UNIVERSITY of York

Discussion Papers in Economics

No. 16/08

Rationalizable Persuasion

Makoto Shimoji

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

Rationalizable Persuasion*

Makoto Shimoji[†]

July 15, 2016

Abstract

We analyze multi-receiver Bayesian persuasion games with heterogeneous beliefs, originating from Kamenica and Gentzkow (2011). We directly examine the sender's messages, which are supported by rationalizability. With no strategic interactions at the stage game, the sender's optimization problem can be viewed as a set of linear programming problems. We also show some generic properties of solutions. With strategic interactions at the stage game, we provide examples on two aspects of communication (only arising with the receivers' strategic interactions): "talking about others privately" and "tacit understandings", of which the latter is implied by forward induction.

Keywords: Bayesian Persuasion Games, Multiple Receivers, Heterogeneous Beliefs

JEL Classification: C72, D83

*I would like to thank Pierpaolo Battigalli and Joel Sobel for their comments and suggestions. I would also like to thank Subir Chattopadhyay, Anton Kolotilin, Alan Krause, Shintaro Miura, Tatsuyoshi Okimoto and Wataru Tamura for their comments and suggestions and Ricardo Alonso for sharing the older version of Alonso and Câmara (2016a).

†Department of Economics and Related Studies, University of York, Heslington, York, YO10 5DD, UK. Email: makoto.shimoji@york.ac.uk

1 Introduction

In a seminal study, Kamenica and Gentzkow (2011) added a new dimension to the literature on strategic information transmission by introducing a novel type of message space.[1] In the pioneering work by Crawford and Sobel (1982), the privately informed sender chooses different messages depending on her types. The departure in Kamenica and Gentzkow (2011) is that, while assuming no private information, the message which the sender chooses involves randomness conditional on the unknown state. Kamenica and Gentzkow (2011) analyzed games with this type of messages which we call *Bayesian persuasion games*.

In Kamenica and Gentzkow (2011), there is one sender and one receiver who share the common belief regarding the states. The sender chooses a message consisting of (i) the set of signals, and (ii) the signal distribution conditional on the states.[2] After observing the message and the realization of signal, the receiver updates her belief about the states and chooses a best response. By identifying the receiver's best response for each message and signal realization, the sender chooses the message which maximizes her expected payoff. The sender influences the receiver's behavior through the receiver's belief update. Kamenica and Gentzkow (2011) provided a condition under which the sender benefits from persuasion and analyzed optimal messages. Their technical contribution is that for the sender's optimization problem, she can only focus on the interim (updated) belief which needs to satisfy a consistency condition called Bayes plausibility.[3]

Our main goal in this paper is to show that we can analyze *multi-receiver* Bayesian persuasion games with *heterogeneous beliefs*.[4] In order to do so, while the receivers' best responses still need to be examined in the stage game (the

[1]While Rayo and Segal (2010) employ the same type of message space, the receiver in their setting has private information and binary actions. See Sobel (2013) for a survey.

[2]Kamenica and Gentzkow (2011) use the word "signal" for our definition of message.

[3]This is also known as the concavification approach. See Aumann and Maschler (1995).

[4]Even with heterogeneous beliefs, one can solve the sender's optimization problem as if they share the common prior since in the sender's and the receivers' payoff functions, each π – sender's message – is multiplied by their priors. There is no need for such a detour in our approach.

game after signal realization) for each possible message and signal realization, we explicitly examine the sender's messages (instead of the interim belief) for her optimization problem. We provide two examples to motivate our study.

Heterogeneous Beliefs. You and I are at a cafe which only offers two items: Apple Pie (A) and Banana Cake (B). The person sitting next to us placed an order. We only know that "she" ordered either A or B. If I ask a friend of mine working at this restaurant, "he" would tell us what she ordered. We decide to bet on whether she ordered A or B, which we refer to as the states. You name a state. If you are correct (incorrect), you (I) win and I (you) lose. The payoff for win is one and the payoff for loss is zero for each of us. You believe that it is more likely that she ordered B. Let p_1^0 be the probability you assign to A and hence $p_1^0 \in (0, \frac{1}{2})$; you would choose B with the initial belief p_1^0.

Before you name A or B, I send a message, π, which takes the form of a distribution of signals $\{a, b\}$ conditional on the states $\{A, B\}$; I choose two values $\pi(a \mid A) = \alpha \in [0, 1]$ and $\pi(a \mid B) = \beta \in [0, 1]$. For example, $\alpha = 1$ and $\beta = 0$ means that signal realization reveals the state. After choosing α and β, my friend working at the cafe checks what she ordered and generates a signal according to π. After observing the signal, you update your belief via Bayes rule:

$$p_1^\pi(A \mid a) = \frac{\alpha p_1^0}{\alpha p_1^0 + \beta(1-p_1^0)} \quad \text{and} \quad p_1^\pi(A \mid b) = \frac{(1-\alpha)p_1^0}{(1-\alpha)p_1^0 + (1-\beta)(1-p_1^0)}.$$

Given your updated belief, you choose A or B to maximize your expected payoff. Two observations: (i) I do not choose a signal, and (ii) your choice will be influenced by changing your belief, i.e., by choosing α and β.

I can predict how you would behave for each possible message and signal realization. I can therefore identify the message maximizing my expected payoff. This decision also depends on my belief. Let p_S^0 be the probability I assign to A.[5] Suppose first that $p_S^0 > p_1^0$. I then simply let you choose B by

[5]See Example 2 in Subsection 3.1 for the formal analysis.

sending no message. Suppose instead $p_S^0 < p_1^0$. In this case, I would want you to choose A instead. The optimal message is $\alpha = 1$ and $\beta = \frac{p_1^0}{1-p_1^0}$, and you then have $p_1^\pi(A \mid a) = \frac{1}{2}$ and $p_1^\pi(A \mid b) = 0$; you choose A after observing a or B after observing b. I attempt to confuse you; given your behavior, I prefer higher β and $1 - \alpha$. Since p_S^0 is small, I can accept $\alpha = 1$. Instead, I choose the highest β (i.e., $\frac{p_1^0}{1-p_1^0}$); you do not choose A after observing a if $\beta > \frac{p_1^0}{1-p_1^0}$.

Multiple Receivers. A teacher (sender) has a class with two students, 1 and 2 (receivers). They only have the final exam, before which the teacher decides to offer a mock exam. A student chooses either "high effort" or "low effort" for the exams. The two grades are "pass" and "fail". Two states are such that one student is "H" (pass is guaranteed) and the other student is "L" (pass is achieved if and only if she puts in high effort). Although the teacher and two students all believe that it is more likely that student 2 is H (student 1 is L), their confidences are different; the probabilities assigned to this state are 0.8 by the teacher, 0.7 by student 1 and 0.6 by student 2.

The mock exam is a surprise to the students (hence low effort from both) so that their performance only reflects their types. The teacher announces her grading policy (message) – a joint-distribution of two grades conditional on the states – when they take the mock exam. Given the grade and the grading policy, each student updates her belief regarding her type and decides how much effort to put in *for the final.* While effort is costly, a good performance on the final is rewarding. If a student is L, her payoff is zero if her effort (for the final) is low and one if her effort is high. If a student is H, her payoff is one if her effort is low and zero if her effort is high. The teacher believes that *the students should have the same academic workload irrespective of their types*; the teacher's payoff is one if they choose the same effort (low or high) and zero otherwise.

Figure 1 shows the optimal grading policy (message) for the teacher.[6] The

[6]The corresponding analysis can be found as Example 3 in Subsection 3.2, which has a slightly different payoff for the sender; $\varepsilon > 0$. The same solution is obtained even if $\varepsilon = 0$.

		2	
		fail	pass
1	fail	$\frac{2}{3}$	0
	pass	0	$\frac{1}{3}$

(a) $1 = L$ and $2 = H$

		2	
		fail	pass
1	fail	$\frac{2}{9}$	0
	pass	$\frac{7}{9}$	0

(b) $1 = H$ and $2 = L$

Figure 1: Grading Policy

interim beliefs for students 1 and 2 are respectively:

$$p_1^\pi(L \mid \text{fail}) = \frac{\frac{7}{10}\frac{2}{3}}{\frac{7}{10}\frac{2}{3}+\frac{3}{10}\frac{2}{9}} = \frac{7}{8} \qquad p_1^\pi(L \mid \text{pass}) = \frac{\frac{7}{10}\frac{1}{3}}{\frac{7}{10}\frac{1}{3}+\frac{3}{10}\frac{7}{9}} = \frac{1}{2}$$
$$p_2^\pi(L \mid \text{fail}) = \frac{\frac{4}{10}}{\frac{6}{10}\frac{2}{3}+\frac{4}{10}} = \frac{1}{2} \qquad p_2^\pi(L \mid \text{pass}) = \frac{\frac{4}{10}\cdot 0}{\frac{6}{10}\frac{1}{3}+\frac{4}{10}\cdot 0} = 0.$$

Each student chooses high (low) effort if she receives a fail (pass). Signals are jointly distributed, implying that individual persuasion is not optimal.[7]

Remark 1 *In the examples above, each receiver is indifferent between two actions after observing a certain signal. One may wonder what would happen if she chooses the other action. We will discuss this in Subsection 5.2.*

Previous studies have analyzed these two settings separately.[8] Kamenica and Gentzkow (2011) suggested that their approach can be extended to unlinked Bayesian persuasion games with multiple receivers where (i) the sender's preferences are separable with respect to the receivers' actions, and (ii) each receiver only cares about her action. Alonso and Câmara (2016b) and Wang

[7]In this particular setting, (i) if the teacher would like each to choose high effort, she would only persuade student 2, and (ii) if the teacher would like each to have a pass on the final, she would only need to reveal the state.

[8]Heterogeneous beliefs and multiple receivers have been considered in the context of cheap talk. For example, see Miura and Yamashita (2014) for the former, and Farrell and Gibbons (1989) and Goltsmand and Pavlov (2011) for the latter.

(2013) analyze voting games with multiple receivers. Alonso and Câmara (2016a) consider heterogeneous beliefs with one receiver.[9] It is my understanding that the current study is the first to analyze these two settings jointly.

Our approach is supported by *rationalizability*. Although our solutions in general would also be supported by equilibrium concepts which admit heterogeneous beliefs, we use rationalizability for three reasons. First, while the standard notion of perfect Bayesian equilibrium uses the common prior, rationalizability naturally admits heterogeneous beliefs.[10] Second, for the Bayesian persuasion games where the stage games involve no strategic interaction (which we call *unlinked*), the identification of a best response for each player directly corresponds to rationalizability. Third, for the Bayesian persuasion games with strategic interactions in the stage games (which we call *linked*), rationalizability may have a sharper prediction due to forward induction when the sender also plays the stage game. We demonstrate this by means of an example.

Our contribution to *unlinked* Bayesian persuasion games has two parts. First, after providing some examples (Section 3), we show that the sender's optimization problem for unlinked Bayesian persuasion games can be seen as a set of *linear programming problems* to the sender (Section 4).[11] We take the examination of the receivers' best responses as "constraints" for the sender's optimization problem. The sender's expected payoff is linear in signal distribution. Although the receivers use their interim beliefs to identify their best responses, these constraints can be modified as linear expressions (with weak inequalities) in the signal distribution.

Second, we show several results (generic properties) based on the observations we extract from the examples shown in Section 3 (Section 5). The

[9] Alonso and Câmara (2016a) focus on language-invariant Bayesian perfect equilibrium: the receiver's behavior only depends on her interim belief regarding the states. This allows the sender to focus on the receiver's interim belief regarding the states.

[10] Perfect Bayesian equilibrium does not have to use the common prior. Battigalli (1999) discusses the relationship between weak rationalizability and weak perfect Bayesian equilibrium.

[11] The adoption of the linear programing approach is not as straightforward as it may sound. In unlinked games, we first specify the receivers' strategies which depend on signal realization, and yet the sender can still choose a message with which some signals will never be realized. This means that the constraints are not closed. See Section 4 for details.

first two observations concern the receivers. We show (i) that for any two constraints (corresponding to two signals), one implies the other generically, and (ii) the results concerning the receivers' multiple best responses, including Remark 1. The last result concerns the sender's rationalizable message. We show (i) that the rationalizable message is generically unique, and therefore (ii) that it is generically the case that the sender strictly prefers either "doing nothing" or "persuading receivers".

Communications in *linked* Bayesian persuasion games are richer and more complex. Although we cannot directly apply the linear programming approach to linked Bayesian persuasion games, we can still analyze any Bayesian persuasion game. We highlight two aspects of communication which cannot be observed in unlinked Bayesian persuasion games, for which we will provide examples (Section 6).

First, we privately talk about others, creating asymmetric information. In unlinked Bayesian persuasion games, the sender does not have an incentive to share with any receiver the information regarding the signals the other receivers observe since doing so would reveal more information regarding the states.[12] In linked Bayesian persuasion games, the sender may find it optimal to do so since such information would also reveal actions taken by others. The sender, however, would have to choose how much information to share with each receiver. The first example shows that signal realization may contain information regarding other receivers' signals to encourage (dis)coordination among the receivers.

Second, we often have tacit understandings without explicit communication, implying that "silence" may have some meanings. However, this cannot happen in unlinked Bayesian persuasion games, since "silence" (doing nothing) simply implies that the sender cannot influence the receivers' behavior. The second example, where the sender also plays the stage game as a receiver, suggests that "silence" may nevertheless persuade the receivers successfully. This is due to forward induction implied by rationalizability.[13]

[12]Even if it is optimal to reveal the state to each receiver, the sender does not have to use the public message.

[13]Heifetz, Meier, and Schipper (2013) extended the notion of extensive-form rationaliz-

There are several other recent studies on information disclosure. Brocas and Carrillo (2007) consider a model with two states where instead of choosing signal distributions conditional on the states, the sender chooses the number of times (binary) signals are revealed given the pre-fixed signal distribution conditional on the states. Rayo and Segal (2010) consider a model where the set of actions is binary; "accept" or "reject". The sender chooses the distribution of signals conditional on the states (prospects). If the receiver accepts, players' payoffs depend on the state. If the receiver rejects, she receives the realization of her uniformly distributed reservation payoff (she knows the value when she chooses her action) while the sender's payoff is pre-determined. Ostrovsky and Schwarz (2010) consider job matchings between students and potential employers. It is the school who knows the students' types and controls the information revelation to the potential employers. Hörner and Skrzypacz (2014) analyze the model with multiple rounds of persuasion stages. Kolotilin, Li, Mylovanov, and Zapechelnyuk (2015) study persuasion when (i) the receiver privately observes her type, and (ii) the receiver's payoff is linear in the state and the type. Bergemann and Morris (2016) consider correlated equilibrium for games with incomplete information under a common prior. Bergemann and Morris (2016) show that their approach can analyze Bayesian persuasion games with multiple receivers. The decision rule (mediator) which recommends actions would correspond to the message in the current paper.[14]

ability with the presence of asymmetric unawareness. Their example (in a different setting) showed that without asymmetric awareness, not taking certain action (silence) led to the unique outcome due to forward induction.

[14] Tamura (2014) assumes (i) that the receiver's action is continuous, (ii) that the receiver's utility function is strictly concave, and (iii) that both the sender's and the receiver's preferences are quadratic. With these assumptions, the sender cannot affect the expected value of the receiver's optimal action, and the sender's expected utility is linear in the variance-covariance matrix of the (multi-dimensional) states. In this setting, Tamura (2014) uses the semidefinite programming approach. Kolotilin (2016) assumes (i) both the sender and the receiver have their own types with supports that are compact (and the sender's type is realized only after her message is sent), (ii) the receiver's action space is binary, (iii) one action leads to a payoff of zero (normalization) for both the sender and the receiver, and (iv) the single-crossing assumption for the receiver's preferences from the other action; there exists a threshold of her private information with which her expected payoff is zero. The sender only needs to make sure that the receiver's expected payoff from the other action is zero. See also Chan, Li, and Wang (2016), Lipnowski and Mathevet (2015) and Taneva

2 Preliminaries

Our formal treatment is geared towards unlinked Bayesian persuasion games; i.e., no strategic interaction at the stage game (the game which the receivers play after signal realization). We have one sender. Let N be the finite set of the receivers with $|N| = n \geq 1$. The finite set of states is Θ with θ being a typical element. Let p_S^0 be the sender's (commonly known) initial belief over Θ. Likewise, for each receiver $i \in N$, let p_i^0 be her (commonly known) initial belief over Θ. We assume that $p_S^0(\theta) > 0$ for each $\theta \in \Theta$ and $p_i^0(\theta) > 0$ for each $\theta \in \Theta$ and $i \in N$. We allow heterogeneous beliefs.

A Bayesian persuasion game has two stages:

First Stage. Let Ξ_i be the finite set of *signals* for each receiver $i \in N$ with ξ_i being a typical element and $\Xi = \prod_{j \in N} \Xi_j$. The sender chooses a *message*, a joint distribution over Ξ conditional on Θ, which we denote π. We assume that the sender's message is such that for every $i \in N$ and $\xi_i \in \Xi_i$, there exists $\theta \in \Theta$ such that $\pi(\xi_i \mid \theta) > 0$. That is, each signal is realizable with a positive probability. Let Π be the set of messages. The message itself is observable to the receivers. As an extreme case, π can be such that the signals of all receivers are perfectly correlated, which we interpret as the case where signals are publicly observable.

Second Stage. Given a message π, the marginal $\pi(\xi_i \mid \theta) = \sum_{j \neq i} \pi((\xi_i, \xi_j) \mid \theta)$ for each $\xi_i \in \Xi_i$ and $\theta \in \Theta$ is computed for each receiver $i \in N$. After observing π and ξ_i, each receiver $i \in N$ revises her belief regarding each $\theta \in \Theta$ via Bayes' rule:

$$p_i^\pi(\theta \mid \xi_i) = \frac{\pi(\xi_i \mid \theta) p_i^0(\theta)}{\sum_{\tilde{\theta} \in \Theta} \pi(\xi_i \mid \tilde{\theta}) p_i^0(\tilde{\theta})}. \tag{1}$$

A message is called *null for player* $i \in N$ if it induces the interim belief identical to the initial belief for each signal; i.e., Ξ_i is singleton, or for each $\xi_i \in \Xi_i$,

(2016) for other applications.

$\pi(\xi_i \mid \theta') = \pi(\xi_i \mid \theta'')$ for each $\theta', \theta'' \in \Theta$. If a message is null for each player $i \in N$, we simply say that *the message is null*. Let π^0 be a null message.

After privately observing ξ_i before playing the stage game, each receiver $i \in N$ chooses her action with her *interim* (updated) belief, $p_i^\pi(\cdot \mid \xi_i)$. Let A_i be the finite set of actions for each receiver $i \in N$ in the stage game with a_i being a typical element and $A = \prod_{j \in N} A_j$.[15] Since each signal itself has no meaning, we assume without loss of generality that $|\Xi_i| = |A_i|$ for each $i \in N$; i.e., the sender recommends an action to each receiver. For the same reason, we do not distinguish the permutations of Ξ_i.[16] Let $s_i : \Pi \times \Xi_i \to A_i$ be a pure strategy, S_i be the set of pure strategies for each player $i \in N$ in the stage game, and $S = \prod_{j \in N} S_j$. Let $s_i(\pi, \xi_i) \in A_i$ be receiver i's action after observing $\xi_i \in \Xi_i$ given π.

Remark 2 *With $|\Xi_i| = |A_i|$ for each $i \in N$, a rationalizable message may be such that two different signals induce the same action for some receiver. In such cases, we can reduce Ξ_i by eliminating "redundant" signals; the corresponding discussion can be found in Subsection 4.3. Most of our results in Section 5 use such reduced sets of signals.*

The sender's payoff function takes the form of $u_S(a, \theta)$; the sender's payoff depends on receivers' action profile and the state. For unlinked Bayesian persuasion games, we have $u_i(a_i, \theta)$ for each receiver $i \in N$; each receiver i's payoff depends only on her own action and the state. This will be relaxed for linked Bayesian persuasion games.

The solution concept we use is Δ-*rationalizability* by Battigalli and Siniscalchi (2003), where Δ corresponds to the set of restrictions on players' (first-order) beliefs. In the current setting, Δ corresponds to p_S^0 for the sender and p_i^0 for each receiver $i \in N$. The procedure for unlinked Bayesian persuasion games goes analogous to backward induction:

[15] Our focus is on static stage games. The receivers may play dynamic games after observing signals.

[16] Kamenica and Gentzkow (2011) call such messages (with $\Xi_1 \subseteq A_1$) *straightforward*. The sender may not use some signals, which we will discuss later.

1. Given that Ξ_i and A_i are finite for each $i \in N$, for each strategy profile s, (i) identify the set of messages to which s_i is a best response for each $i \in N$ (this set could be empty for some s as demonstrated in Section 3), and (ii) within this set, identify the message leading to the highest expected payoff for the sender.

2. Among the messages identified above, identify a best response for the sender.

We will discuss in Section 4 how the procedure is linked to the linear programming approach for unlinked Bayesian persuasion games.

3 Examples

In this section, we demonstrate how the linear programming approach works. We first use the example from Kamenica and Gentzkow (2011) and its variation. We then provide an example with two receivers.

3.1 Examples – One Receiver

We first use the example from Kamenica and Gentzkow (2011) to demonstrate that our approach chooses the same prediction. In the modified example, we demonstrate that predictions could be substantially different depending on players' heterogeneous beliefs. We first discuss the receiver's behavior, which is shared by the two examples. Differences in the two examples are (i) the sender's preferences and (ii) heterogeneity in beliefs.

Receiver 1 has two actions, $A_1 = \{a_1', a_1''\}$. There are two states, $\Theta = \{\theta', \theta''\}$. Let p_S^0 and p_1^0 be the prior beliefs for the sender and receiver 1 respectively. There are two signals $\Xi_1 = \{\xi_1', \xi_1''\}$ and the sender chooses $\pi \in \Pi$. We assume that $u_1(a_1', \theta') = u_1(a_1'', \theta'') = 1$ and $u_1(a_1', \theta'') = u_1(a_1'', \theta') = 0$. After observing π and $\xi_1 \in \Xi_1$, receiver 1's expected payoffs are

$$\left\{ \begin{array}{c} p_1^\pi(\theta' \mid \xi_1)u_1(a_1', \theta') + p_1^\pi(\theta'' \mid \xi_1)u_1(a_1', \theta'') = p_1^\pi(\theta' \mid \xi_1) \\ p_1^\pi(\theta' \mid \xi_1)u_1(a_1'', \theta') + p_1^\pi(\theta'' \mid \xi_1)u_1(a_1'', \theta'') = p_1^\pi(\theta'' \mid \xi_1) \end{array} \right\} \text{ for } \left\{ \begin{array}{c} a_1' \\ a_1'' \end{array} \right\}.$$

Given $\pi(\xi_1 \mid \theta) > 0$ for each $\xi_1 \in \Xi_1$ and $\theta \in \Theta$, we consider three pure strategies for receiver 1 as her best responses:

(a) $a_1' \in A_1$ for each $\xi_1 \in \Xi_1$, which requires $p_1^\pi(\theta' \mid \xi_1) \geq p_1^\pi(\theta'' \mid \xi_1)$ for each $\xi_1 \in \Xi_1$, or

$$\frac{\pi(\xi_1|\theta')p_1^0(\theta')}{\pi(\xi_1|\theta')p_1^0(\theta')+\pi(\xi_1|\theta'')[1-p_1^0(\theta')]} \geq \frac{\pi(\xi_1|\theta'')[1-p_1^0(\theta')]}{\pi(\xi_1|\theta')p_1^0(\theta')+\pi(\xi_1|\theta'')[1-p_1^0(\theta')]}$$

or

$$\pi(\xi_1 \mid \theta'') \leq \frac{p_1^0(\theta')}{1-p_1^0(\theta')}\pi(\xi_1 \mid \theta')$$

for each $\xi_1 \in \Xi_1$.

(b) $a_1'' \in A_1$ for each $\xi_1 \in \Xi_1$, which requires $p_1^\pi(\theta' \mid \xi_1) \leq p_1^\pi(\theta'' \mid \xi_1)$ for each $\xi_1 \in \Xi_1$, or

$$\pi(\xi_1 \mid \theta'') \geq \frac{p_1^0(\theta')}{1-p_1^0(\theta')}\pi(\xi_1 \mid \theta')$$

for each $\xi_1 \in \Xi_1$.

(c) a_1' for ξ_1' and a_1'' for ξ_1'', which requires $p_1^\pi(\theta' \mid \xi_1') \geq p_1^\pi(\theta'' \mid \xi_1')$ and $p_1^\pi(\theta' \mid \xi_1'') \leq p_1^\pi(\theta'' \mid \xi_1'')$, or

$$\pi(\xi_1' \mid \theta'') \leq \frac{p_1^0(\theta')}{1-p_1^0(\theta')}\pi(\xi_1' \mid \theta') \quad \text{and} \quad \pi(\xi_1'' \mid \theta'') \geq \frac{p_1^0(\theta')}{1-p_1^0(\theta')}\pi(\xi_1'' \mid \theta').$$

These relationships are visualized in Figure 2. For each case above, there are three possibilities (from left to right in Figure 2): $p_1^0(\theta') < \frac{1}{2}$, $p_1^0(\theta') = \frac{1}{2}$, and $p_1^0(\theta') > \frac{1}{2}$. Note (i) that a point in the box diagram uniquely identifies π, and (ii) that the cases above do not consider the origins. Figure 2 shows that while (c) above is possible for any value of $p_1^0(\theta')$, (a) is possible only for $p_1^0(\theta') \geq \frac{1}{2}$ and (b) is possible only for $p_1^0(\theta') \leq \frac{1}{2}$.

We now discuss the sender's behavior in two examples.

Example 1. In the original example from Kamenica and Gentzkow (2011),

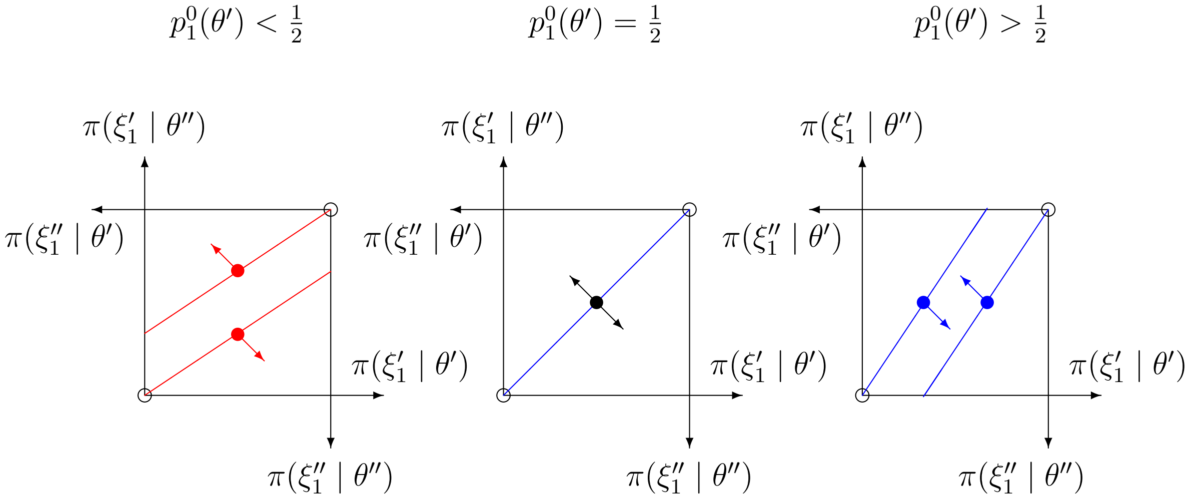

(a) a_1' for each $\xi_1 \in \Xi_1$

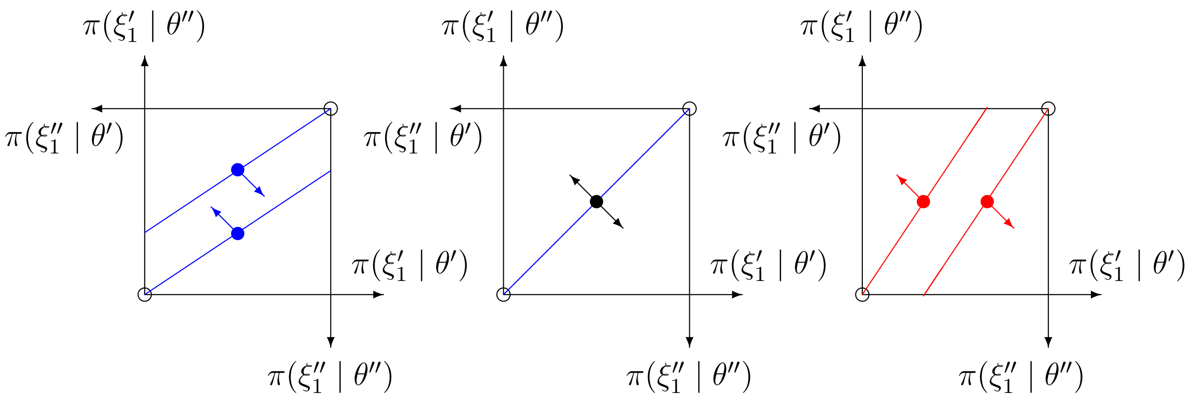

(b) a_1'' for each $\xi_1 \in \Xi_1$

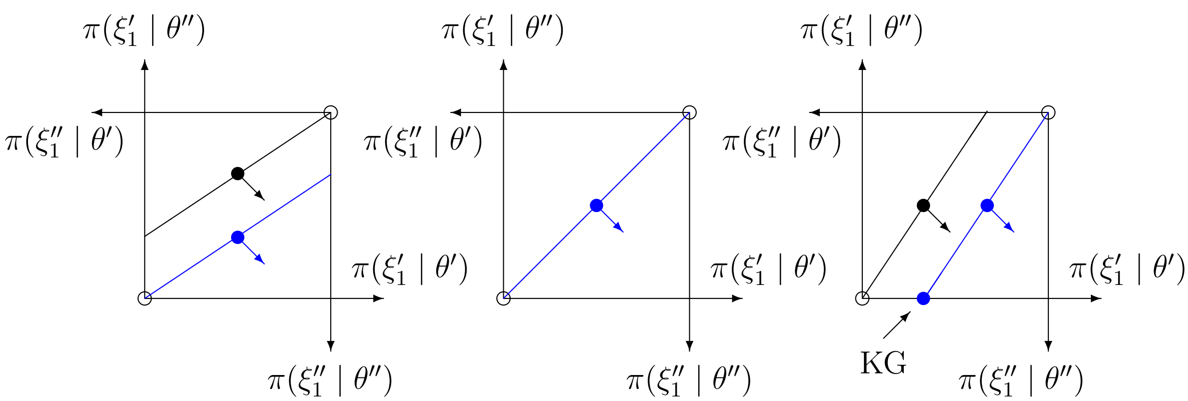

(c) a_1' for ξ_1' and a_1'' for ξ_1''

Figure 2: Receiver 1's best responses

the sender is the prosecutor and receiver 1 is the judge; $\Theta = \{innocent, guilty\}$, $A_1 = \{acquit, convict\}$, and $\Xi_1 = \{innocent, guilty\}$. Let $\theta' = \xi'_1 = innocent$, $\theta'' = \xi''_1 = guilty$, $a'_1 = acquit$ and $a''_1 = convict$. Let $p^0_S(\theta') = p^0_1(\theta') = 0.7$ (and hence receiver 1 chooses a'_1 without persuasion). The sender's payoffs are such that $u_S(a'_1, \theta) = 0$ and $u_S(a''_1, \theta) = 1$ for each $\theta \in \Theta$.

The sender avoids receiver 1's strategy "a'_1 for each $\xi_1 \in \Xi_1$" (Figure 2 (a)). Given $p^0_1(\theta') > \frac{1}{2}$, it is not possible for the sender to make receiver 1 choose a''_1 for each $\xi_1 \in \Xi_1$ (Figure 2 (b)). The remaining possibility is that receiver 1 chooses a'_1 for ξ'_1 and a''_1 for ξ''_1 (Figure 2 (c)). The sender's expected payoff is higher if the chance of ξ''_1 is higher; i.e, towards the southwest. The sender therefore chooses $\pi(\xi''_1 \mid \theta') = \frac{1-0.7}{0.7} = \frac{3}{7}$ and $\pi(\xi''_1 \mid \theta'') = 1$, consistent with the solution in Kamenica and Gentzkow (2011).

Two observations are highlighted by Kamenica and Gentzkow (2011). First, if receiver 1 observes ξ'_1, corresponding to a'_1 which is the *worst action* from the sender's point of view, she knows that the state is θ'. This observation is reflected in Proposition 4 of Kamenica and Gentzkow (2011). We will discuss this for the case of multiple receivers in Example 3. Second, when receiver 1 observes ξ''_1, she is indifferent between two actions; the constraint for ξ''_1 binds. Proposition 5 of Kamenica and Gentzkow (2011) provides the characterization of the receiver's interim belief for such cases. We will discuss the receivers' multiple best responses in Subsection 5.2.

Example 2. The sender's payoffs are such that $u_S(a'_1, \theta'') = u_S(a''_1, \theta') = 1$ and $u_S(a'_1, \theta') = u_S(a''_1, \theta'') = 0$; two players have completely opposite preferences. This corresponds to Example 1 by letting $\theta' = a'_1 = A$, $\theta'' = a''_1 = B$, $\xi'_1 = a$, and $\xi''_1 = b$. Assume that $p^0_1(\theta') < \frac{1}{2}$. Receiver 1 would choose a''_1 without persuasion. It is not possible for the sender to induce receiver 1 to choose a'_1 independent of a signal (Figure 2 (a)). There are two other possibilities: (i) a''_1 independent of a signal (Figure 2 (b)), and (ii) a'_1 for ξ'_1 and a''_1 for ξ''_1 (Figure 2 (c)). Note that the sender's expected payoff with the former is constant at $p^0_S(\theta')$. By looking at different values of $p^0_S(\theta')$, we compare these strategies for receiver 1 and identify the rationalizable message for the sender.

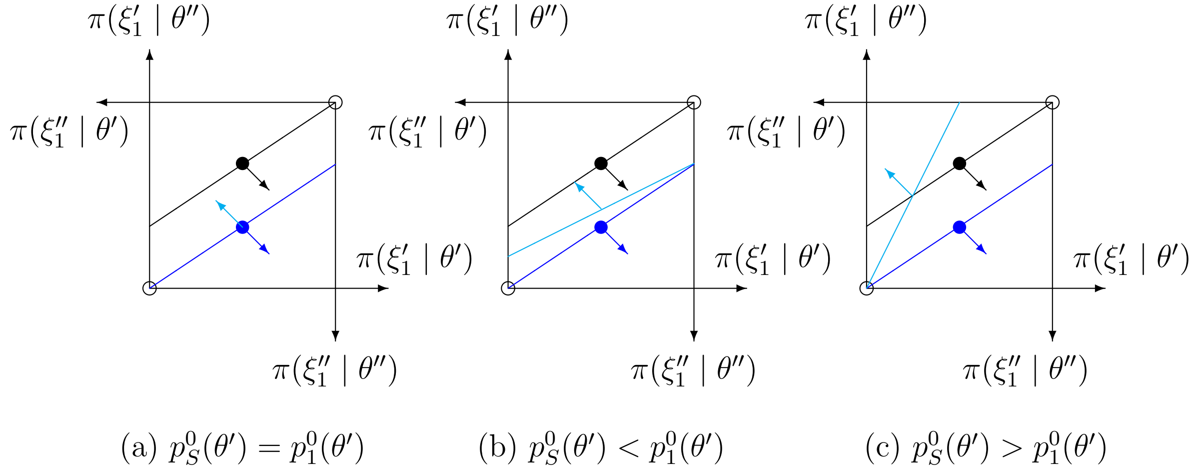

Figure 3: Heterogeneous Beliefs; a_1' for ξ_1' and a_1'' for ξ_1''

With receiver 1's strategy "a_1' for ξ_1' and a_1'' for ξ_1''" as the constraint, the sender chooses π to maximize

$$\pi(\xi_1' \mid \theta'')p_S^0(\theta'') + \pi(\xi_1'' \mid \theta')p_S^0(\theta') \;\;\Leftrightarrow\;\; \pi(\xi_1' \mid \theta'')[1 - p_S^0(\theta')] + [1 - \pi(\xi_1' \mid \theta')]p_S^0(\theta').$$

Note (i) that the sender's expected payoff is higher if both $\pi(\xi_1' \mid \theta'')$ and $\pi(\xi_1'' \mid \theta')$ increase (i.e., moving towards the northwest), and (ii) the slope of the sender's "indifference curve" is $\frac{d\pi(\xi_1'|\theta'')}{d\pi(\xi_1'|\theta')} = \frac{p_S^0(\theta')}{1-p_S^0(\theta')}$.

Three possibilities regarding the sender's prior are visualized in Figure 3:

1. $p_S^0(\theta') = p_1^0(\theta')$ or $\frac{p_S^0(\theta')}{1-p_S^0(\theta')} = \frac{p_1^0(\theta')}{1-p_1^0(\theta')}$: The slope of the indifference curve and the slope of the constraints coincide. There is a continuum of solutions. Given $\pi(\xi_1' \mid \theta'') = \frac{p_S^0(\theta')}{1-p_S^0(\theta')}\pi(\xi_1' \mid \theta')$, the sender's expected payoff is $p_S^0(\theta')$. The sender is indifferent between two messages – even if the sender attempts to influence receiver 1's behavior, the expected payoff remains the same.

2. $p_S^0(\theta') < p_1^0(\theta')$ or $\frac{p_S^0(\theta')}{1-p_S^0(\theta')} < \frac{p_1^0(\theta')}{1-p_1^0(\theta')}$: The sender's indifference curve is flatter than the constraints. Since receiver 1 overestimates the possibility of θ', if θ' is the true state, the sender lets receiver 1 know it, $\pi(\xi_1' \mid \theta') = 1$. In return, the sender makes sure that $\pi(\xi_1' \mid \theta'')$ takes the highest

possible value, $\pi(\xi_1' \mid \theta'') = \frac{p_1^0(\theta')}{1-p_1^0(\theta')}$. The expected payoff for the sender is $\frac{p_1^0(\theta')[1-p_S^0(\theta')]}{1-p_1^0(\theta')}$. Note

$$\frac{p_1^0(\theta')[1-p_S^0(\theta')]}{1-p_1^0(\theta')} - p_S^0(\theta') = \frac{p_1^0(\theta')-p_S^0(\theta')}{1-p_1^0(\theta')} > 0$$

and therefore the sender induces receiver 1 to choose "a_1' for ξ_1' and a_1'' for ξ_1''" by choosing the message described above. Note that the constraint for ξ_1' binds.

3. $p_S^0(\theta') > p_1^0(\theta')$ or $\frac{p_S^0(\theta')}{1-p_S^0(\theta')} > \frac{p_1^0(\theta')}{1-p_1^0(\theta')}$: The sender's expected payoff is strictly less than $p_S^0(\theta')$ if π satisfies the constraints (note that the origin is excluded). Hence, the sender induces receiver 1 to choose "a_1'' independent of signal". The sender achieves this with the null message (i.e., no persuasion).

3.2 Example – Two Receivers

In the following example, we have two receivers. Remember that the original example from Kamenica and Gentzkow (2011) shows that if the receiver chooses the worst *action* (from the sender's point of view), she knows the state.[17] This is reflected in Proposition 4 in Kamenica and Gentzkow (2011). The following example shows that for the case of multiple receivers and heterogeneous beliefs, even if the sender lets the receivers choose the worst *action profile*, the receivers may not know the state.

Example 3. There are two possible states $\Theta = \{A, B\}$. Receivers 1 and 2 have two actions $A_i = \{a, b\}$ for each $i \in \{1, 2\}$. The payoffs are summarized in (a) of Figure 4 where the first payoff in each cell is for the sender while the second and third payoffs are for receivers 1 and 2 respectively. Receiver 1's utility is one if she matches her action to the state, a if A or b if B, and zero otherwise. Receiver 2's utility is one if she avoids the match, a if B and b if A,

[17] Strictly speaking, the sender may not know the state if there are multiple states where the corresponding (worst) action is the unique best response for the receiver.

A		2: a	2: b
1	a	1, 1, 0	0, 1, 1
1	b	$-\varepsilon$, 0, 0	1, 0, 1

B		2: a	2: b
1	a	1, 0, 1	0, 0, 0
1	b	$-\varepsilon$, 1, 1	1, 1, 0

(a) Payoffs

A	$\pi(\xi_2' \mid A)$	$\pi(\xi_2'' \mid A)$
$\pi(\xi_1' \mid A)$	w_a	x_a
$\pi(\xi_1'' \mid A)$	y_a	z_a

B	$\pi(\xi_2' \mid B)$	$\pi(\xi_2'' \mid B)$
$\pi(\xi_1' \mid B)$	w_b	x_b
$\pi(\xi_1'' \mid B)$	y_b	z_b

(b) Conditional Probabilities

Figure 4: Multiple Receivers: Example

and zero otherwise. Note that there is no strategic interaction between them. The sender's utility is one if receivers 1 and 2 choose the same action, and zero if they choose (a, b) and $-\varepsilon$ if they choose (b, a). We assume that $\varepsilon > 0$ is an arbitrary small number. Note (i) that the sender's payoffs are state-independent, and (ii) that (b, a) is the worst action profile from the sender's point of view.

Let $p_S^0(A) = p_S$ and $p_i^0(A) = p_i$ for each $i \in \{1, 2\}$. Receivers 1 and 2 will choose (a, a) if $p_1 > \frac{1}{2} > p_2$ or (b, b) if $p_1 < \frac{1}{2} < p_2$, in which case the sender chooses a null message. If $\min\{p_1, p_2\} > \frac{1}{2}$ or $\frac{1}{2} > \max\{p_1, p_2\}$, the sender attempts to change their beliefs so that there would be a chance that they will coordinate. We focus on the former, i.e., $\min\{p_1, p_2\} > \frac{1}{2}$. Without persuasion, receiver 1 chooses a while receiver 2 chooses b.

Let $\Xi_i = \{\xi_i', \xi_i''\}$ for each $i \in \{1,2\}$. Consider the following strategy profile for receivers 1 and 2: each receiver $i \in \{1,2\}$ chooses either a after observing ξ_i' or b after observing ξ_i''. The required conditions are

$$\begin{array}{ccc} \frac{\pi(\xi_1'|A)p_1}{\pi(\xi_1'|A)p_1+\pi(\xi_1'|B)(1-p_1)} & \geq & \frac{\pi(\xi_1'|B)(1-p_1)}{\pi(\xi_1'|A)p_1+\pi(\xi_1'|B)(1-p_1)} \\ \frac{\pi(\xi_1''|A)p_1}{\pi(\xi_1''|A)p_1+\pi(\xi_1''|B)(1-p_1)} & \leq & \frac{\pi(\xi_1''|B)(1-p_1)}{\pi(\xi_1''|A)p_1+\pi(\xi_1''|B)(1-p_1)} \end{array}$$

or

$$\pi(\xi_1' \mid B) \leq \tfrac{p_1}{1-p_1}\pi(\xi_1' \mid A) \quad \text{and} \quad \pi(\xi_1'' \mid B) \geq \tfrac{p_1}{1-p_1}\pi(\xi_1'' \mid A) \tag{2}$$

for receiver 1 and

$$\begin{array}{ccc} \frac{\pi(\xi_2'|B)(1-p_2)}{\pi(\xi_2'|A)p_2+\pi(\xi_2'|B)(1-p_2)} & \geq & \frac{\pi(\xi_2'|A)p_2}{\pi(\xi_2'|A)p_2+\pi(\xi_2'|B)(1-p_2)} \\ \frac{\pi(\xi_2''|B)(1-p_2)}{\pi(\xi_2''|A)p_2+\pi(\xi_2''|B)(1-p_2)} & \leq & \frac{\pi(\xi_2''|A)p_2}{\pi(\xi_2''|A)p_2+\pi(\xi_2''|B)(1-p_2)} \end{array}$$

or

$$\pi(\xi_2' \mid B) \geq \tfrac{p_2}{1-p_2}\pi(\xi_2' \mid A) \quad \text{and} \quad \pi(\xi_2'' \mid B) \leq \tfrac{p_2}{1-p_2}\pi(\xi_2'' \mid A) \tag{3}$$

for receiver 2. The notation in (b) of Figure 4 modifies (2) and (3) as

$$(w_b + x_b) \leq \tfrac{p_1}{1-p_1}(w_a + x_a) \quad \text{and} \quad \tfrac{p_1}{1-p_1}(y_a + z_a) \leq (y_b + z_b) \tag{4}$$

$$\tfrac{p_2}{1-p_2}(w_a + y_a) \leq (w_b + y_b) \quad \text{and} \quad (x_b + z_b) \leq \tfrac{p_2}{1-p_2}(x_a + z_a). \tag{5}$$

The sender's expected payoff given the strategy profile shown above is

$$\begin{array}{rl} & (w_a + z_a)p_S + (w_b + z_b)(1 - p_S) - \varepsilon(y_a p_S + y_b(1 - p_S)) \\ = & (w_a + z_a - \varepsilon y_a)p_S + (w_b + z_b - y_b)(1 - p_S) \end{array}$$

The sender maximizes this expression subject to (4), (5), and the following constraints:

$$w_a + x_a + y_a + z_a = 1 \quad \text{and} \quad w_b + x_b + y_b + z_b = 1$$
$$w_a, x_a, y_a, z_a \geq 0 \quad \text{and} \quad w_b, x_b, y_b, z_b \geq 0$$

Suppose $p_S = 0.8$, $p_1 = 0.7$, $p_2 = 0.6$, and $\varepsilon = 0.01$. The solution is $(w_a, x_a, y_a, z_a) = (\frac{2}{3}, 0, 0, \frac{1}{3})$ and $(w_b, x_b, y_b, z_b) = (\frac{2}{9}, 0, \frac{7}{9}, 0)$. Note that since $y_b > 0$, it is possible that receiver 1 chooses b (after observing ξ_1'') and receiver 2 chooses a (after observing ξ_2') when the state is B, leading to the worst payoff for the sender, $-\varepsilon$. The receivers always coordinate if the state is A $(w_a + z_a = 1)$ while the chance that they will coordinate when the state is B is $\frac{2}{9}$ $(w_b + z_b = \frac{2}{9})$. The sender's expected payoff is $\frac{8}{10} + \frac{2}{10}[\frac{2}{9} + \frac{7}{9}(-\varepsilon)] = \frac{1}{45}(38 - 7\varepsilon) \approx 0.84$.[18]

We have

$$(\pi(\xi_1' \mid A), \pi(\xi_1'' \mid A)) = (\tfrac{2}{3}, \tfrac{1}{3}) \qquad (\pi(\xi_1' \mid B), \pi(\xi_1'' \mid B)) = (\tfrac{2}{9}, \tfrac{7}{9})$$
$$(\pi(\xi_2' \mid A), \pi(\xi_2'' \mid A)) = (\tfrac{2}{3}, \tfrac{1}{3}) \qquad (\pi(\xi_2' \mid B), \pi(\xi_2'' \mid B)) = (1, 0).$$

Two observations. First, receiver 1 never knows the state. Second, receiver 2 knows the state if she observes ξ_2'' (the state is A). Unlike what Proposition 4 in Kamenica and Gentzkow (2011) suggests, when they choose the worst action profile (b, a), the state is B and they do not know this. It can also be shown that the constraints for ξ_1'' and ξ_2' bind. Note that by letting respectively (i) θ' and θ'' correspond to two states "student 1 (2) is L (H)" and "student 1 (2) is H (L), (ii) ξ_i' and ξ_i'' be "fail" and "pass" for each $i \in N$, and (iii) a and b be "high effort" and "low effort", the example corresponds to the multiple-receiver example we discussed in the Introduction.[19]

[18] Persuading only one receiver is an option for the sender. Appendix A shows that the corresponding expected payoffs for the sender do not exceed her expected payoff shown above.

[19] We let $\varepsilon = 0$ in the Introduction with which the solution remains the same.

3.3 Observations

We point out several observations from the previous examples, which we will generalize in Section 5.

1. In the box diagrams, when each signal corresponds to a different action, not only are the two constraints parallel to each other, but also one implies the other except in the case of $p_1^0(\theta') = \frac{1}{2}$ in Figure 2. We will show that, by taking into account the probabilities π assigns to the two signals, this relationship holds generically.

2. Binding constraints mean that the receivers have multiple best responses to some signals. As noted in Remark 1, we investigate what would happen if the receivers choose different actions under the presence of multiple best responses. We show that the result concerning the first observation above implies that the sender does not have the corresponding rationalizable message, and hence we do not have to consider this type of possibility for our approach.

3. Rationalizable messages are unique in these examples and the sender prefers either "doing nothing" or "persuading receivers" – Example 2 (with common prior) is an exception. We show that both hold generically.[20]

4 Linear Programming Approach

In this section, we discuss the linear programming approach for unlinked Bayesian persuasion games. The approach has two steps as we discussed above, confirming that the linear programming approach is supported by rationalizability.

[20]This result depends on the assertion that the map from the set of signals (which are assigned positive probabilities) to the set of actions for each receiver is bijective. If the sender uses multiple signals for the same action (which is possible but unnecessary), the result does not hold. See Subsection 4.3.

4.1 Second Stage

The second stage only concerns the receivers' decision problems. Given $\pi \in \Pi$ and $\xi_i \in \Xi_i$ for each $i \in N$, let $s_i^+(\pi, \xi_i)$ be such that

$$\sum_{\theta \in \Theta} u_i(s_i^+(\pi, \xi_i), \theta) p_i^\pi(\theta \mid \xi_i) \geq \sum_{\theta \in \Theta} u_i(a_i, \theta) p_i^\pi(\theta \mid \xi_i)$$

or

$$\sum_{\theta \in \Theta} \left[u_i(s_i^+(\pi, \xi_i), \theta) - u_i(a_i, \theta) \right] \pi(\xi_i \mid \theta) p_i^0(\theta) \geq 0 \tag{6}$$

for each $a_i \in A_i$. Since we assume that there does not exist $\xi_i \in \Xi_i$ such that $\pi(\xi_i \mid \theta) = 0$ for each $i \in N$ and $\theta \in \Theta$, $\sum_{\theta \in \Theta} \pi(\xi_i \mid \theta) p_i^0(\theta) > 0$. Let s_i^+ be a corresponding best response for receiver $i \in N$.

As the examples in Section 3 illustrate, different messages can lead to the same strategy for each receiver. Define $t_i : \Xi_i \to A_i$, $t = (t_1, \ldots, t_n)$ and T be the collection of t's. Define an equivalence class as

$$[t] = \left\{ s^+ \in S \mid \begin{array}{c} \text{there exists } \pi \text{ such that } t_i(\xi_i) = s_i^+(\pi, \xi_i) \\ \text{for each } i \in N \text{ and } \xi_i \in \Xi_i \end{array} \right\}.$$

Note that each strategy in the same equivalence class has the same payoff implication for all players.

As demonstrated in Section 3 (e.g., (a) of Figure 2 when $p_1^0(\theta') < \frac{1}{2}$), some equivalence classes can be empty. One may wonder whether it is possible that all equivalence classes are empty, in which case the sender's optimization problem is not well defined. We now show that there always exists a non-empty set.

Lemma 1 *There exists a non-empty equivalence class.*

Proof. Suppose that the sender chooses a null message. Since A_i is finite, there exists a best response to receiver i's prior belief for each $i \in N$. The result follows. $\square$

Let $T^+ \subseteq T$ be the collection of non-empty equivalence classes, with t^+ being a typical element. Given $t^+ \in T^+$, let

$$\Pi(t^+) = \left\{ \pi \in \Pi \mid \begin{array}{c} \text{there exists } s_i^+ \in S_i \text{ such that } s_i^+(\pi, \xi_i) = t_i^+(\xi_i) \\ \text{for each } i \in N \text{ and } \xi_i \in \Xi_i \end{array} \right\}.$$

In other words, given $t^+ \in T^+$, if the sender chooses $\pi \in \Pi(t^+)$, for each $i \in N$ and each $\xi_i \in \Xi_i$, receiver i maximizes her expected payoff by choosing $t_i^+(\xi_i)$. Given $t^+ \in T^+$, for each $\pi \in \Pi(t^+)$ and $\xi_i \in \Xi_i$, we treat $s_i^+(\pi, \xi_i)$ and $t_i^+(\xi_i)$ interchangeably.

4.2 First Stage

For each $i \in N$ and $t^+ \in T^+$, although the actions chosen by t_i^+ are fixed (i.e., $t_i^+(\xi_i)$ for each $\xi_i \in \Xi_i$), by changing the corresponding message $\pi \in \Pi(t^+)$ and thus the likelihood of each $\xi_i \in \Xi_i$, the sender can change her expected payoff. Hence, for each $t^+ \in T^+$, the sender can identify $\pi \in \Pi(t^+)$ which maximizes her expected payoff. Remember, however, that we do not consider the "origins" in the box diagrams: there is no ξ_i such that $\pi(\xi_i \mid \theta) = 0$ for each $\theta \in \Theta$. This implies that $\Pi(t^+)$ is not closed, and no solution may exist for the sender (e.g., (c) of Figure 3). For the sender's optimization problem, we instead look at the *closure* of $\Pi(t^+)$, denoted by $\bar{\Pi}(t^+)$. We treat $\bar{\Pi}(t^+)$ as the sender's constraint corresponding to each t^+.

At the first stage, for each $t^+ \in T^+$ with the corresponding $\bar{\Pi}(t^+)$, the sender can compute the set of messages which maximizes her expected payoff:

$$\bar{\Pi}^+(t^+) = \operatorname*{argmax}_{\pi \in \bar{\Pi}(t^+)} \sum_{\theta \in \Theta} \sum_{\xi \in \Xi} u_S(t^+(\xi), \theta) \pi(\xi \mid \theta) p_S^0(\theta). \tag{7}$$

This gives the sender's optimal message with respect to $t^+ \in T^+$. Note that both (6) and (7) are *linear* in π. In other words, this maximization problem can be viewed as a *linear programing problem* given $t^+ \in T^+$. Given $t^+ \in T^+$ and $\bar{\Pi}^+(t^+)$, let $u_S^+(t^+)$ be the corresponding maximized expected payoff for the sender.

As seen in (c) of Figure 3, however, there may exist $t^+ \in T^+$ such that given $\bar{\Pi}^+(t^+)$, the optimal message for the sender corresponds to an origin and therefore is not an element of $\Pi(t^+)$ (i.e., there exists $\xi_i' \in \Xi_i$ such that $\pi(\xi_i' \mid \theta) = 0$ for each $\theta \in \Theta$), suggesting the possibility that no solution exists. The following result shows that there always exists $\hat{t}^+ \in T^+$ which replaces t^+ as seen in Example 2.

Proposition 1 *Given $t^+ \in T^+$, suppose that the solution for the sender, $\pi^+ \in \bar{\Pi}^+(t^+)$, is such that $\pi^+ \notin \Pi(t^+)$. Then, there exist $\hat{t}^+ \in T^+$ such that $u_S^+(\hat{t}^+) = u_S^+(t^+)$ and the corresponding optimal message, $\hat{\pi}^+ \in \bar{\Pi}^+(\hat{t}^+)$, is such that $\hat{\pi}^+ \in \Pi(\hat{t}^+)$.*

See Appendix B for the proof. It is important to note that $\hat{\pi}^+ \in \Pi(\hat{t}^+)$ and hence $\hat{\pi}^+$ is a feasible message. Proposition 1 suggests that if we encounter t^+ with $\pi^+ \in \bar{\Pi}^+(t^+) \backslash \Pi(t^+)$, we can always look for another $\hat{t}^+ \in T^+$ and $\hat{\pi}^+ \in \bar{\Pi}^+(\hat{t}^+) \cap \Pi(\hat{t}^+)$ with $u_S^+(\hat{t}^+) = u_S^+(t^+)$.

The sender can then identify a rationalizable message. The sender's rationalizable message π^* is such that $\pi^* \in \Pi(t^*)$ where

$$t^* \in \underset{t^+ \in T^+}{\operatorname{argmax}}\, u_S^+(t^+). \tag{8}$$

Accordingly, the sender's optimization problem can be viewed as a *set of linear programming problems*. That A_i and Ξ_i for each $i \in N$ are finite implies the following result.

Proposition 2 *A rationalizable message π^* exists.*

4.3 Redundant Signals

As the construction of message used in Proposition 1 suggests, given a rationalizable message for the sender, it is possible that there are multiple signals which induce the same action for some receiver (e.g., a null message with multiple signals as seen in the case of $p_S^0(\theta') > p_1^0(\theta')$ in Example 2). For such cases, the sender can simply reduce the number of signals by removing "redundant" signals.

Proposition 3 *Given the sender's rationalizable message π^*, suppose that there exist receiver $i \in N$ and signals $\tilde{\xi}_i, \hat{\xi}_i \in \Xi_i$ such that $t_i^*(\tilde{\xi}_i) = t_i^*(\hat{\xi}_i)$. Then, there exists $\Xi_i^* \subset \Xi_i$ such that (i) for any $\tilde{\xi}_i, \hat{\xi}_i \in \Xi_i^*$, $t^*(\tilde{\xi}_i) \neq t^*(\hat{\xi}_i)$ and (ii) the sender's maximized expected payoff with respect to $\Xi_i^* \times \Xi_{-i}$ remains the same.*

See Appendix C for the proof. Our focus is on Ξ_i^* for each $i \in N$ for our results (except Proposition 4 and Corollary 1) hereafter. Note that if the sender uses a null message for receiver $i \in N$, Ξ_i^* is singleton.

5 Generalization

In this section, we generalize the observations from the previous examples discussed in Subsection 3.3.

5.1 On Constraints

In the box diagrams of the examples in Subsection 3.1 (except when $p_1^0(\theta') = \frac{1}{2}$ in Figure 2), not only are the two constraints parallel to each other, but also one implies the other when each signal corresponds to a different action. We now show that this is generically true by taking into account the probabilities that the message assigns to the two signals.

Given s^+ and $\pi \in \Pi(s^+)$, consider receiver $i \in N$ with $\tilde{\xi}_i, \hat{\xi}_i \in \Xi_i$ such that $s_i^+(\pi, \tilde{\xi}_i) \neq s_i^+(\pi, \hat{\xi}_i)$.[21] We have

$$\sum_{\theta \in \Theta} \left[u_i(s_i^+(\pi, \tilde{\xi}_i), \theta) - u_i(s_i^+(\pi, \hat{\xi}_i), \theta) \right] \pi(\tilde{\xi}_i \mid \theta) p_i^0(\theta) \geq 0 \tag{9}$$

$$\sum_{\theta \in \Theta} \left[u_i(s_i^+(\pi, \tilde{\xi}_i), \theta) - u_i(s_i^+(\pi, \hat{\xi}_i), \theta) \right] \pi(\hat{\xi}_i \mid \theta) p_i^0(\theta) \leq 0. \tag{10}$$

We say that (9) implies (10) if any $\pi \in \Pi(s^+)$ satisfying (9) also satisfies (10). The definition for "(10) implies (9)" is defined analogously. Let $c(\tilde{\xi}_i, \hat{\xi}_i \mid \theta) =$

[21] Note that we are still looking at Ξ_i, not Ξ_i^*.

$\pi(\tilde{\xi}_i \mid \theta) + \pi(\hat{\xi}_i \mid \theta)$ for each $\theta \in \Theta$. We consider the following expression:

$$\sum_{\theta \in \Theta} \left[u_i(s_i^+(\pi, \tilde{\xi}_i), \theta) - u_i(s_i^+(\pi, \hat{\xi}_i), \theta) \right] c(\tilde{\xi}_i, \hat{\xi}_i \mid \theta) p_i^0(\theta). \tag{11}$$

The following result shows that the relationship between (9) and (10) depends on the value of (11).

Proposition 4 *Given s^+ and $\pi \in \Pi(s^+)$, consider receiver $i \in N$ with $\tilde{\xi}_i, \hat{\xi}_i \in \Xi_i$ such that $s_i^+(\pi, \tilde{\xi}_i) \neq s_i^+(\pi, \hat{\xi}_i)$. Then,*

- *(11) is negative if and only if (9) implies (10),*
- *(11) is positive if and only if (10) implies (9), and*
- *(11) is zero if and only if (9) and (10) hold simultaneously with equalities.*

The proof can be found in Appendix D. That small perturbations in the sender's payoffs make (11) non-zero implies our claim.

Corollary 1 *Given s^+ and $\pi \in \Pi(s^+)$, consider receiver $i \in N$ with $\tilde{\xi}_i, \hat{\xi}_i \in \Xi_i$ such that $s_i^+(\pi, \tilde{\xi}_i) \neq s_i^+(\pi, \hat{\xi}_i)$. Then, generically either (9) implies (10) or (10) implies (9).*

Proposition 4 also has a strong implication for each receiver $i \in N$ with $|\Xi_i^*| = 2$ when the sender chooses a rationalizable message π^*. Let $\Xi_i^* = \{\tilde{\xi}_i, \hat{\xi}_i\}$. This means (i) that there are only two actions receiver i would take, and (ii) that $c(\tilde{\xi}_i, \hat{\xi}_i \mid \theta) = 1$ for each $\theta \in \Theta$. Consider Example 1 again. Note that the constraint for ξ_1'' is the one which binds (Figure 2 (c)), corresponding to a_1'' which receiver 1 would not have taken if there is no belief update. Indeed, for the case of $|\Xi_i^*| = 2$, if one of the two constraints binds, this corresponds to the action which would not have been chosen under the initial belief. In this sense, receiver 1 "changes her mind" after observing $\hat{\xi}_1''$ (corresponding to the binding constraint) and chooses a_1''.

Corollary 2 *Given π^*, consider receiver $i \in N$ with $|\Xi_i^*| = 2$. Only the constraint corresponding to the action which would not have been chosen under p_i^0 binds if and only if (11) is non-zero.*

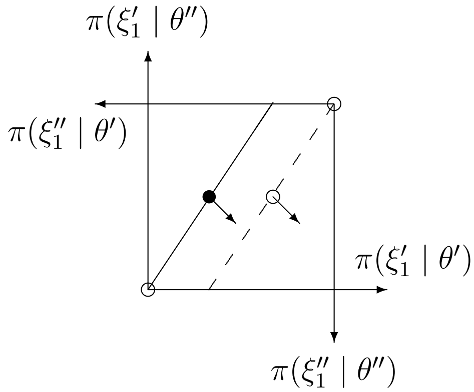

(a) Example 1 Revisited

	a_1'	a_1''	a_1'''
θ'	2, 1	2, 0	0, $\frac{1}{2}$
θ''	1, 0	1, 1	0, $\frac{1}{2}$

(b) Unintended Best Response

Figure 5: Multiple Best Responses

Proof. Suppose without loss of generality that (11) is positive, and hence $s_i^+(\pi^*, \tilde{\xi}_i)$ leads to a higher expected payoff for receiver i compared to $s_i(\pi^*, \hat{\xi}_i)$ under p_i^0. Proposition 4 then says that if one of (9) and (10) binds, it is (10) (and (9) holds with strict inequality), implying the result. $\square$

5.2 Multiple Best Responses for Receivers

The receivers' multiple best responses may arise due to two possibilities:

1. The rationalizable messages needs some binding constraints, and
2. Expected payoffs for the receivers simply take the same value.

Kamenica and Gentzkow (2011) sidestep such considerations by using the notion of *sender-preferred equilibrium* – if there are multiple best responses for the receiver, she will take the one with which the sender receives the highest payoff.

We begin by considering the first possibility. Take again Example 1 in Subsection 3.1. When the receiver observes ξ_1'', both actions are best responses. The solution from Kamenica and Gentzkow (2011) and Example 1 suggest that the receiver chooses a_1'' after observing ξ_1''. Suppose instead that the receiver chooses a_1' after observing ξ_1'' *if she is indifferent* (Remark 1).[22] This is

[22] Note that this is different from another strategy "a_1' for each $\xi_1 \in \Xi_1$".

visualized in (a) of Figure 5. Under this particular strategy for receiver 1, there exists *no* rationalizable message for the sender; while strictly positive expected payoff is achievable for the sender, such messages are all strictly dominated.[23] This is because the sender's expected payoff is not uppersemicontinuous at the points on the boundary of the constraint.[24]

Proposition 5 *Given π^*, consider receiver $i \in N$ with $\tilde{\xi}_i, \hat{\xi}_i \in \Xi_i^*$ such that (9) holds with equality while (10) holds with strict inequality. Consider receiver i's alternative strategy such that after observing $\tilde{\xi}_i$, she chooses $s_i^+(\pi^*, \hat{\xi}_i)$ instead of $s_i^+(\pi^*, \tilde{\xi}_i)$, and everything else remains the same. Then, there exists no rationalizable message with respect to the receiver's strategy profile corresponding to this change.*

Proof. That (9) binds simply means that any other message satisfying (9) with strict inequality is not rationalizable. This leads to the claim. □

We now consider the second possibility. Consider the following example.

Example 4. We have one receiver, three actions $A_1 = \{a_1', a_1'', a_1'''\}$, and two states $\Theta = \{\theta', \theta''\}$. Players' payoffs are shown in (b) of Figure 5 where for each cell, the first element is the sender's payoff and the second element is receiver 1's payoff. If receiver 1 assigns equal probability to each state after signal realization, each action is a best response.

Suppose that the receiver's initial belief assigns equal probability to each state. In this case, our approach uses two signals, and the receiver chooses either a_1' or a_1'' depending on the realization of signal. Moreover, any message (with two signals) is rationalizable. As noted, if the receiver's interim belief assigns equal probability to each state, a_1''' becomes a best response as well. If the sender expects that the receiver chooses a_1''' in this case, she will simply avoid the message which results in this particular interim belief. Unlike the first possibility

[23] A similar observation can be found in the ultimatum game with the continuous action space for the first mover, for example.

[24] See, for example, Dufwenberg and Stegeman (2002).

considered above, the presence of multiple best responses for the receivers may create multiple rationalizable messages in this type of circumstance.[25] The second possibility is however generically avoided (via perturbations in the receivers' payoffs), and hence we have the following result.

Proposition 6 *If the sender is indifferent between two actions after signal realization, it is generically because rationalizable message needs binding constraints.*

5.3 Beneficial Persuasion and Uniqueness of Rationalizable Message

Let a_i^0 be a best response to a null message π^0 (and hence her initial belief) for each $i \in N$. Given a rationalizable message π^*, we say that the persuasion from a^0 to t^* is *beneficial* if $u_S^+(t^*) > u_S^+(a_0)$. If a^0 and t^* are both unique best responses (to the initial beliefs and the interim beliefs, respectively), we simply say that the persuasion is beneficial. Note that if π^* is null, the persuasion is not beneficial. Therefore, we immediately have the following result.

Lemma 2 *Given π^*, if the persuasion is beneficial, there exists receiver $i \in N$ such that $|\Xi_i^*| \geq 2$.*

The converse does not hold, however. That there exists receiver $i \in N$ such that $|\Xi_i^*| \geq 2$ could suggest that receiver i may choose an action different from a_i^0 depending on signal realization. However, Example 4 suggests this is not the case.[26] Even if the sender avoids a_1''', it is still the case that a rationalizable message may use two signals. Although these two actions have the same payoff implication for the sender and hence the sender can actually use a null message (unless the receiver's initial belief says that two states are equally likely), our argument in Subsection 4.3 does not rule out non-null messages. It is, however, the case that perturbations in the sender's payoffs generically avoid this type of redundancy.

[25] Sender-preferred equilibrium simply eliminates a_1'''. Since receiver 1 chooses an action from $\{a_1', a_1''\}$ independent of her interim belief, there is no role for persuasion.

[26] Example 2 (with common prior) can be used as well.

	a_1'	a_1''	a_1'''
θ'	2, 1	0, 0	0, $\frac{1}{2}+\varepsilon$
θ''	0, 0	1, 1	0, $\frac{1}{2}+\varepsilon$

Figure 6: ξ_1^0 does not exist

Proposition 7 Ξ^* *is generically unique.*

We make three points regarding a_i^0. First, since A_i is finite, a_i^0 is generically unique for each $i \in N$. Second, if π^* is null for receiver $i \in N$ (and hence Ξ_i^* is a singleton), receiver i chooses a_i^0. Third, given π^*, there may be no signal in Ξ_i^* which induces a_i^0, as shown in the following example.

Example 5. We have the sender and receiver 1. The payoffs are summarized in Figure 6, where (i) the first and second elements are the sender's and receiver 1's payoffs respectively, and (ii) $\varepsilon > 0$ is arbitrarily small. Assume that $p_1^0(\theta') \approx \frac{1}{2}$ and the unique best response to the initial belief is a_1'''. The rationalizable message is such that the sender reveals the state. After observing any signal, receiver 1 has a unique best response, which is different from a_1'''.

As shown in the previous subsection, even if a receiver has multiple best responses, (i) it is generically due to the sender's strategic motive, and (ii) in such cases, there is no corresponding rationalizable message if the sender chooses another best response. Moreover, Proposition 7 says that Ξ_i^* is generically unique for each $i \in N$. Perturbations in the sender's payoff therefore lead to the following result:[27]

Proposition 8 *The following two statements hold generically:*

[27] Regarding part 1 of Proposition 8, Proposition 1 in Mangasarian (1979) provides the formal treatment. Regarding part 2 of Proposition 8, the same result for the case of one receiver with the common prior can be deduced from Proposition 2 of Kamenica and Gentzkow (2011) by refining their notion of "information the sender would share".

1. π^* is unique, and

2. persuasion is beneficial or π^* is null.

Part 1 of Proposition 8 immediately implies that while π^* is sensitive to the receivers' beliefs, small changes in the sender's belief do not affect π^* if π^* is unique. In other words, two senders with slightly different beliefs would choose the same rationalizable message. In addition, part 2 of Proposition 8 implies that if π^* is null, persuasion is not beneficial. This however only holds for unlinked Bayesian persuasion games. For a linked Bayesian persuasion game where the sender also plays the stage game as a receiver, the sender may persuade the receivers with the null message, which is demonstrated by our example in the next section. This suggests that the notion of beneficial persuasion is appropriate only for unlinked Bayesian persuasion games.

6 Linked Bayesian Persuasion Games

The presence of strategic interactions at the stage game makes it more challenging to analyze linked Bayesian persuasion games. Several points:

- The informational content of signals becomes richer – the sender may reveal to the receivers some information regarding others' signal observations.
- That the sender can play the stage game may sharpen the predictions – there is a possibility of tacit understanding.[28]
- The linear programming approach does not work since players' beliefs concern not only the states, but also the opponents' behavior. However, we can still analyze any linked Bayesian persuasion game.

[28] Note that if the sender plays the stage game in an unlinked Bayesian persuasion game, it is a single-person decision problem by definition.

θ'		a_2' (receiver 2)	a_2'' (receiver 2)
receiver 1	a_1'	0, 1, 1	0, 1, 0
	a_1''	0, 0, 0	1, 0, 1

θ''		a_2' (receiver 2)	a_2'' (receiver 2)
receiver 1	a_1'	0, 0, 0	1, 0, 1
	a_1''	0, 1, 1	0, 1, 0

Figure 7: Privately Talking about Others

- Unlike unlinked Bayesian persuasion games, the application of rationalizability may involve more than two steps of elimination procedures.[29] Generic properties discussed in Section 5 do not necessarily hold.

We already discussed the first two points in the Introduction. In the following subsections, we elaborate on them by means of examples.

6.1 Privately Talking about Others

In unlinked Bayesian persuasion games, the sender does not have an incentive to provide with the receivers information regarding others' signal observations, since it would provide more information regarding the states. But in a linked Bayesian persuasion game, the sender may find it optimal. The sender, however, would have to choose how much information to provide to each receiver.

Example 6. We have the sender and receivers 1 and 2. The incomplete information game is shown in Figure 7 where for each cell, the first, the second, and the third elements are the payoffs for the sender, receiver 1, and receiver 2 respectively. Receiver 2 would like to coordinate with receiver 1, either (a_1', a_2') or (a_1'', a_2''), when the state is θ', while she would like to avoid coordination when the state is θ''. The sender obtains a positive payoff for (a_1'', a_2'') for θ' or (a_1', a_2'') for θ''. Receiver 1 not only has the unique dominant action for each state, a_1'

[29]For linked games, it is often easier to use the dominance argument. The dominance counterpart of Δ-rationalizability can be viewed as the incomplete information version of *conditional dominance* by Shimoji and Watson (1998).

for θ' and a_1'' for θ'', but also does not care about receiver 2's action. We assume that $p_S^0(\theta') < p_1^0(\theta')$. Note that this example is similar to Example 2 regarding the sender and receiver 1, except that (dis)coordination between the receivers is required in this example.

Consider the following message. We have $\Xi_1 = \{\xi_1', \xi_1''\}$ and $\Xi_2 = \{\xi_2', \xi_2''\}$. For receiver 1, follow the solution identified in (b) of Figure 3 so that receiver 1 chooses a_1' for ξ_1' and a_1'' for ξ_1''. For receiver 2, signal realization depends on both the state and receiver 1's signal: (i) ξ_2' will be realized if $\{\theta', \xi_1'\}$ or $\{\theta'', \xi_1''\}$, and (ii) ξ_2'' will be realized if $\{\theta', \xi_1''\}$ or $\{\theta'', \xi_1'\}$. Given receiver 1's signal, the sender recommends a_2' with ξ_2' and a_2'' with ξ_2''. The (dis)coordination of actions between the receivers is perfect from the sender's point of view. Remember that receiver 1's payoff is independent of receiver 2's action. Given the argument provided for the solution in Example 2 which is shown in (b) of Figure 3, the sender's overall message is optimal.

In this example, the sender lets only receiver 2 know receiver 1's signal observation. Note that if receiver 1 knows the signal that receiver 2 observes, combined with her own signal observation, she will know the state. The sender would never provide such information to receiver 1.

6.2 Tacit Understanding

In the following example, the sender also plays at the stage game as a receiver. In unlinked Bayesian persuasion games, if the sender chooses a null message, it is necessarily the case that persuasion is not beneficial. The example shows that this is no longer the case: "silence" (null message) may have some meanings.

Example 7. We have the sender and receiver 1. The sender not only chooses a message, but also plays the stage game (as a receiver). The incomplete information game is shown in (a) of Figure 8. They have different views regarding the likelihood of each state. While receiver 1 believes that each state is equally likely, the sender is pessimistic and believes that the probability of

Receiver 1

θ'		L	R
Sender	U	$\underline{4}$, $\underline{2}$	$\underline{3}$, 0
	D	2, $\underline{1}$	2, 0

Receiver 1

θ''		L	R
Sender	U	$\underline{1}$, -1	-1, $\underline{0}$
	D	0, -2	$\underline{0}$, $\underline{0}$

(a) Stage Game

Receiver 1

		L	R
Sender	U	$\underline{\frac{1}{2}}$	0
	D	$-\frac{1}{2}$	$\underline{0}$

(b) Expected Payoffs for Receiver 1

Receiver 1

		L	R
Sender	U	$\underline{\frac{11}{5}}$	$\frac{3}{5}$
	D	$\frac{4}{5}$	$\underline{\frac{4}{5}}$

(c) Expected payoffs from Sender's point of view

Figure 8: Tacit Understanding

θ' is only $\frac{2}{5}$.

As a benchmark, consider the message which reveals the actual state to both. In this case, receiver 1 has a dominant action for each state: L for θ' and R for θ'', implying that the outcomes are (U, L) for θ' and (D, R) for θ''. The sender's expected payoff is then $\frac{8}{5}(= 4 \times \frac{2}{5} + 0 \times \frac{3}{5})$. It is however desirable for the sender that receiver 1 chooses L independent of the state (and she chooses high effort).

Suppose instead that the sender chooses the null message and they play the second-stage game with their priors. The corresponding expected payoffs for receiver 1 are shown in (b) of Figure 8. It is no longer the case that receiver 1 has a dominant action; she would choose L for U and R for D. The corresponding expected payoffs for the sender are shown in (c) of Figure 8. The sender would choose U for L and D for R – coordination problem.

However, the fact that the sender could have revealed the state solves this coordination problem. If the sender reveals the state, her expected payoff is $\frac{8}{5}$ as shown above. Note that this is strictly greater than her expected payoffs with D shown in (c) of Figure 8. Hence, if the sender chooses not to reveal any information, receiver 1 infers that the sender will not choose D and will choose L. Hence, by choosing not to influence their beliefs, the sender successfully persuades receiver 1 to choose L even though the state could be θ''.

This example shows a possibility that the sender may not reveal any additional information to persuade the receivers; having a chance to disclose the state is sufficient. The conclusion depends on the receiver's counterfactual reasoning regarding the sender's behavior. This is the implication of forward induction by rationalizability.[30]

[30] In this example, there are other rationalizable messages. If, for example, we introduce a small cost of communication, the message we identified is the unique rationalizable message. Moreover, if iterative weak dominance is used instead, any message surviving iterative weak dominance will lead to the outcome where receiver 1 always chooses L. For the epistemic foundation of weak dominance, see Brandenburger, Friedenberg, and Keisler (2008). For the relationship between sequential best response and weak dominance, see Shimoji (2004).

7 Conclusion

In this paper, we analyzed Bayesian persuasion games from Kamenica and Gentzkow (2011). We showed that it is possible to analyze multiple-receiver Bayesian persuasion games with heterogeneous beliefs. Our departure from Kamenica and Gentzkow (2011) is that we directly analyze the sender's messages. With our approach, the sender's optimization problem is the examination of a set of linear programming problems for unlinked Bayesian persuasion games. This is supported by rationalizability.

For linked Bayesian persuasion games, we also showed (i) that the sender strictly prefers either "doing nothing (null message)" or "persuading receivers", and (ii) that the rationalizable message is generically unique. We also showed some results on the constraints. In particular, for a receiver with two actions, if she is indifferent between the two actions, it is after observing the signal inducing the action which she will not choose without persuasion.

In addition, we provided two examples for linked Bayesian persuasion games. The first example showed that the message may contain richer information (i.e., not only about the states, but also other receivers' signals). The second example, where the sender also plays the stage game, showed that "silence" may have some meanings. This implies that the notion of beneficial persuasion would be appropriate only for unlinked Bayesian persuasion games.

Our linear programming approach implies that any unlinked Bayesian persuasion game can be "mechanically" analyzed. Many interesting examples would be linked Bayesian persuasion games. There is no reason to exclude dynamic games as stage games. Although the analyses of linked Bayesian persuasion games may not be straightforward, we hope that there will be more studies of such games.[31]

[31] If the stage game is a *Bayesian potential game*, the sender can utilize the potential function to analyze the receivers' behavior via Bayes Nash equilibrium. The definition of potential games is by Monderer and Shapley (1996). For Bayesian potential games, see van HeuMen, Peleg, Tijs, and Born (1996) and Ui (2009).

References

[1] Alonso, Ricardo, and Odilon Câmara. 2016a. "Bayesian Persuasion with Heterogeneous Priors." *Journal of Economic Theory*, forthcoming.

[2] Alonso, Ricardo, and Odilon Câmara. 2016b. "Persuading Voters." *American Economic Review*, forthcoming.

[3] Aumann, Robert J., and Michael B. Maschler. 1995. Repeated Games with Incomplete Information. MIT Press.

[4] Battigalli, Pierpaolo. 1999. "Rationalizability in Incomplete Information Games." EUI Working Paper (ECO No.99/17) available at http://cadmus.eui.eu/handle/1814/702.

[5] Battigalli, Pierpaolo, and Marchiano Siniscalchi. 2003. "Rationalization and Incomplete Information." *Advances in Theoretical Economics*, 3(1).

[6] Bergemann, Dirk, and Stephen Morris. 2016. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics*, 11(2), 487-522.

[7] Brandenburger, Adam, Amanda Friedenberg, and H. Jerome Keisler. (2008) "Admissibility in Games." *Econometrica*, 76(2), 307-352.

[8] Brocas, Isabelle, and Juan D. Carrillo. 2007. "Influence through Ignorance." *Rand Journal of Economics*, 38(4), 931-947.

[9] Chan, Jimmy, Fei Li, and Yun Wang. 2016. "Discriminatory Persuasion: How to Convince a Group." Working Paper available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2688044.

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

[11] Dufwenberg, Martin, and Mark Stegeman. 2002. "Existence and Uniqueness of Maximal Reductions under Iterated Strict Dominance." *Econometrica*, 70(5), 2007-2023.

[12] Farrell, Joseph, and Robert Gibbons. 1989. "Cheap Talk with Two Audiences." *American Economic Review.* 79(5), 1214-1223.

[13] Goltsman, Maria, and Gregory Pavlov. 2011. "How to Talk to Multiple Audiences." *Games and Economic Behavior.* 72(1), 100-122.

[14] Heifetz, Aviad, Martin Meier, and Burkhard C. Schipper. 2013. "Dynamic Unawareness and Rationalizable Behavior." *Games and Economic Behavior*, 81, 50-68.

[15] van Heumen, Robert, Bezalel Peleg, Stef Tijs, and Peter Born. 1996. "Axiomatic Characterizations of Solutions for Bayesian Games." *Theory and Decision*, 40(2), 103-129.

[16] Hörner, Johannes, and Andrzej Skrzypacz. 2014. "Selling Information." *Journal of Political Economy*, forthcoming.

[17] Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590-2615.

[18] Kolotilin, Anton. 2016. "Optimal Information Disclosure: Quantity vs. Quality." Working Paper available at https://sites.google.com/site/akolotilin/home

[19] Kolotilin, Anton, Ming Li, Tymofiy Mylovanov, Andriy Zapechelnyuk. 2015. "Persuasion of a Privately Informed Receiver." Working Paper available at https://sites.google.com/site/akolotilin/home, https://sites.google.com/site/tmylovanov/home, and https://sites.google.com/site/azapech/.

[20] Lipnowski, Elliot, and Laurent Mathevet. 2015. "Disclosure to a Psychological Audience." Working Paper available at http://elliotlipnowski.com and http://laurentmathevet.com.

[21] Mangasarian, O. L. 1979. "Uniqueness of Solution in Linear Programming." *Linear Algebra and its Applications*, 25: 151-162.

[22] Miura, Shintaro, and Takuro Yamashita. 2014. "On the Possibility of Information Transmission." Working Paper available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2433186.

[23] Monderer, Dov, and Lloyd S. Shapley. 1996. "Potential Games." *Games and Economic Behavior*, 14(1), 124-143.

[24] Ostrovsky, Michael, and Michael Schwarz. 2010. "Information Disclosure and Unraveling in Matching Markets." *American Economic Journal: Microeconomics*, 2(2), 34-63.

[25] Rayo, Luis, and Ilya Segal. 2010. "Optimal Information Disclosure." *Journal of Political Economy*, 118(5), 949-987.

[26] Shimoji, Makoto. 2004. "On the Equivalence of Weak Dominance and Sequential Best Response." *Games and Economic Behavior*, 48(2), 385-402.

[27] Shimoji, Makoto, and Joel Watson. 1998. "Conditional Dominance, Rationalizability, and Game Forms." *Journal of Economic Theory*, 83(2), 161-195.

[28] Sobel, Joel. 2013. "Giving and Receiving Advice." In *Advances in Economics and Econometrics, Tenth World Congress, Volume I, Economic Theory*, ed. D. Acemoglu, M. Arellano, and E. Dekel, 305-341. New York: Cambridge University Press.

[29] Taneva, Ina A. 2016. "Information Design." Working Paper available at https://sites.google.com/site/itaneva13/home.

[30] Ui, Takashi. 2009. "Bayesian Potentials and Information Structures: Team Decision Problems Revisited." *International Journal of Economic Theory*, 5(3), 271-291.

[31] Wang, Yun. 2013. "Bayesian Persuasion with Multiple Receivers." available at https://sites.google.com/site/yunwang0606/.

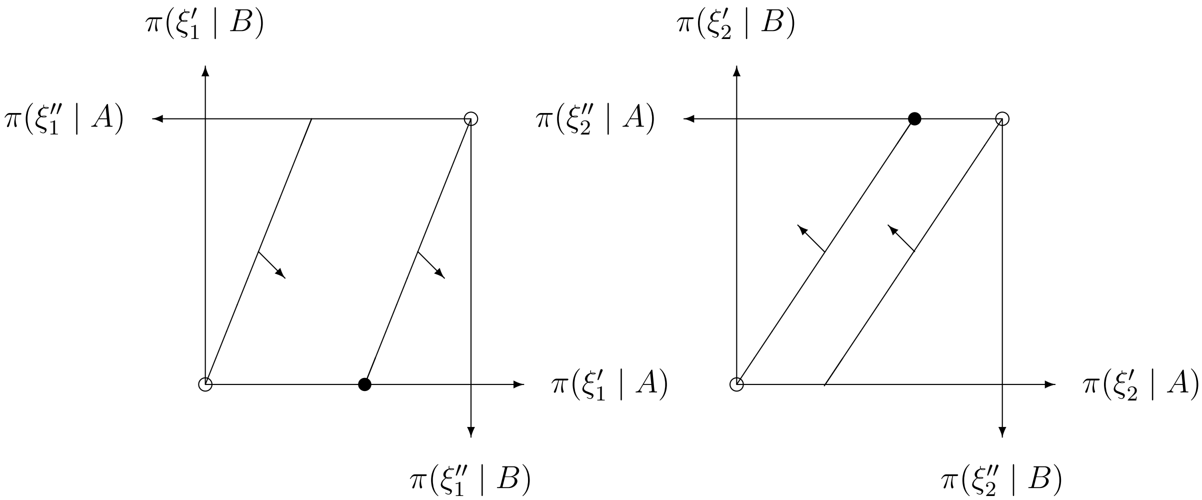

(a) Receiver 1

(b) Receiver 2

Figure 9: Persuading Only One Receiver

A Multiple Receivers: Example

Suppose that the sender attempts to persuade only one receiver. Remember that receiver 1 chooses a with her initial belief while receiver 2 chooses b with her initial belief. Suppose first that the sender only persuades receiver 1. Consider the strategy for receiver 1 "a after ξ_1' and b after ξ_1''." The corresponding constraints are in (2), also visualized in (a) of Figure 9. Given that receiver 2 will choose b with her initial belief, the sender wants to increase the chance that ξ_1'' is observed (southwest). We then have $\pi(\xi_1' \mid A) = \frac{4}{7}$, $\pi(\xi_1'' \mid A) = \frac{3}{7}$, $\pi(\xi_1' \mid B) = 0$ and $\pi(\xi_1'' \mid B) = 1$. The sender's expected payoff is $\frac{8}{10}\frac{3}{7} + \frac{2}{10} = \frac{19}{35} \approx 0.54$.

Likewise, suppose that the sender only persuades receiver 2. For the strategy "a after ξ_2' and b after ξ_2''", the corresponding constraints are in (3), also visualized in (b) of Figure 9. Given that receiver 1 will choose a with her initial belief, the sender wants to increase the chance that ξ_2' is observed (northeast). We have $\pi(\xi_2' \mid A) = \frac{2}{3}$, $\pi(\xi_2'' \mid A) = \frac{1}{3}$, $\pi(\xi_2' \mid B) = 1$ and $\pi(\xi_2'' \mid B) = 0$. The sender's expected payoff is $\frac{8}{10}\frac{2}{3} + \frac{2}{10} = \frac{11}{15} \approx 0.73$.

$$
\begin{array}{ccccccccccc}
\vdots & \vdots & & \vdots & & \vdots & & \vdots & \vdots & \vdots \\
\tilde{\xi}_i & & & \pi^+(\tilde{\xi}_i \mid \theta) & & & & & \tilde{\pi}^+(\tilde{\xi}_i \mid \theta) = \beta\pi^+(\tilde{\xi}_i \mid \theta) & \\
\vdots & \vdots & & \vdots & & \vdots & & \vdots & \vdots & \vdots \\
\xi_i' & 0 & \cdots & 0 & \cdots & 0 & & & \tilde{\pi}^+(\xi_i' \mid \theta) = (1-\beta)\pi^+(\tilde{\xi}_i \mid \theta) & \\
\vdots & \vdots & & \vdots & & \vdots & & \vdots & \vdots & \vdots \\
 & & \cdots & \theta & \cdots & & & \cdots & \theta & \cdots \\
 & & & \pi^+ & & & & & \tilde{\pi}^+ &
\end{array}
$$

Figure 10: Construction of $\tilde{\pi}$

B Proof of Proposition 1

Lemma 3 *Consider $t^+ \in T^+$. Suppose that the solution for the sender, $\pi^+ \in \bar{\Pi}^+(t^+)$, is such that $\pi^+ \notin \Pi(t^+)$. Then, there exists $\tilde{t}^+ \in T^+$ and $\tilde{\pi}^+ \in \bar{\Pi}^+(\tilde{t}^+)$ such that*

$$\tilde{\pi}^+(\xi_i \mid \theta) > 0 \textit{ for every } \xi_i \in \Xi_i \textit{ and } \theta \in \Theta$$

and

$$\sum_{\theta \in \Theta} \sum_{\xi \in \Xi} u_S(\tilde{t}^+(\xi), \theta)\tilde{\pi}(\xi \mid \theta) p_S^0(\theta) = u_S^+(t^+).$$

Proof. We start with $\pi^+(\xi_i \mid \theta)$ while fixing $\pi^+(\xi_j \mid \theta)$ for every $j \neq i$. Take $\tilde{\xi}_i \in \Xi_i$ such that $\pi^+(\tilde{\xi}_i \mid \tilde{\theta}) > 0$ for some $\tilde{\theta} \in \Theta$. Take any $\beta \in (0, 1)$ and let (i) $\tilde{\pi}^+(\tilde{\xi}_i \mid \theta) = \beta\pi^+(\tilde{\xi}_i \mid \theta)$ for each $\theta \in \Theta$ and (ii) $\tilde{\pi}^+(\xi_i' \mid \theta) = (1-\beta)\pi^+(\tilde{\xi}_i \mid \theta)$ for each $\theta \in \Theta$ (see Figure 10). Given this, let $\tilde{t}_i^+$ be the corresponding best response for receiver i. Note that while $t_i^+(\xi_i')$ is not well defined, $\tilde{t}_i^+(\xi_i') = \tilde{t}_i^+(\tilde{\xi}_i) = t_i^+(\tilde{\xi}_i)$ since the interim belief is homogeneous of degree zero in π. Hence, we have $\tilde{t}^+ \in T^+$ and $\tilde{\pi}^+ \in \bar{\Pi}^+(\tilde{t}^+)$. Since the outcome does not change, they lead to the same expected payoff for the sender. $\square$

If it is still the case that $\tilde{\pi}^+ \notin \Pi(\tilde{t}^+)$, we repeat the same argument. Since A_i is finite, there exist $\hat{t}^+ \in T^+$ and $\hat{\pi}^+ \in \bar{\Pi}^+(\hat{t}^+) \cap \Pi(\hat{t}^+)$.

C Proof of Proposition 3

The following result immediately implies Proposition 3.

Lemma 4 *Given π^*, suppose that there exist $i \in N$ and $\tilde{\xi}_i, \hat{\xi}_i \in \Xi_i$ such that $t_i^*(\tilde{\xi}_i) = t_i^*(\hat{\xi}_i)$. Consider $\tilde{\pi}^*$ with $\tilde{\Xi}_i = \Xi_i \backslash \{\hat{\xi}_i\}$ such that*

1. *$\tilde{\pi}^*(\tilde{\xi}_i \mid \theta) = \pi^*(\tilde{\xi}_i \mid \theta) + \pi^*(\hat{\xi}_i \mid \theta)$ for each $\theta \in \Theta$, and*
2. *$\tilde{\pi}^*(\xi_i \mid \theta) = \pi^*(\xi_i \mid \theta)$ for each $\xi_i \in \tilde{\Xi}_i \backslash \{\tilde{\xi}_i\}$ and $\theta \in \Theta$.*

Let $\tilde{\Xi} = \tilde{\Xi}_i \times \Xi_{-i}$. Then, $\tilde{s}^*$ with $\tilde{s}^*(\tilde{\pi}^*, \xi) = s^*(\pi^*, \xi)$ for each $\xi \in \tilde{\Xi}$ is a best response profile with respect to $\tilde{\pi}^*$.

Proof. Note that the comparison of any two actions after observing any signal in $\tilde{\Xi}_i \backslash \{\tilde{\xi}_i\}$ remains the same. The comparison of the expected payoffs from $s_i^*(\pi^*, \tilde{\xi}_i)$ and $a_i \in A_i \backslash \{s_i^*(\pi^*, \tilde{\xi}_i)\}$ after observing $\tilde{\xi}_i$ gives

$$
\begin{aligned}
& \textstyle\sum_{\theta\in\Theta} u_i(s_i^*(\pi^*, \tilde{\xi}_i), \theta) \frac{\tilde{\pi}^*(\tilde{\xi}_i|\theta) p_i^0(\theta)}{\sum_{\theta'\in\Theta} \tilde{\pi}^*(\tilde{\xi}_i|\theta') p_i^0(\theta')} - \sum_{\theta\in\Theta} u_i(a_i, \theta) \frac{\tilde{\pi}^*(\tilde{\xi}_i|\theta) p_i^0(\theta)}{\sum_{\theta'\in\Theta} \tilde{\pi}^*(\tilde{\xi}_i|\theta') p_i^0(\theta')} \\
= \;& \left[\frac{1}{\sum_{\theta'\in\Theta} \tilde{\pi}^*(\tilde{\xi}_i|\theta') p_i^0(\theta')}\right] \\
& \times \left[\textstyle\sum_{\theta\in\Theta} u_i(s_i^*(\pi^*, \tilde{\xi}_i), \theta) \tilde{\pi}^*(\tilde{\xi}_i \mid \theta) p_i^0(\theta) - \sum_{\theta\in\Theta} u_i(a_i, \theta) \tilde{\pi}^*(\tilde{\xi}_i \mid \theta) p_i^0(\theta)\right] \\
= \;& \left[\frac{1}{\sum_{\theta'\in\Theta} \tilde{\pi}^*(\tilde{\xi}_i|\theta') p_i^0(\theta')}\right] \\
& \times \Big[\textstyle\sum_{\theta\in\Theta} u_i(s_i^*(\pi^*, \tilde{\xi}_i), \theta) [\pi^*(\tilde{\xi}_i \mid \theta) + \pi^*(\hat{\xi}_i \mid \theta)] p_i^0(\theta) \\
& - \textstyle\sum_{\theta\in\Theta} u_i(a_i, \theta) [\pi^*(\tilde{\xi}_i \mid \theta) + \pi^*(\hat{\xi}_i \mid \theta)] p_i^0(\theta)\Big] \\
= \;& \left[\frac{1}{\sum_{\theta'\in\Theta} \tilde{\pi}^*(\tilde{\xi}_i|\theta') p_i^0(\theta')}\right] \\
& \times \Big\{\left[\textstyle\sum_{\theta\in\Theta} u_i(s_i^*(\pi^*, \tilde{\xi}_i), \theta) \pi^*(\tilde{\xi}_i \mid \theta) p_i^0(\theta) - \sum_{\theta\in\Theta} u_i(a_i, \theta) \pi^*(\tilde{\xi}_i \mid \theta) p_i^0(\theta)\right] \\
& \times \left[\textstyle\sum_{\theta\in\Theta} u_i(s_i^*(\pi^*, \tilde{\xi}_i), \theta) \pi^*(\hat{\xi}_i \mid \theta) p_i^0(\theta) - \sum_{\theta\in\Theta} u_i(a_i, \theta) \pi^*(\hat{\xi}_i \mid \theta) p_i^0(\theta)\right]\Big\}
\end{aligned}
$$

$$
\begin{aligned}
= \;& \left[\frac{\sum_{\theta'\in\Theta}\pi^*(\tilde{\xi}_i|\theta')p_i^0(\theta')}{\sum_{\theta'\in\Theta}\tilde{\pi}^*(\tilde{\xi}_i|\theta')p_i^0(\theta')}\right] \\
& \times\left[\sum_{\theta\in\Theta}u_i(s_i^*(\pi^*,\tilde{\xi}_i),\theta)\frac{\pi^*(\tilde{\xi}_i|\theta)p_i^0(\theta)}{\sum_{\theta'\in\Theta}\pi^*(\tilde{\xi}_i|\theta')p_i^0(\theta')}-\sum_{\theta\in\Theta}u_i(a_i,\theta)\frac{\pi^*(\tilde{\xi}_i|\theta)p_i^0(\theta)}{\sum_{\theta'\in\Theta}\pi^*(\tilde{\xi}_i|\theta')p_i^0(\theta')}\right] \\
& +\left[\frac{\sum_{\theta'\in\Theta}\pi^*(\hat{\xi}_i|\theta')p_i^0(\theta')}{\sum_{\theta'\in\Theta}\tilde{\pi}^*(\hat{\xi}_i|\theta')p_i^0(\theta')}\right] \\
& \times\left[\sum_{\theta\in\Theta}u_i(s_i^*(\pi^*,\tilde{\xi}_i),\theta)\frac{\pi^*(\hat{\xi}_i|\theta)p_i^0(\theta)}{\sum_{\theta'\in\Theta}\pi^*(\hat{\xi}_i|\theta')p_i^0(\theta')}-\sum_{\theta\in\Theta}u_i(a_i,\theta)\frac{\pi^*(\hat{\xi}_i|\theta)p_i^0(\theta)}{\sum_{\theta'\in\Theta}\pi^*(\hat{\xi}_i|\theta')p_i^0(\theta')}\right].
\end{aligned}
$$

Remember that $s_i^*(\pi^*,\tilde{\xi}_i)=s_i^*(\pi^*,\hat{\xi}_i)$. Hence, $\tilde{s}_i^*(\tilde{\pi}^*,\xi_i)=s_i^*(\pi^*,\xi_i)$ for each $\xi_i\in\tilde{\Xi}_i$. The result follows.

D Proof of Proposition 4

Given $\tilde{\xi}_i,\hat{\xi}_i\in\Xi_i^*$, let $\Pi(s^+,\tilde{\xi}_i,\hat{\xi}_i)\subset\Pi(s^+)$ be such that for any $\pi',\pi''\in\Pi(s^+,\tilde{\xi}_i,\hat{\xi}_i)$, (i) $\pi'(\xi_j\mid\theta)=\pi''(\xi_j\mid\theta)$ for each $j\neq i$ and $\theta\in\Theta$ and (ii) $\pi'(\xi_i\mid\theta)=\pi''(\xi_i\mid\theta)$ for each $\xi_i\neq\tilde{\xi}_i,\hat{\xi}_i$ and $\theta\in\Theta$. In other words, π' and π'' can be different only with respect to $\tilde{\xi}_i$ and $\hat{\xi}_i$. Note that this implies that $c(\tilde{\xi}_i,\hat{\xi}_i\mid\theta)$ is fixed for each $\theta\in\Theta$. Let

$$
\begin{aligned}
\tilde{\Pi}(s^+,\tilde{\xi}_i,\hat{\xi}_i) &= \left\{\pi\in\Pi(s^+,\tilde{\xi}_i,\hat{\xi}_i)\mid(9)\text{ holds with equality}\right\} \\
\hat{\Pi}(s^+,\tilde{\xi}_i,\hat{\xi}_i) &= \left\{\pi\in\Pi(s^+,\tilde{\xi}_i,\hat{\xi}_i)\mid(10)\text{ holds with equality}\right\}
\end{aligned}
$$

The following result shows that (9) and (10) hold simultaneously with equalities if and only if (11) is zero.

Lemma 5 *It is either (i) $\tilde{\Pi}(s^+,\tilde{\xi}_i,\hat{\xi}_i)\cap\hat{\Pi}(s^+,\tilde{\xi}_i,\hat{\xi}_i)=\emptyset$ or (ii) $\tilde{\Pi}(s^+,\tilde{\xi}_i,\hat{\xi}_i)=\hat{\Pi}(s^+,\tilde{\xi}_i,\hat{\xi}_i)$. Moreover, the latter holds if and only if (11) is zero.*

Proof. (9) and (10) can be rewritten as

$$
\begin{aligned}
& \sum_{\theta\in\Theta}\left[u_i(s_i^+(\pi,\tilde{\xi}_i),\theta)-u_i(s_i^+(\pi,\hat{\xi}_i),\theta)\right]\left[c(\tilde{\xi}_i,\hat{\xi}_i\mid\theta)-\pi(\hat{\xi}_i\mid\theta)\right]p_i^0(\theta)\geq 0 \\
\Leftrightarrow\; & \sum_{\theta\in\Theta}\left[u_i(s_i^+(\pi,\tilde{\xi}_i),\theta)-u_i(s_i^+(\pi,\hat{\xi}_i),\theta)\right]c(\tilde{\xi}_i,\hat{\xi}_i\mid\theta)p_i^0(\theta)
\end{aligned}
$$

$$\geq \sum_{\theta \in \Theta} \left[u_i(s_i^+(\pi, \tilde{\xi}_i), \theta) - u_i(s_i^+(\pi, \hat{\xi}_i), \theta) \right] \pi(\hat{\xi}_i \mid \theta) p_i^0(\theta) \tag{12}$$

and

$$\begin{aligned}
&\sum_{\theta \in \Theta} \left[u_i(s_i^+(\pi, \tilde{\xi}_i), \theta) - u_i(s_i^+(\pi, \hat{\xi}_i), \theta) \right] \left[c(\tilde{\xi}_i, \hat{\xi}_i \mid \theta) - \pi(\tilde{\xi}_i \mid \theta) \right] p_i^0(\theta) \leq 0 \\
\Leftrightarrow\ &\sum_{\theta \in \Theta} \left[u_i(s_i^+(\pi, \tilde{\xi}_i), \theta) - u_i(s_i^+(\pi, \hat{\xi}_i), \theta) \right] c(\tilde{\xi}_i, \hat{\xi}_i \mid \theta) p_i^0(\theta) \\
&\leq \sum_{\theta \in \Theta} \left[u_i(s_i^+(\pi, \tilde{\xi}_i), \theta) - u_i(s_i^+(\pi, \hat{\xi}_i), \theta) \right] \pi(\tilde{\xi}_i \mid \theta) p_i^0(\theta).
\end{aligned} \tag{13}$$

Suppose that (9) and (10) (and hence (12) and (13)) hold with equalities. Suppose that (11) is positive. Then the right-hand side of (13) is positive, contradicting (9) with equality. Likewise, suppose that (11) is negative. Then, the right-hand side of (12) is negative, contradicting (10) with equality. Hence, if (11) is non-zero, $\tilde{\Pi}(s^+, \tilde{\xi}_i, \hat{\xi}_i) \cap \hat{\Pi}(s^+, \tilde{\xi}_i, \hat{\xi}_i) = \emptyset$.

Suppose that (11) is zero. Then, the comparisons of (i) (9) and (13) and (ii) (10) and (12) show that (9) and (10) hold with equalities simultaneously, implying that $\tilde{\Pi}(s^+, \tilde{\xi}_i, \hat{\xi}_i) = \hat{\Pi}(s^+, \tilde{\xi}_i, \hat{\xi}_i)$. This implies the "if" part of the second part. Suppose that $\tilde{\Pi}(s^+, \tilde{\xi}_i, \hat{\xi}_i) = \hat{\Pi}(s^+, \tilde{\xi}_i, \hat{\xi}_i)$, meaning that (9) and (10) hold with equalities for any $\pi \in \tilde{\Pi}(s^+, \tilde{\xi}_i, \hat{\xi}_i) = \hat{\Pi}(s^+, \tilde{\xi}_i, \hat{\xi}_i)$. This simply means that (11) is zero, implying the "only if" part of the second part. $\square$

Lemma 5 implies the last part of Proposition 4. If (11) is negative, (12) (and hence (9)) implies that (10) holds with strict inequality while (11) is positive, (13) (and hence (10)) implies that (9) holds with strict inequality, implying the "only if" part of the first two of Proposition 4. Suppose that (10) holds with equality. If (9) and hence (12) hold, the left-hand side of (12) is positive. Suppose that (9) holds with equality. If (10) and hence (13) hold, the left-hand side of (13) is negative. They imply the "if" part of the first two of Proposition 4.