amaldi, E.; Coniglio, S.; and Gualandi, S. 2010. Improving cutting plane generation with 0-1 inequalities by bi-criteria separation. In *International Symposium on Experimental Algorithms (SEA)*, 266–275. Springer.

Amaldi, E.; Coniglio, S.; and Gualandi, S. 2014. Coordinated cutting plane generation via multi-objective separation. *Mathematical Programming* 143(1-2):87–110.

Arieli, I., and Babichenko, Y. 2016. Private Bayesian persuasion. *Available at SSRN 2721307*.

Babichenko, Y., and Barman, S. 2016. Computational aspects of private Bayesian persuasion. *arXiv:1603.01444*.

Badanidiyuru, A.; Bhawalkar, K.; and Xu, H. 2018. Targeting and signaling in ad auctions. In *ACM-SIAM SODA*, 2545–2563. SIAM.

Bardhi, A., and Guo, Y. 2018. Modes of persuasion toward unanimous consent. *THEOR ECON* 13(3):1111–1149.

Bergemann, D., and Morris, S. 2013. Robust predictions in games with incomplete information. *ECONOMETRICA* 81(4):1251–1308.

Bergemann, D., and Morris, S. 2016a. Bayes correlated equilibrium and the comparison of information structures in games. *THEOR ECON* 11(2):487–522.

Bergemann, D., and Morris, S. 2016b. Information design, Bayesian persuasion, and Bayes correlated equilibrium. *AM ECON REV* 106(5):586–91.

Cai, Y., and Papadimitriou, C. 2014. Simultaneous bayesian auctions and computational complexity. In *ACM EC*, 895–910.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; Kyropoulou, M.; Lucier, B.; Renato, P.; Tardos, E.; et al. 2015. Bounding the inefficiency of outcomes in generalized second price auctions. *J ECON THEORY* 156(C):343–388.

Castiglioni, M.; Celli, A.; and Gatti, N. 2019. Persuading voters: It's easy to whisper, it's hard to speak loud. *arXiv preprint arXiv:1908.10620*.

Celli, A.; Coniglio, S.; and Gatti, N. 2019a. Bayesian persuasion with sequential games. *arXiv preprint arXiv:1908.00877*.

Celli, A.; Coniglio, S.; and Gatti, N. 2019b. Computing optimal ex ante correlated equilibria in two-player sequential games. In *AAMAS*, 909–917.

Coniglio, S., and Tieves, M. 2015. On the generation of cutting planes which maximize the bound improvement. In *International Symposium on Experimental Algorithms*, 97–109. Springer.

Dughmi, S., and Xu, H. 2016. Algorithmic Bayesian persuasion. In *ACM STOC*, 412–425. ACM.

Dughmi, S., and Xu, H. 2017. Algorithmic persuasion with no externalities. In *ACM EC*, 351–368.

Dughmi, S.; Immorlica, N.; and Roth, A. 2014. Constrained signaling in auction design. In *ACM-SIAM SODA*, 1341–1357.

Dughmi, S. 2017. Algorithmic information structure design: a survey. *ACM SIGEC EX* 15(2):2–24.

Dughmi, S. 2018. On the hardness of designing public signals. *GAME ECON BEHAV*.

Emek, Y.; Feldman, M.; Gamzu, I.; Leme, R. P.; and Tennenholtz, M. 2012. Signaling schemes for revenue maximization. In *ACM EC*, 514–531.

Forges, F. 1993. Five legitimate definitions of correlated equilibrium in games with incomplete information. *THEOR DECIS* 35(3):277–310.

Goldstein, I., and Leitner, Y. 2018. Stress tests and information disclosure. *J ECON THEORY* 177:34–69.

Grötschel, M.; Lovász, L.; and Schrijver, A. 1981. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* 1(2):169–197.

Harsanyi, J. C. 1967. Games with incomplete information played by "bayesian" players. *MANAGE SCI* 14(3):159–182, 320–334, 486–502.

Hartline, J.; Syrgkanis, V.; and Tardos, E. 2015. No-regret learning in Bayesian games. In *NeurIPS*, 3061–3069.

Kamenica, E., and Gentzkow, M. 2011. Bayesian persuasion. *AM ECON REV* 101(6):2590–2615.

Kamenica, E. 2018. Bayesian persuasion and information design. *ANNU REV ECON* 11.

Kaplan, T. R., and Zamir, S. 2000. The strategic use of seller information in private-value auctions. Technical report, Hebrew University, Center For Rationality.

Koller, D.; Megiddo, N.; and von Stengel, B. 1996. Efficient computation of equilibria for extensive two-person games. *GAME ECON BEHAV* 14(2):247–259.

Leme, R. P.; Syrgkanis, V.; and Tardos, É. 2012. Sequential auctions and externalities. In *ACM-SIAM SODA*, 869–886. SIAM.

Maschler, M.; Solan, E.; and Zamir, S. 2013. *Game Theory*. Cambridge University Press.

Myerson, R. 1979. Incentive compatibility and the bargaining problem. *ECONOMETRICA* 47(1):61–73.

Rabinovich, Z.; Jiang, A.; Jain, M.; and Xu, H. 2015. Information disclosure as a means to security. In *AAMAS*, 645–653.

Rayo, L., and Segal, I. 2010. Optimal information disclosure. *J POLIT ECON* 118(5):949–987.

Taneva, I. 2019. Information design. *AM ECON J-MICROECON*. Forthcoming.

von Stengel, B., and Forges, F. 2008. Extensive-form correlated equilibrium: Definition and computational complexity. *MATH OPER RES* 33(4):1002–1022.

von Stengel, B. 1996. Efficient computation of behavior strategies. *GAME ECON BEHAV* 14(2):220–246.

Wang, Y. 2013. Bayesian persuasion with multiple receivers. *Available at SSRN 2625399*.

Xu, H.; Rabinovich, Z.; Dughmi, S.; and Tambe, M. 2015. Exploring information asymmetry in two-stage security games. In *AAAI*, 1057–1063.

Xu, H.; Freeman, R.; Conitzer, V.; Dughmi, S.; and Tambe, M. 2016. Signaling in Bayesian Stackelberg games. In *AAMAS*, 150–158.

Xu, H. 2019. On the tractability of public persuasion with no externalities. *CoRR* abs/1906.07359.