



Edinburgh School of Economics
Discussion Paper Series
Number 256

Competing with Asking Prices

Ina Taneva
(The University of Edinburgh)

Date
February 2015

Published by
School of Economics
University of Edinburgh
30 -31 Buccleuch Place
Edinburgh EH8 9JT
+44 (0)131 650 8361
<http://edin.ac/16ja6A6>



THE UNIVERSITY *of* EDINBURGH

Information Design

INA TANEVA *

February 11, 2015

Abstract

There are two ways of creating incentives for interacting agents to behave in a desired way. One is by providing appropriate payoff incentives, which is the subject of mechanism design. The other is by choosing the information that agents observe, which we refer to as information design. We consider a model of symmetric information where a designer chooses and announces the information structure about a payoff relevant state. The interacting agents observe the signal realizations and take actions which affect the welfare of both the designer and the agents. We characterize the general finite approach to deriving the optimal information structure for the designer — the one that maximizes the designer's ex ante expected utility subject to agents playing a Bayes Nash equilibrium. We then apply the general approach to a symmetric two state, two agent, and two actions environment in a parameterized underlying game and fully characterize the optimal information structure: it is never strictly optimal for the designer to use conditionally independent private signals; the optimal information structure may be a public signal or may consist of correlated private signals. Finally, we examine how changes in the underlying game affect the designer's maximum payoff. This exercise provides a joint mechanism/information design perspective.

Keywords: Information design, implementation, incomplete information, Bayes correlated equilibrium, sender-receiver games.

JEL Classification: C72, D72, D82, D83.

*School of Economics, The University of Edinburgh, Edinburgh, UK. E-mail Address: ina.taneva@ed.ac.uk
I am profoundly grateful to my advisor, Laurent Mathevet, for his continuous and invaluable help, support and encouragement. I am greatly indebted to Maxwell Stinchcombe and Thomas Wiseman for their thoughtful comments, advice and guidance. I also wish to thank numerous seminar and conference audiences for helpful feedback. All remaining errors are my own.

1 Introduction

In many economic and social settings one person or institution communicates with multiple agents engaged in a strategic interaction. In court, a prosecutor presents the results of her investigation to a jury consisting of several members. In advertising, a company chooses how much and what type of information to reveal about its new product to target different groups of customers through samples, demo versions, and information brochures. In politics, election platforms are designed to appeal to constituents, government officials, as well as leaders of other countries. In financial markets, a firm discloses information about its profitability that is relevant to both shareholders and competitors. In economic policy, the Fed releases information about its stimulus campaign, which affects the economic outlook and decisions of consumers, as well as domestic and foreign investors.

These are but a few settings of economic importance that provide context for our general questions: What is the optimal mode of information transmission between a self-interested sender (designer) and a group of strategically interacting receivers (agents), who form their beliefs and take actions based on the information provided? If agents are rational Bayesian players, can the designer select the information structure in a way that makes them play an equilibrium profile most beneficial to her? Is it always optimal for the sender to send a public message observed by all receivers? Or is it sometimes optimal to send privately observed signals? If so, what is the optimal degree of correlation between these private signals? Further, when is the optimal information structure symmetric, and when is it optimal to design an asymmetric information structure? We think of this as the subject of *information design*.

Consider a general environment with multiple interacting agents who choose within a set of possible actions. Their payoffs are determined by their own action, the actions of their opponents, and the realization of a payoff relevant state with a commonly known prior distribution. We refer to this as the *basic game*. In order to analyze the strategic interactions in this setting, we need to also specify what the agents believe about the payoff state, what they believe about their opponents' beliefs, what they believe their opponents believe about their beliefs, and so on. This is captured by the *information structure*. Consider a designer who has preferences over the payoff state and the actions taken by the agents. Mechanism design takes the information structure as given and modifies the basic game so that the agents achieve the designer's desired objective in equilibrium. In contrast, information design takes the basic game as given and selects the information structure which maximizes the designer's objective in equilibrium.

We study the general problem of a self-interested designer communicating with multiple agents engaged in a strategic interaction. Before observing the state of the world, the designer chooses the information structure which maximizes her objective in expectation. The designer's objective is an arbitrary function of the state and the agents' actions. The choice of

information structure can be viewed as a choice of joint distributions over signals conditional on different states. Once the designer chooses the information structure, it becomes common knowledge. The agents then observe the signal realizations and formulate their beliefs about the state of the world as well as their higher order beliefs. After this, they take actions which affect their own, their opponent's, and the designer's payoffs.

The main contribution of this paper is laying out the methodology of information design in finite settings. The general problem is that of engineering the information structure which for the given basic game supports a Bayes Nash equilibrium that maximizes the designer's objective in expectation. In order to do that, we need to first characterize the set of all Bayes Nash equilibria for all possible information structures. This seems like a daunting task, especially in view of the many different beliefs and higher order beliefs we would need to keep track of. Bergemann and Morris (5) provide a tool that allows us to accomplish this objective. They introduce a definition of *Bayes correlated equilibrium* under which we show how to characterize the set of all Bayes Nash equilibria associated with all possible information structures for a given basic game. We do that by constructing the set of Bayes correlated equilibria under the assumption that agents have no further information but their prior beliefs. Using this approach, we are able to characterize the set of all Bayes Nash equilibria without explicitly using information structures. Then we maximize the designer's objective function over this set to find the optimal Bayes Nash equilibrium. After that, we back out the information structure which supports it as a Bayes Nash equilibrium for the given basic game.

We then illustrate the general methodology outlined above in a class of symmetric problems with two agents, two actions and two states, for which it provides us with crisp results and conclusions. We work with a parameterized basic game, which is broad enough to capture many different strategic interactions. Moreover, the parameterization allows for comparative statics with respect to the degree of strategic complementarity and substitutability between agents and between each agent and the state. To the best of our knowledge, this is the first application to consider arbitrary objective functions without any a priori assumptions on the form of the information structure.

We provide a complete characterization of the optimal information structure in the symmetric binary environment. The characterization encompasses all possible designer objective functions. The optimal information structure is a function of the underlying game parameters and of the designer's objective. Not surprisingly, when the preferences of the designer and the agents are completely aligned, full information revelation is optimal. However, we also find that making preferences more aligned may in fact decrease the optimal degree of information transmission. This contrasts with results from the literature on cheap talk without commitment. For this symmetric binary setting our results demonstrate that in almost all of the cases the designer benefits from information design as opposed to revealing no information

and letting the agents interact under their prior beliefs. We further show that conditionally independent private signals are never optimal, irrespective of the designer’s objective function.

Finally, to bridge the information design and mechanism design literatures, we address the question of when modifying the payoffs may be beneficial to the information designer. We obtain clear-cut answers for some of the parameter values. This analysis can be viewed both as comparative statics with respect to the underlying game or as a joint mechanism-design/information-design perspective.

Several assumptions are crucial to our model and analysis. The first one is that the designer chooses the information structure before observing the state of the world and is able to perfectly commit to it.¹ The second one is that once the signal realizations occur, the designer cannot change or obfuscate them. Therefore, the agents know that what they observe are undistorted signal realizations from the commonly known conditional distributions. This ensures that they can update their beliefs without considerations of the designer’s incentive compatibility constraints. Third, we abstract away from any communication between the agents. In certain instances, this is a reasonable and realistic assumption. However, in other cases, it might be strategically beneficial for the agents to reveal their signals to each other. We provide some discussion regarding these issues and other possible extensions after presenting our main results.

The remainder of the paper is organized as follows. Section 2 provides an overview of the related literature, followed by the motivating example presented in Section 3. Section 4 introduces the framework and outlines the general approach to information design. In Section 5 we apply the general approach to a particular tractable environment. We provide a complete analysis and characterization of the optimal information structure in the symmetric binary case. Section 6 presents some important extensions, as well as a discussion of how these can be incorporated into the model. Section 7 concludes with some directions for future research. All proofs are relegated to the Appendix.

2 Literature Review

This paper is related to the literature on cheap talk communication. The cheap-talk framework analyzes the optimal information structure when the sender knows the realized state of the world and can send costless, non-verifiable messages. Alternatively, it can be viewed as the case when the designer cannot credibly commit to the ex ante chosen information structure and abide by it once the state of the world has been realized. Most related to our framework are the papers by Farrell and Gibbons (9) and Goltsman and Pavlov (10), which extend the cheap talk model of Crawford and Sobel (7) to an environment with two receivers/audiences.

¹In their Section 1.C, Kamenica and Gentzkow (12) provide an excellent discussion on why this assumption is in fact not as restrictive as it may appear at first.

They study the impact of costless, non-verifiable claims on the beliefs and therefore the actions of the receivers, which in turn affect the utility of both the sender and the receivers.

There are two significant differences between these papers and ours. First, in our environment, the sender has full commitment power and credibly chooses the information structure, a collection of signal distributions conditional on the state, before the state of the world has been realized. Second, the receivers in the cheap talk literature are independent decision makers and their payoffs depend only on their own action and the state of the world. Therefore, if the sender were to communicate via privately observed signals, the problem reduces to solving the single receiver case individually for each of the receivers. In our framework, in contrast, the receivers are involved in a strategic interaction with each other, i.e. they play a game. In this case, even if the sender were to communicate via private signals, the information structure affects the higher order beliefs of the agents, which in turn impact the equilibrium actions.

This paper is also closely related to the literature on Bayesian persuasion, which is sometimes referred to as *cheap talk with commitment*. A central paper in that literature is Kamenica and Gentzkow (12), which is equivalent to information design with one agent. They characterize the optimal signal for any given set of preferences and initial beliefs with techniques from convex analysis. However, the tools used by Kamenica and Gentzkow (12) are not sufficient to address the question in an environment with multiple interacting receivers, as the authors themselves point out: *“There is an important third class of multiple-receiver models, however, where our results do not extend easily: those where the receivers care about each other’s actions and Sender can send private signals to individual receivers.”*²

We use a definition of Bayes correlated equilibrium proposed by Bergemann and Morris (5) to answer this open question and show how things differ in the multiple-interacting-receivers environment. We suggest a general approach to the optimal design of information structures, which can be applied in very general environments. We also provide insights regarding the characterization of optimal information structures for different properties of the designer’s objective function and of the underlying game played by the interacting agents. While Bergemann and Morris (5) provide the tool that enables our analysis, we use it for very different purposes. They focus on characterizing the set of possible Bayes Nash equilibrium outcomes that can arise when players have observed at least a certain level of information and potentially more. They further describe a partial order on information structures under which the size of the equilibrium set varies monotonically. There is no “designer” in their paper, who chooses the information structure with her objective maximization in mind.

Wang (19) also examines the question of “Bayesian persuasion with multiple receivers”. She looks at a specific voting environment, in which the sender has a state-independent utility with a preference for the same alternative. Moreover, she only allows for conditionally independent private signals or purely public signals and compares these two structures. We, on the other

²See Kamenica and Gentzkow (12), p. 2609.

hand, allow for a general form of the designer’s objective and impose no a priori assumptions on the types of information structures we consider. Public and conditionally independent private signals are special cases contained in our specification. More importantly, we show that restricting attention to these two special categories of persuasion mechanisms is not without loss of generality, as the optimal information structure does not always belong to one of them. Also in a voting environment and similarly restricting attention to symmetric public signals, Alonso and Câmara (1) look at an “information controller” who chooses the information content of a public signal to maximize the probability of a proposal being approved. They characterize the optimal signal choice of the controller, as well as the voters’ preferences over electoral rules which can be adopted to induce provision of a more informative signal. This ties in nicely with the results of Lester, Persico and Visschers (15), who propose a formal model that rationalizes the exclusion of probative evidence in trials in order to increase the accuracy of fact-finding. In their framework the information designer is a judge who prescreens the evidence the jury gets to see in order to provide the proper incentives for them and improve the quality of their decision.

Eliaz and Forges (8) consider a specific environment in which a principal chooses what information to reveal to two symmetric agents whose actions are strategic substitutes. In their framework, the disclosure policy is restricted to verifiable evidence where the sender reports the set of possible states and must include the true state. The sender can control the precision of information by controlling the number of elements she includes in that set. In this setting the authors find that when the sender can commit to a disclosure policy before observing the state, it is optimal to reveal the state perfectly to one agent and disclose nothing to the other. However, this result crucially relies on the specific designer objective function they look at, the strategic substitutes assumption on the players’ actions and the hard evidence assumption. The main part of the analysis in Eliaz and Forges (8) however deals with the case of an informed principal who chooses the information disclosure in the absence of commitment, which is different from our framework. Further, we allow for the agents’ actions to be both strategic substitutes and strategic complements.

There is an extensive list of papers studying the comparison of information structures in strategic interactions: Bergemann and Morris (5), Gossner (11), Lehrer, Rosenberg and Shmaya (13) and (14), Peski (17), etc. Closest to ours is Lehrer, Rosenberg and Shmaya (13). They restrict attention to symmetric games of common interest and rank information structures according to highest player payoffs they induce under different solution concepts. In contrast, we characterize the optimal information structure under Bayes Nash equilibrium and in view of the designer’s welfare rather than the agents’ equilibrium payoffs.

A number of papers analyze the equilibrium behavior and socially optimal use of information in a tractable class of environments with quadratic payoffs and a normally distributed state of the world (Angeletos and Pavan (2), Bergemann and Morris (6), Bergemann et al.

(4)). These papers assume a specific information structure under which each player observes two normally distributed signals: a public signal common to all players and a conditionally independent signal that is privately observed. They characterize the equilibrium use of information and compare that to some efficiency benchmark. In contrast, we do not assume a particular information structure a priori. Our focus is the reverse-engineering aspect of the problem, which concerns the choice of an information structure that will decentralize the most desirable distribution over actions and states of the world as a Bayes Nash equilibrium. We are also interested in how this optimal choice changes with the designer’s objective function, which is not necessarily socially optimal.

3 Motivating Example

Consider a prosecutor who conducts an investigation and reports the resulting findings to a jury.³ The prosecutor’s objective is to convince the jury that the defendant is guilty and to achieve conviction. She chooses the investigation process and is obligated by law to fully and truthfully report the findings to the jury. The choice of investigation process can be viewed as the prosecutor’s decisions regarding which witnesses to subpoena, what questions to ask them, which forensic and other tests to order, how to structure her arguments, etc. If the defendant is guilty, then choosing a more informative investigation will tend to help the prosecutor’s case and increase the likelihood of conviction. However, if the defendant is innocent, a more informative investigation will impede the prosecutor’s case. The question we focus on is whether the prosecutor can gain by choosing the investigation process optimally, in a way that maximizes the overall probability of conviction by a jury consisting of rational Bayesian agents.

To formalize the example, suppose the jury consists of two members, indexed by i and j . There are two states of the world: the defendant is either innocent (θ_0) or guilty (θ_1). The prosecutor (designer) and the jurors (agents) share a common prior belief, which assigns probability to the defendant being innocent 70 percent of the time, $\Pr(\theta_0) = 0.7$, and guilty 30 percent of the time, $\Pr(\theta_1) = 0.3$. The jurors get utility from choosing the just action: vote to acquit (a_0) when innocent and vote to convict (a_1) when guilty. Let us assume that unanimity of the jurors’ votes is required for a verdict to be reached. If the votes are not unanimous, the case is declared a mistrial due to a deadlocked jury. The payoffs of the jury members are given by the following matrix:

³This is a multiple-agent version of the prosecutor-judge example of Kamenica and Gentzkow (12). In fact, we purposefully assume the same prior distribution and prosecutor objective, which allows for direct comparisons with the single receiver case.

$\theta = \theta_0$	a_0	a_1	$\theta = \theta_1$	a_0	a_1
a_0	2, 2	1, 0	a_0	0, 0	0, 1
a_1	0, 1	0, 0	a_1	1, 0	2, 2

Each juror receives payoff of at least 1 if he chooses the just vote irrespective of what the other juror does. If the other juror votes justly as well, the payoff is increased to 2, since then a just verdict is reached. Whenever a juror votes unjustly (to convict when innocent and to acquit when guilty) he gets a payoff of 0 irrespective of what the other juror does. The objective of the prosecutor is to achieve a conviction, irrespective of the state. Her utility function is thus given by

$$V(a^i, a^j, \theta) = \begin{cases} 1 & \text{if } a^i = a^j = a_1 \\ 0 & \text{otherwise.} \end{cases}$$

The choice of investigation can be formally represented by conditional distributions $\pi(\cdot|\theta_0)$ and $\pi(\cdot|\theta_1)$ over a set of signal realizations. The prosecutor chooses π , which then becomes common knowledge, and the jury observes the undistorted signal realizations from the investigation. This specification implies that the prosecutor is able to produce arbitrary precision, including signals that are perfectly informative of the true state. This may seem unrealistic, as often times the true state of the world cannot be revealed irrespective of the thoroughness and diligence of the investigator or the sophistication of the employed techniques. However, it is without loss of generality if we interpret the state of the world to be the most informative signal the prosecutor is able to produce. Choosing π would then determine how much and in what way the most informative signal gets obscured.

If the prosecutor chooses a completely uninformative investigation or equivalently if she chooses not to conduct one, then both jurors will vote to acquit. This is their default action profile since innocence is more likely than guilt. The prosecutor will in turn receive a certain payoff of $V(a_0, a_0) = 0$. At the other extreme, if she were to choose a completely informative investigation process, i.e. one that reveals the state perfectly, the jurors will both vote for conviction only when the defendant is indeed guilty. This happens 30 percent of the time and results in an expected payoff of 0.3 for the prosecutor.

However, the prosecutor can do better than that. The optimal investigation is in fact given by the following signal structure:

$\theta = \theta_0$	t_0	t_1	$\theta = \theta_1$	t_0	t_1
t_0	1/7	0	t_0	0	0
t_1	0	6/7	t_1	0	1

which is asymmetric with respect to the state. Under this information structure it is a BNE for each jury member to vote to acquit (a_0) when he observes t_0 and to convict (a_1) when he observes t_1 . The incentive BNE compatibility constraints are as follows:

$$u(a_0|t_0) = 1 \cdot 2 = 2 > u(a_1|t_0) = 1 \cdot 0 = 0$$

and

$$u(a_1|t_1) = \frac{2}{3} \cdot 0 + \frac{1}{3} \cdot 2 = \frac{2}{3} = u(a_0|t_1) = \frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0 = \frac{2}{3}.$$

Under this information structure and BNE, the expected value of the prosecutor’s objective function is 0.9. The jury members know that 70 percent of the defendants are innocent, yet they end up convicting 90 percent of them. They are completely aware that the investigation was chosen in a way to maximize the probability of conviction; yet they react in a rational Bayesian way given the signal realizations they observe.

We observe the same two characteristics here that Kamenica and Gentzkow (12) derive for the optimal information structure in the single receiver case. First, when each juror votes to acquit, the prosecutor’s least favorite option, he is certain that the defendant is innocent. In other words, we have $\pi(t_0, t_0|\theta_1) = \pi(t_0, t_1|\theta_1) = \pi(t_1, t_0|\theta_1) = 0$. If these probabilities were positive, the prosecutor could decrease them in favor of increasing $\pi(t_1, t_1|\theta_1)$. This will increase both the marginal probability of the signal realization (t_1, t_1) and the willingness of each juror to convict when observing t_1 . Both of these effects increase the expected payoff of the prosecutor. Hence, the optimal information structure in this setting will always have $\pi(t_1, t_1|\theta_1) = 1$.

Second, when a juror votes for conviction, he is exactly indifferent between the two votes. If he were strictly in favor of convicting, then the prosecutor could increase the probability of $\pi(t_1, t_1|\theta_0)$ and decrease the probability of $\pi(t_0, t_0|\theta_0)$, to the point at which the juror becomes indifferent. That will not change the juror’s optimal choice given t_1 — he will still choose to convict — but will increase the probability of (t_1, t_1) and hence, also the probability of conviction. The designer could increase $\pi(t_1, t_1|\theta_0)$ to the point where, conditional upon receiving t_1 , the posterior probability put on θ_0 becomes so high that the juror would choose to acquit. This turning point for the posterior on θ_0 is $\frac{2}{3}$ in this example.

A fundamental difference between the single receiver case of Kamenica and Gentzkow (12) and the current framework concerns the posterior beliefs. The judge in their example needs to have a posterior belief (on “guilty” or θ_1) of at least $\frac{1}{2}$ in order to convict. Here, in contrast, each juror convicts as long as his posterior belief on θ_1 is at least $\frac{1}{3}$. This happens because of the complementarities in the strategic interaction and the optimal choice of information structure. Since unanimity is needed for a verdict to be reached and the structure is such that both jurors always observe the same signal realization, each juror knows that voting for conviction will only really make a difference if the other juror were to vote in the same way. Therefore, receiving a signal indicative of a guilty defendant, i.e. t_1 , will make a juror more willing to vote for conviction for two reasons. First, if the defendant is guilty, then his vote is needed for a just verdict to be reached. Second, if the defendant is innocent, then voting to acquit will not help achieve the right verdict (a payoff of 2) and instead will only give him the payoff from unilaterally choosing the right action (a payoff of 1). This is because, each juror knows that conditional on receiving a signal t_1 , the other juror has received the same signal

and thus is voting to convict.⁴

We can also show that in the analogous case of strategic substitutes, the posterior on θ_1 that is necessary for conviction increases to $\frac{2}{3}$. The payoffs in this case are given by

$\theta = \theta_0$	a_0	a_1	$\theta = \theta_1$	a_0	a_1
a_0	1, 1	2, 0	a_0	0, 0	0, 2
a_1	0, 2	0, 0	a_1	2, 0	1, 1

While the prosecutor-jury framework does not make sense with this parameterization, changing the game to one with strategic substitutes provides intuition as to how that affects the results for the same objective function of the designer. In particular, the optimal information structure now has $\pi(t_0, t_0 | \theta_0) = \frac{11}{14}$ and $\pi(t_1, t_1 | \theta_0) = \frac{3}{14}$, with the distribution conditional on θ_1 remaining the same as before. Conditional on observing t_1 the posterior necessary for a juror to convict is now $\frac{2}{3}$, that is twice as high as in the case of strategic complements. This is due to the strong incentive for each juror to individually choose the just vote, irrespective of what the other juror decides. Hence, a designer who wants the agents to coordinate on the non-default action will have a harder time doing so when the underlying game is one of strategic substitutes as opposed to strategic complements.

4 The General Approach

This section describes the general approach to information design in finite environments.

4.1 Setup

There are N agents engaged in a strategic interaction. The set of agents is denoted by I and we index a generic player by $i = 1, \dots, N$. Each player has a finite set of actions A_i and we write $A = A_1 \times \dots \times A_N$ for the set of action profiles and a for a generic element of that set. There is a finite set of states Θ with θ denoting a generic element of that set. Each agent has a utility function $u_i : A \times \Theta \rightarrow \mathbb{R}$ that depends on the played action profile and on the ex ante unknown state of the world. The designer has a utility function $V : A \times \Theta \rightarrow \mathbb{R}$, so that her payoff is affected by the action profile that agents play and the state of the world. Designer and agents share a common full support prior $\psi \in \text{int}(\Delta(\Theta))$ and that is common knowledge. Let $G = ((A_i, u_i)_{i=1}^N, \psi)$. We refer to G as the *basic game*.

An *information structure* $S = ((T_i)_{i=1}^N, \pi)$ consists of a finite set of signals T_i for each player i and a signal distribution $\pi : \Theta \rightarrow \Delta(T)$ where $T = T_1 \times \dots \times T_N$. We denote by t_i

⁴In the case when there is no benefit to choosing the just vote, i.e. when instead of 1 the payoffs to mis-coordinated votes are always 0, the posterior can be as low as the prior for an equilibrium of both jurors always convicting to be achieved. This is because under the null information structure, if the other juror were to always vote to convict, it is a best response to do the same.

a generic element of T_i and similarly by t , a generic element of T . Together, the tuple (G, S) defines a game of incomplete information.⁵

Given a known basic game G , the designer chooses and publicly announces an information structure S , which becomes common knowledge. The agents then observe the choice S and the subsequent signal realizations. Depending on the choice of information structure, these signal realizations may be only privately observable, they may be common to different subsets of agents, or they might be public to everyone. Upon observing his signal realization, each agent formulates his first order and higher order beliefs taking into account the common knowledge of the information structure S . Then, each agent selects an action, which maximizes his interim expected utility. The resulting action profile is a Bayes Nash equilibrium (BNE) of the incomplete information game (G, S) at the interim level. The designer's problem is to choose an information structure which induces agents to play a BNE that maximizes her ex ante expected utility. That is, the designer selects among the BNE of (G, S) at the ex ante level, the one that is most beneficial to her. If there are multiple equilibria of (G, S) , we take a best-case approach and consider the one which yields the highest ex ante expected utility to the designer. If in turn there are multiple equilibria that maximize the designer's ex ante expected utility, we select arbitrarily among them. We follow Kamenica and Gentzkow (12) and take this best-case approach since it provides a meaningful benchmark in case of equilibrium multiplicity.

4.2 Designer's Problem

For a given basic game G , an information structure S induces a BNE of the incomplete information game (G, S) , which in turn determines a distribution over action profiles and states of the world. Hence, the designer's problem can be organized as follows: 1) Characterize the set of all BNE of G that could emerge under all possible information structures. We refer to this as the *constraint set* of the optimization problem. 2) Among all BNE, select (the) one which generates a distribution over actions and states that maximizes the designer's ex ante expected utility. We refer to the latter as the *objective function* of the designer's optimization. 3) Find the information structure which induces this BNE for the given basic game G . In this section we will show that steps 1) and 2) reduce to a linear programming problem. We will also show that without loss of generality we can focus on a particular class of information structures when approaching step 3).

⁵This division of a game of incomplete information into a basic game and an information structure has been previously used in the literature; see for example Bergemann and Morris (5) and Lehrer, Rosenberg, and Shmaya (13).

4.2.1 Constraint Set

To determine the constraint set, we need to characterize the set of all BNE that could emerge under all possible information structures, of which there are infinitely many. To accomplish this task, we use a definition of correlated equilibrium introduced by Bergemann and Morris (5). We show below that using their definition of correlated equilibrium under a special information structure, we can characterize the set of BNE that could emerge under all possible information structures. With this purpose in mind, we introduce a few definitions to establish the necessary terminology.

A (behavioral) strategy for player i in (G, S) is a mapping $\beta_i : T_i \rightarrow \Delta(A_i)$.

DEFINITION 1. (Bayes Nash Equilibrium)

A strategy profile β is a Bayes Nash equilibrium (BNE) of (G, S) if for each $i \in I$, $t_i \in T_i$ and $a_i \in A_i$ with $\beta_i(a_i|t_i) > 0$, we have

$$\begin{aligned} \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi(t_i, t_{-i} | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi(t_i, t_{-i} | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (4.1)$$

for all $a'_i \in A_i$.

We next state the definition of Bayes correlated equilibrium as introduced by Bergemann and Morris (5). Let $\sigma : T \times \Theta \rightarrow \Delta(A)$ be a distribution over action profiles conditional on type profiles and states.

DEFINITION 2. (Bayes Correlated Equilibrium)

A distribution σ is a Bayes correlated equilibrium (BCE) of (G, S) if for each $i \in I$, $t_i \in T_i$ and $a_i \in A_i$, we have

$$\begin{aligned} \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi(t_i, t_{-i} | \theta) \sigma((a_i, a_{-i}) | (t_i, t_{-i}), \theta) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi(t_i, t_{-i} | \theta) \sigma((a_i, a_{-i}) | (t_i, t_{-i}), \theta) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (4.2)$$

for all $a'_i \in A_i$.

A BCE distribution σ reflects the assumption of common certainty of rationality and the common prior assumption in the basic game G when the players have observed at least information structure S .

The designer is ultimately interested in what can be said about the equilibrium distributions of action profiles conditional on states of the world, as that determines the expected

value of her objective function. She is not interested in the distributions conditional on the signals, as the information structure is simply a tool and not the end goal. Therefore, she would like to find the most beneficial equilibrium distribution of actions conditional on states, which maximizes the expected value of her objective function, without assuming a specific information structure to start with. Let mapping $\nu : \Theta \rightarrow \Delta(A)$ be a distribution over action profiles conditional on states.

DEFINITION 3. *A distribution ν is a BNE of (G, S) if β is a BNE of (G, S) and*

$$\sum_{t \in T} \pi(t|\theta) \left(\prod_{j=1}^N \beta_j(a_j|t_j) \right) = \nu(a|\theta) \quad (4.3)$$

for each $a \in A$ and $\theta \in \Theta$. A distribution ν is a BCE of (G, S) if σ is a BCE of (G, S) and

$$\sum_{t \in T} \pi(t|\theta) \sigma(a|t, \theta) = \nu(a|\theta) \quad (4.4)$$

for each $a \in A$ and $\theta \in \Theta$.

Subsequently, for a basic game G and information structure S , we use $BNE(G, S)$ to denote the set of BNE distributions ν and $BCE(G, S)$ to denote the set of BCE distributions ν . In the designer's problem, the constraint set is the largest set of distributions ν that could emerge if agents play a BNE for a basic game G under any possible information structure. To characterize this set, we show that it is easier to work with the set of BCE for G under a particular information structure.

We next define the information structure, which plays an important role in the upcoming analysis. The *null* information structure \underline{S} has $\underline{T}_i = \{t_i\}$ for all i and $\underline{\pi}(t|\theta) = 1$ for all $\theta \in \Theta$. Thus, the null information structure $\underline{S} = (\underline{T}, \underline{\pi})$ provides no information at all about the state of the world. The next results established the characterization of the largest set of BNE distributions through its equivalence to the set of BCE under the null information structure.

PROPOSITION 1. *The following holds: $BCE(G, \underline{S}) = \cup_S BNE(G, S)$.*

The above result established the equivalence between the largest set of BNE distributions for a basic game G and the set of BCE distributions for G under the null information structure. It is a version of Theorem 2 by Bergemann and Morris (5). We will work with the constraints defining the set $BCE(G, \underline{S})$ to characterize the constraint set of the designer $\cup_S BNE(G, S)$.

Intuitively, the result can be interpreted as follows. In a BCE distribution, the correlation between the actions given the state is arbitrary. In a BNE distribution, the correlation between the actions given the state can be generated only through independent probability distributions of individual actions given signals according to (4.3). To generate every possible

distribution in $BCE(G, S)$ as a BNE, i.e. with behavioral strategies, the additional coordination with the state must come through the conditioning on the signals. Therefore, the information structure has to provide the necessary correlation of the independently chosen actions and the state. Every BCE distribution can thus be replicated as a BNE distribution for an appropriately chosen information structure S . S should provide enough information about the state to generate the required correlation in the equilibrium distribution. To summarize, a BCE distribution under \underline{S} can be viewed as a stochastic device which is sophisticated in terms of how much correlation it can generate between the actions, but does not use the information structure at all. A BNE distribution under S , on the other hand, can be viewed as a stochastic device which generates all the correlation between the actions through S . Therefore more intricate information structures are required for the latter to replicate any distribution of the former.

By Proposition 1 we can characterize the set of all BNE by means of the BCE incentive constraints (4.2) under the null information structure \underline{S} . We need to combine these with the constraints ensuring ν is a proper probability distribution. Hence, the set $BCE(G, \underline{S})$ is the collection of $\nu(a|\theta)$ such that:

- i) $\nu(a|\theta) \geq 0$ for all $a \in A$ and $\theta \in \Theta$,
- ii) $\sum_{a \in A} \nu(a|\theta) = 1$ for all $\theta \in \Theta$, and
- iii) $\sum_{a_{-i}, \theta} \psi(\theta) \nu((a_i, a_{-i})|\theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \nu((a_i', a_{-i})|\theta) u_i((a_i', a_{-i}), \theta)$ for all $i \in I$, $a_i \in A_i$ and $a_i' \in A_i$.

The above constraints are all linear in $\nu(a|\theta)$. Therefore, the set $BCE(G, \underline{S})$ is a convex polygon. By Theorem A of Stinchcombe (18), the set of BCE is non-empty. Hence, the constraint set of the designer is a non-empty convex polygon.

4.2.2 Objective Function

The designer's utility when the agents play action profile a and the state is θ is given by $V(a, \theta)$. The designer's objective is to maximize the ex ante expected value of her utility, which can be written as

$$\mathbb{E}_\nu[V] = \sum_{a, \theta} V(a, \theta) \nu(a|\theta) \psi(\theta).$$

Notice that this objective is also linear in $\nu(a|\theta)$. Hence, the designer is maximizing a linear objective function over a non-empty convex polygon and the tools of linear programming can be utilized to find the optimal solution. By the fundamental theorem of linear programming, a solution ν^* exists and is at one of the corners of the constraint set. Therefore, $\nu^* \in \cup_S BNE(G, S)$ is the BNE the designer would like to induce. We next character use the information structure S^* which supports ν^* as a BNE, i.e. for which $\nu^* \in BNE(G, S^*)$.

4.2.3 Optimal Information Structure

We first simplify the problem by showing that, without loss of generality, we can restrict attention to a certain class of information structures, which we call *direct*.

DEFINITION 4. *Given a basic game G , an information structure $S = (T, \pi)$ is direct if $T = A$ and there exists $\nu \in BNE(G, S)$ such that $\nu(a|\theta) = \pi(a|\theta)$ for all $a \in A$ and $\theta \in \Theta$.*

DEFINITION 5. *Given basic game G , we say that an information structure S has value \tilde{V} if there exists a distribution $\nu \in BNE(G, S)$ such that $\mathbb{E}_\nu[V] = \tilde{V}$.*

PROPOSITION 2. *The following are equivalent:*

- (i) *There exists an information structure with value V^* ;*
- (ii) *There exists a direct information structure with value V^* ;*
- (iii) *There exists a BNE distribution ν such that $\mathbb{E}_\nu[V] = V^*$.*

The main implication of Proposition 2 is that we can work with direct information structures only. The equivalence of (i) and (ii) is in spirit very similar to the revelation principle (e.g., Myerson (16)). The equivalence between (ii) and (iii) uses Proposition 1 and a *truthful* equilibrium strategy. The intuition behind this is simple. If there is a BNE distribution ν over action profiles conditional on states, then it must be that ν is also a BCE distribution under the null information structure. Thus, if the designer uses a direct information structure with the same probability distribution ν , it is Bayes incentive compatible for each agent to follow the action recommendation implied by the observed signal realization assuming that the other agents do so as well. This generates an BNE equilibrium distribution ν under a direct information structure, which in turn results in the same ex ante expected payoff for the designer.

COROLLARY 1. *The optimal information structure is given by $S^* = (A, \pi^*)$, where $\pi^*(a|\theta) = \nu^*(a|\theta)$ and $\nu^* = \arg \max_\nu \mathbb{E}_\nu[V]$ s.t. $\nu \in BCE(G, \underline{S})$.*

This corollary establishes the equivalence between the optimal information structure and the optimal BCE distribution under the null information structure. Once we find ν^* , we create a direct information structure with the same probability distribution over signal realizations conditional on states. The signal realizations are in fact the action recommendations, which agents have incentive to follow in equilibrium. This direct information structure is optimal.

In our setting, there is nothing that precludes a designer, who has chosen a partially informative information structure, from deciding to release more information after certain “unfavorable” signal realizations. The following result establishes that regardless of the actual signal realizations, the designer would never want to deviate and generate additional signals

(drawn from the same or from a new distribution). That is, if the initial information structure has been optimally chosen to begin with and if all players have observed the realization of the recommended action profile \underline{a} , the designer cannot do better by releasing more information in the form of new draws of signals. That is, the designer cannot benefit from persuading the players to switch to a different action profile by giving them more information.

PROPOSITION 3. *If a realized action profile $\underline{a} \in A$ observed by all players has been induced by an optimal signal, the designer has no incentive to release more information.*

5 Application: Symmetric Binary Environments

In this section we apply the general information design approach outlined above to a symmetric binary environment.

5.1 Setup

Consider a two-player, two-state, two-action environment with symmetric payoffs. There are $N = 2$ players, and we use i as an index for the typical player, and j for his opponent. The set of states of the world is $\Theta = \{\theta_0, \theta_1\}$. The set of actions is the same for both players and given by $A = \{a_0, a_1\}$. The payoffs (or utility functions) $u : A \times \Theta \rightarrow \mathbb{R}$ are also the same for both players. Further, we assume a common prior ψ , which is uniform on the two states, i.e. $\psi(\theta_0) = \psi(\theta_1) = \frac{1}{2}$. Hence, we have specified the basic game $G = (A^2, u, \psi)$. We will refer to this as the *symmetric $2 \times 2 \times 2$ environment*.

Consider the following parameterized framework, where the payoffs in each state are given by: with $c \geq 0$ and $d \geq 0$. The assumption that the payoff parameters are weakly positive

$\theta = \theta_0$	a_0	a_1	$\theta = \theta_1$	a_0	a_1
a_0	c, c	$d, 0$	a_0	$0, 0$	$0, d$
a_1	$0, d$	$0, 0$	a_1	$d, 0$	c, c

Table 1: Parameterized Basic Game

ensures that the participation constraints of the agents to engage in the strategic interaction are always satisfied. This two-parameter representation is rich enough to capture many different environments of interest. We will refer to the basic game with parameters c and d as $G_{c,d}$.

The above payoff matrices assume that players have a preference for playing different actions in the different states of the world. This is an important assumption. Notice that if the same action were preferred in both states, there would be a dominant strategy equilibrium. In this case, the information that players receive is irrelevant for their strategies, and the designer cannot use information design to achieve her desired objective. Hence, information design

becomes relevant only when the players have preferences for coordinating each action with a different state. We denote by a_k the action preferred in state θ_k for $k = 0, 1$. Additionally, we use superscript to signify the agent that takes the action, i.e. a_k^i stands for agent i taking action a_k .

In addition to the preference for aligning their action with the state, the players may exhibit either a preference for coordination (strategic complementarity) or mis-coordination (strategic substitutability) of their action with the action of their opponent. The strength of the preference for alignment with the state versus alignment with one's opponent depends on the relative magnitude of c and d .

The preference of each player for coordination with the state, for any given action of the other player, is represented by $c + d$, which we will refer to as the *unilateral complementarity* (U). This is given by the difference:

$$U = u(a_1, a^j, \theta_1) - u(a_0, a^j, \theta_1) - u(a_1, a^j, \theta_0) + u(a_0, a^j, \theta_0) = c + d \quad (5.1)$$

for each $a^j \in A$. Due to the symmetry, we obtain the same expression for each player and each possible opponent action. The larger (5.1), the stronger the preference for alignment between each player's own action and the state.

In each state, the preference of each player for coordination with the other player is captured by $c - d$. This is given by:

$$T = u(a_1, a_1, \theta_k) - u(a_0, a_1, \theta_k) - u(a_1, a_0, \theta_k) + u(a_0, a_0, \theta_k) = c - d \quad (5.2)$$

for $k = 0, 1$. We will refer to this as the *strategic complementarity* (T). If this difference is positive and large, there is a strong preference for coordination with one's opponent, that is, strong strategic complementarity. On the other hand, if this difference is negative and large, there is a strong preference for mis-coordination between the players and thus, strong strategic substitutability. Consequently, we say that the basic game $G_{c,d}$ exhibits *strategic complements* if $c > d$ and *strategic substitutes* if $c < d$.

This two-parameter payoff representation captures many strategic interactions of interest and different preferences for (mis)coordination. For example, $c > d > 0$ represents the beauty contest game: players want to match the state and have an added benefit if their actions match. This may correspond to a situation of two people deciding to invest in one of two projects. The profitability of the projects depends on an unknown state and on the total investment, with higher investment leading to a more profitable project. Therefore, choosing the right project is associated with a higher payoff if the opponent also invests in the same project. When $d > c > 0$, the payoffs represent the situation of two competitors trying to match the consumer preference for a certain product. If they both match it, they split the market. However, if one of them fails to produce the product with desired features, while

the other one succeeds, then the second firm captures the whole market and obtains a higher payoff.

For any value T of the strategic complementarity (5.2) and any value $U \geq |T|$ of the unilateral complementarity (5.1), we can choose payoff parameters c and d to yield these coordination preferences by setting $c = \frac{U+T}{2}$ and $d = \frac{U-T}{2}$.

5.2 Designer's Problem

5.2.1 Constraint Set

To determine the constraint set we need to characterize the set of all possible BNE for basic game $G_{c,d}$ under all possible information structures. By Proposition 1 we know that for a given basic game G , the largest set of distributions over actions and states of the world, which can be sustained as BNE under some information structure, is given by $BCE(G, \underline{S})$. We restrict attention to distributions which are symmetric both in terms of the players and with respect to the state. These can be fully described by two parameters — q and r — and denoted as $\nu(q, r)$. Hence, a symmetric distribution over action profiles conditional on state can be represented as follows

$\theta = \theta_0$	a_0	a_1	$\theta = \theta_1$	a_0	a_1
a_0	r	$q - r$	a_0	$1 - 2q + r$	$q - r$
a_1	$q - r$	$1 - 2q + r$	a_1	$q - r$	r

The parameter r represents the probability with which in each state both agents simultaneously match the state with their actions: $\Pr(a_0, a_0 | \theta_0) = \Pr(a_1, a_1 | \theta_1) = r$. Hence, it measures the likelihood with which the players coordinate both with each other and with the state. On the other hand, q denotes the probability with which in each state each agent matches the state with his action, irrespective of whether the other agent does so as well or not. For agent i and state θ_0 , this probability is given by $\Pr(a_0^i, a_0^j | \theta_0) + \Pr(a_0^i, a_1^j | \theta_0) = q$.

We choose to work with symmetric distributions for a number of reasons. First, this is without loss of generality when the utility function of the designer is also symmetric in agents and in states — we will define explicitly what this means in the following subsection. However, in general, this will be a constrained optimal BNE distribution. Second, we use symmetric distributions as that allows for a two-dimensional graphical representation of the constraint set and objective function. Third, sometimes the designer is naturally constrained in her choice to symmetric information structures due to laws and regulations. In the symmetric $2 \times 2 \times 2$ environment that means she is optimizing over the set of symmetric BNE, which are induced by symmetric information structures.

Our next result characterizes the set of symmetric BCE of $(G_{c,d}, \underline{S})$. We consider all possible values of the basic game parameters c and d which do not make the strategic inter-

action trivial. In other words, we consider all possible cases with $c, d \geq 0$ and for which both parameters are not simultaneously equal to zero.

PROPOSITION 4. (BCE Distributions)

Consider the symmetric $2 \times 2 \times 2$ environment.

If $c > d$ (strategic complements), the set of symmetric BCE distributions of $(G_{c,d}, \underline{S})$ is given by $\left\{ (q, r) \in Co \left\{ \left(\frac{d}{c+d}, \frac{d}{c+d} \right), \left(\frac{2c-d}{3c-d}, \frac{c-d}{3c-d} \right), (1, 1) \right\} \right\}$.

If $d > c$ (strategic substitutes), the set of symmetric BCE distributions of $(G_{c,d}, \underline{S})$ is given by $\left\{ (q, r) \in Co \left\{ \left(\frac{d}{c+d}, \frac{d}{c+d} \right), (1, 1), \left(\frac{d}{3d-c}, 0 \right), \left(\frac{1}{2}, 0 \right) \right\} \right\}$.

If $d = c > 0$, the set of symmetric BCE distributions of $(G_{c,d}, \underline{S})$ is given by $\left\{ (q, r) \in Co \left\{ \left(\frac{1}{2}, \frac{1}{2} \right), \left(\frac{1}{2}, 0 \right), (1, 1) \right\} \right\}$.

The proof of the proposition shows that the set of BCE distributions under the null information structure for a basic game $G_{c,d}$ is the constraint set determined by four linear inequalities. Three of these inequalities ensure that the parameters of the distribution satisfy the consistency conditions for probability distributions. The fourth inequality represents the incentive constraints associated with BCE under the null information structure.

We make use of the following example to show the construction of the constraint set. We will return to this example throughout the rest of the section to illustrate the different steps of the information design problem.

SYMMETRIC EXAMPLE. Consider the parameterized basic game in Table 1 with $c = 2$ and $d = 1$. Hence, the agents are involved in a coordination game, where they want to both match each other and the state with their actions. Suppose the designer benefits from mis-coordination between the agents' actions irrespective of the state. That is, her utility function is given by:

$$V(a^i, a^j, \theta) = \begin{cases} 1 & \text{if } a^i \neq a^j \\ 0 & \text{otherwise.} \end{cases} \tag{5.3}$$

The constraint set $BCE(G_{2,1}, \underline{S})$ is depicted in Figure 1. The red line represents the BCE incentive constraint. It always goes through the point $(\frac{1}{2}, \frac{1}{4})$, plotted on the graph, which represents the symmetric mixed strategy BNE when the agents have no information.

5.2.2 Objective function

We consider a general utility function for the designer $V : A \times \Theta \rightarrow \mathbb{R}$. Hence, $V(a, \theta)$ is the designer payoff when a is the action profile played by the agents and θ is the state of the world. For any given objective function $V(a, \theta)$, where a is the action profile played by the

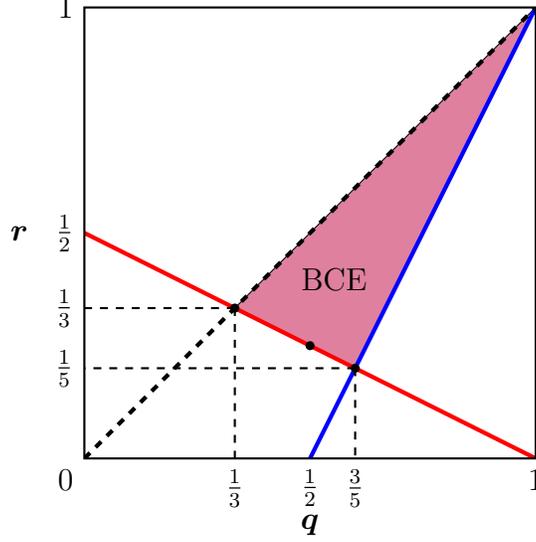


Figure 1: Constraint Set (Symmetric Example)

agents and θ is the realized state of the world, its expectation given a BNE distribution $\nu(q, r)$ is

$$\begin{aligned} \mathbb{E}(V) = & \psi [rV(a_0, a_0, \theta_0) + (q-r)V(a_0, a_1, \theta_0) + (q-r)V(a_1, a_0, \theta_0) + (1-2q+r)V(a_1, a_1, \theta_0)] \\ & + (1-\psi) [rV(a_1, a_1, \theta_1) + (q-r)V(a_0, a_1, \theta_1) + (q-r)V(a_1, a_0, \theta_1) + (1-2q+r)V(a_0, a_0, \theta_1)]. \end{aligned} \quad (5.4)$$

Reorganizing and regrouping terms gets us to:

$$\begin{aligned} \mathbb{E}(V) = & \left[\psi [V(a_1, a_1, \theta_0) - V(a_0, a_1, \theta_0) - V(a_1, a_0, \theta_0) + V(a_0, a_0, \theta_0)] \right. \\ & \left. + (1-\psi) [V(a_1, a_1, \theta_1) - V(a_0, a_1, \theta_1) - V(a_1, a_0, \theta_1) + V(a_0, a_0, \theta_1)] \right] r \\ & + \left[\psi [V(a_0, a_1, \theta_0) - V(a_1, a_1, \theta_0) + V(a_1, a_0, \theta_0) - V(a_1, a_1, \theta_0)] \right. \\ & \left. + (1-\psi) [V(a_0, a_1, \theta_1) - V(a_0, a_0, \theta_1) + V(a_1, a_0, \theta_1) - V(a_0, a_0, \theta_1)] \right] q \\ & + \psi V(a_1, a_1, \theta_0) + (1-\psi)V(a_0, a_0, \theta_1) \\ = & R \cdot r + Q \cdot q + \text{const.} \end{aligned} \quad (5.5)$$

The coefficient in front of r , which we denote by R , captures the “expected” preference for complementarity between the actions in the designer’s objective function. It is indeed a weighted average of the complementarities between the actions in each state, the weights being the prior probabilities for each state. Therefore, the coefficient in front of r measures the average importance of coordination of the agents’ actions in the designer’s objective function.

The coefficient in front of q , which we label Q , is the expected preference for *unilateral*

coordination of each player’s action with the state, assuming the other player mismatches the state. For example, suppose that the state is θ_0 . Then $V(a_0, a_1, \theta_0) - V(a_1, a_1, \theta_0)$ captures the benefit of having the first player unilaterally match the state with his action as opposed to having perfect mis-coordination between both of the actions and the state. For the second player, the relevant expression is $V(a_1, a_0, \theta_0) - V(a_1, a_1, \theta_0)$. So the sum of those two expressions represents the preference of the designer for “unilateral” coordination between the players and the state θ_0 . Therefore, the coefficient in front of q measures the importance of unilateral coordination in the designer’s utility in expectation over the two states.

Lastly, a utility function that is symmetric in both the agents’ actions and the state is characterized by the following equalities: (i) $V(a_0, a_0, \theta_0) = V(a_1, a_1, \theta_1)$, (ii) $V(a_0, a_1, \theta_0) = V(a_1, a_0, \theta_0) = V(a_0, a_1, \theta_1) = V(a_1, a_0, \theta_1)$ and (iii) $V(a_1, a_1, \theta_0) = V(a_0, a_0, \theta_1)$. As mentioned above, when the designer’s utility function is symmetric, restricting attention to symmetric information structures is without loss of generality.

SYMMETRIC EXAMPLE. *The utility function of the designer given by (5.3) is symmetric both in the agents’ actions and in the state. Substituting the values into (5.5), gives*

$$\mathbb{E}(V) = -2r + 2q \tag{5.6}$$

as the ex ante expected objective function. This is represented by a level line with a slope of one, the value of which increases when shifted in the direction of the lower-right corner (see Figure 3).

5.2.3 Optimal Information Structure

In the previous section, we characterized the set $BCE(G_{c,d}, \underline{S})$, which is the budget set of the designer. We can now maximize the designer’s objective function (5.5) over this set. Let us denote by $\nu^*(q, r)$ the distribution which maximizes (5.5) over $BCE(G_{c,d}, \underline{S})$. Once we find $\nu^*(q, r)$, we can reverse-engineer the information structure S^* which decentralizes it as a BNE. By Proposition 2 we know that there exists a direct information structure S^* such that $\nu^*(q, r) \in BNE(G_{c,d}, S^*)$. And by Corollary 1 we know that $S^* = (A, \pi^*)$ with $\pi^*(a|\theta) = \nu^*(a|\theta)$ for all $a \in A$ and $\theta \in \Theta$.

Therefore, the direct information structures which support all distributions $\nu(q, r) \in BCE(G_{c,d}, \underline{S})$ as BNE, can be parameterized in an analogous way with the following conditional probabilities $\pi(\cdot|\theta)$ on signal realizations:

$\theta = \theta_0$	a_0	a_1	$\theta = \theta_1$	a_0	a_1
a_0	r	$q - r$	a_0	$1 - 2q + r$	$q - r$
a_1	$q - r$	$1 - 2q + r$	a_1	$q - r$	r

Table 2: Direct Information Structures

The information structure parameterization in Table 2 is very general as it represents all binary information structures which are symmetric across agents and states. The parameter q is the probability with which each agent receives the action recommendation that “matches” the state, i.e. the “state-matching” action. We refer to it as the *precision* of the information structure. The parameter r is the probability with which both agents simultaneously receive the state-matching action recommendation. We refer to it as the *correlation* of the information structure.

The above parameterization also includes many important special structures. The case of conditionally independent private signals is captured by setting $r = q^2$ for $q \in (0, 1)$. In this case, each agent receives a private signal which is equal to the state-matching action with probability q and is independent of the signal of his opponent. Both agents thus receive the state-matching action recommendation with probability $q \times q = r$ and receive opposite action recommendations with probability $q \times (1 - q) = q - r$. On the other hand, the case of public signals is covered by setting $r = q$. This ensures that both agents always receive the same action recommendation, where q is the probability of having that action match the state. We denote public signals by $S_{q,q}$. For the general case of *private signals* with precision q and correlation r , we write $S_{q,r}$.

Of particular importance is the *null information* structure \underline{S} which provides no information about the state θ . In terms of the above parameterization, the null information structure corresponds to $q = \frac{1}{2}$ and can be denoted as $S_{\frac{1}{2},r}$. In this case, the signals are completely uninformative with respect to the state. Notice also that there are infinitely many null information structures, each one associated with a different degree of correlation between the signals. On the other hand, there is only one *full information* structure \bar{S} which reveals the state of the world perfectly, captured by $q = r = 1$ and written as $S_{1,1}$.

It is useful for the upcoming analysis to graphically represent the set of direct binary information structures in the (q, r) -space (Figure 2). For the conditional probabilities in Table 2 to be positive, we need to have r smaller than q , greater than $2q - 1$ and greater than 0. Thus, the set of possible direct information structures is defined by three linear constraints. The first line is $r = q$, which describes the set of all public signals with different levels of precision. The second line is $r = 2q - 1$, which represents all information structures with minimal levels of correlation for a given level of precision q . And the third line is $r = 0$, which corresponds to all information structures with zero correlation consistent with different levels of precision. That is why these three lines determine the set of direct information structures. It is easy to see that the set of conditionally independent signals, $r = q^2$ with $q \in (0, 1)$ is in the interior of the set of all possible information structures.

Before we move on to the complete characterization, let us demonstrate graphically how we obtain the optimal information structure in the symmetric example we have been using throughout this section. This is shown in Figure 3.

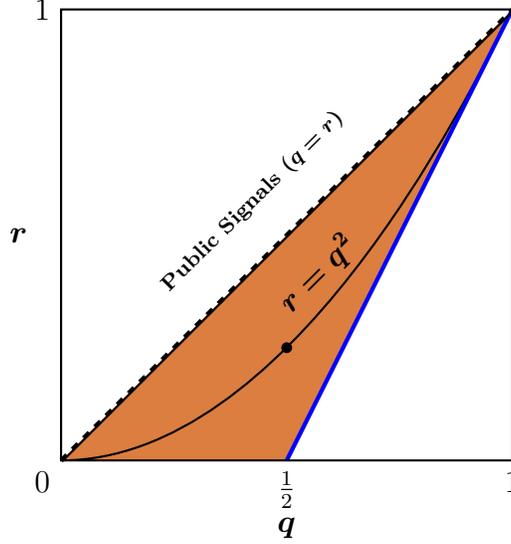


Figure 2: Set of Direct Information Structures

SYMMETRIC EXAMPLE. *The symmetric BCE which maximizes the expectation of the objective function is $\nu^*(\frac{3}{5}, \frac{1}{5})$. The optimal direct information structure is thus given by $S^* = (A, \nu^*)$ and is summarized in the following matrices:*

$\theta = \theta_0$	a_0	a_1
a_0	$\frac{1}{5}$	$\frac{2}{5}$
a_1	$\frac{2}{5}$	0

$\theta = \theta_1$	a_0	a_1
a_0	0	$\frac{2}{5}$
a_1	$\frac{2}{5}$	$\frac{1}{5}$

Under this information structure, the expected value of the designer's objective function is $\frac{4}{5}$. Due to the symmetry of the binary environment and of the designer's utility function, this information structure is a global optimum. In other words, restricting attention to symmetric information structures is, in this case, without loss of generality.

Our next result is a complete characterization of the optimal symmetric information structure for all possible designer's objective functions and basic games $G_{c,d}$. Recall that R and Q are defined as the average preference for coordination of the agents' actions and the average preference for unilateral coordination of each player with the state, respectively (see Section 5.2).

THEOREM 1.

1. *If $R > 0$ and $Q > 0$, the full information structure is always optimal.*
2. *If $R < 0$, $Q > 0$ and the basic game exhibits strategic complements, the optimal information structure is public signals with precision $\frac{d}{c+d}$ if $-\frac{Q}{R} < \frac{c-3d}{2(c-d)}$; private signals with precision $\frac{2c-d}{3c-d}$ and correlation $\frac{c-d}{3c-d}$ if $\frac{c-3d}{2(c-d)} < -\frac{Q}{R} < 2$; and the full information structure if $-\frac{Q}{R} > 2$.*

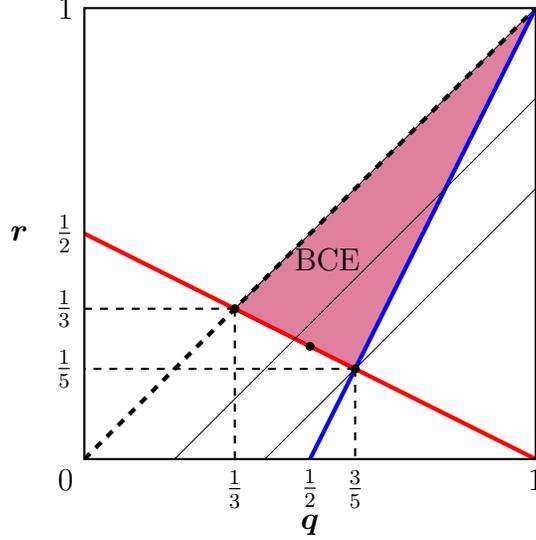


Figure 3: Optimal Information Structure (Symmetric Example)

3. If $R < 0$, $Q > 0$ and the basic game exhibits strategic substitutes, the optimal information structure is the null information structure if $-\frac{Q}{R} < 2$; and the full information structure if $-\frac{Q}{R} > 2$.
4. If $R > 0$, $Q < 0$ and the basic game exhibits strategic complements, the optimal information structure is the full information structure if $-\frac{Q}{R} < 1$; and private signals with precision $\frac{2c-d}{3c-d}$ and correlation $\frac{c-d}{3c-d}$ if $-\frac{Q}{R} > 1$.
5. If $R > 0$, $Q < 0$ and the basic game exhibits strategic substitutes, the optimal information structure is the full information structure if $-\frac{Q}{R} < 1$; public signals with precision $\frac{d}{c+d}$ if $1 < -\frac{Q}{R} < \frac{c-3d}{2(c-d)}$; and private signals with precision $\frac{d}{3d-c}$ and correlation 0 if $-\frac{Q}{R} > \frac{c-3d}{2(c-d)}$.
6. If $R < 0$, $Q < 0$ and the basic game exhibits strategic complements, the optimal information structure is public signals with precision $\frac{d}{c+d}$ if $-\frac{Q}{R} < \frac{c-3d}{2(c-d)}$; and private signals with precision $\frac{2c-d}{3c-d}$ and correlation $\frac{c-d}{3c-d}$ if $-\frac{Q}{R} > \frac{c-3d}{2(c-d)}$.
7. If $R < 0$, $Q < 0$ and the basic game exhibits strategic substitutes, the optimal information structure is private signals with precision $\frac{d}{3d-c}$ and correlation 0.

Our characterization theorem is summarized in Table 3 of Appendix B. The information design problem can be seen as utility maximization given the designer's preferences over distributions of actions conditional on the states, where the budget set is $BCE(G_{c,d}, \underline{S})$. The slope of the designer's level line, $-\frac{Q}{R}$, can be viewed as a marginal rate of substitution. It represents the designer's benefit from an increase in the probability (q) of state coordination relative to

the benefit from an increase in the probability (r) of action coordination. When this slope is negative, the designer benefits from simultaneous movements in both probabilities and so is willing to trade in an increase in one for a decrease in the other. When it is positive, however, the designer is willing to trade between simultaneous movements: an increase (decrease) in one for an increase (decrease) in the other. Thus, the slope of the designer level line represents the tradeoff that she is willing to accept between the two parameters of the equilibrium distributions.

Since $BCE(G_{c,d}, \underline{S})$ is the budget set of the information designer, the slopes of its boundaries represent the tradeoffs between the parameters q and r that need to be maintained so that the BCE incentive compatibility constraints remain satisfied. Therefore, these slopes represent the rates at which the designer may trade changes in one parameter for changes in the other. We next explain the intuition in a few cases.

Consider the case when the designer would like the agents to both coordinate their actions with each other ($R > 0$) and with the state ($Q > 0$). If the game has strategic complements, then the agents have the exact same preferences as the designer. Therefore, it is best to give the agents full information, as they will use it to coordinate perfectly with each other and the state, which is the objective of the designer. Notice that for the same preferences of the designer, the optimal information structure is full information also when the game has strategic substitutes, which is somewhat counterintuitive. The reason is that once the agents have full information, it is always a dominant strategy to play the action that matches the state, because the strategic substitutes are not strong enough. This is due to the fact that we restrict the payoff parameters to be strictly positive, which is a limitation of the model. If we were to allow for c to be negative, this result would change. However, in the current setting, a designer who would like the agents to coordinate with each other and with the state achieves that by giving full information both when the game exhibits strategic complements and strategic substitutes.

Next, consider the case of a designer with preferences described by $R < 0$ and $Q > 0$. She would hence like to choose q to be as high as possible and r to be as low as possible. This translates into a preference for mis-coordination between the agents irrespective of the state, as it maximizes the probability of mismatched actions $\Pr(a_0, a_1) = \Pr(a_1, a_0) = q - r$. If the game has strategic substitutes, there is an underlying incentive for the agents to mismatch their actions. However, each one of them still has an incentive to match the state and the mismatched action profile is never a full information Nash equilibrium. If the preference for coordination with the state is not as strong as the disutility from coordination between the agents, then the designer would choose to reveal no information. This will maximize the probability of mis-coordinated actions. Conversely, if the preference for coordination with the state is stronger than the disutility from coordination between the actions, then it is optimal for the designer to reveal everything. The actions will never be mis-coordinated in this case;

nonetheless, the designer would prefer to have perfect state coordination.

For the case of $R > 0$ and $Q < 0$, the designer wants the agents to coordinate their actions but to not coordinate with the state. If the game exhibits strategic complements, the agents would like to both coordinate with each other and with the state. Therefore, if the designer's preference for action coordination is stronger than the disutility from state coordination, she will choose the full information structure and have them coordinate on both actions and the state. In contrast, if she really dislikes state coordination, she will choose correlated private signals with imperfect precision. By doing this, the designer foregoes the perfect action coordination she could achieve with full information in order to achieve some degree of state mis-coordination. Thus, depending on the strength of those two preferences, the outcome is either full information or correlated private signals with imperfect precision.

Lastly, consider the case of $R < 0$ and $Q < 0$, where the designer wants the agents to mis-coordinate both on the state and the actions. If the game exhibits strategic substitutes, the optimal information structure is private signals with low precision. This ensures that the agents do not obtain enough information about the state, so that they don't coordinate too much with it. At the same time, they have enough information about what the other player has likely observed. Since the game has strategic substitutes, the agents have an incentive to mis-coordinate their actions. Hence the signals are used mainly as a device for the agents to condition their actions on, in order to achieve mis-coordination.

To contrast our results with those from the literature on cheap talk without commitment, we would like to point out that making the preferences of the designer and the agents more aligned, may in fact decrease the optimal precision of information. For example, when we have strategic substitutes, and the preferences of the designer are $R > 0, Q > 0$, full information is optimal. For strategic substitutes and $R < 0, Q > 0$, the designer also wants action mis-coordination and state coordination, just like the agents. So the preferences have become more aligned. However, the null structure is optimal in this case for certain values of the parameters.

COROLLARY 2. *Conditionally independent private signals are never optimal.*

Notice that $r = q^2$, $q \in (0, 1)$ is always in the interior of the set of information structures and is therefore never optimal.

COROLLARY 3. *The only case when the designer may not benefit from information design is when $R < 0$, $Q > 0$, $-\frac{Q}{R} < 2$ and the basic game exhibits strategic substitutes.*

There is no benefit from information design whenever revealing no information and letting the agents operate under their prior beliefs is optimal. Thus, as long as the optimal information structure differs from the null, the designer benefits from information design. No information revelation, i.e. $q = 1/2$, is only ever strictly optimal in the case of strategic substitutes when

the designer has preferences described by $R < 0$, $Q > 0$ and $-\frac{Q}{R} < 2$. The intuition behind this case was described above. In all remaining cases, information design is beneficial.

5.2.4 Indirect Information Structures

In some cases, rather than sending direct action recommendations, the designer may prefer or be confined to using signals, which are intrinsically associated to varying degrees with the different states of the world.⁶ This is the second class of information structures we consider, which we call “indirect”.

When the designer is not able to send direct action recommendations to create the information structure, the interpretation of the signals becomes relevant. In a way, this can be viewed as using a predetermined “language” to create the information structure. In binary environments, the designer needs two different signals to generate the set of information structures that can support all BNE. In order to generate symmetric information structures, one of the signals has to be designated as more indicative of state θ_0 , and the other signal — as more indicative of state θ_1 . Let us denote the former by t_0 and the latter by t_1 . Therefore, conditional on θ_0 (θ_1) the probability that signal t_0 (t_1) is observed has to be at least $\frac{1}{2}$ i.e. $\Pr(t_0|\theta_0) = \Pr(t_1|\theta_1) = r + (q - r) = q$ has to be weakly greater than $\frac{1}{2}$. Otherwise the signals would not be indicative of the states.

We can use the same parameterization as in Table 2, only instead of the action recommendations a_0 and a_1 , we use the signals t_0 and t_1 respectively. Most importantly, we now have the added restriction that the precision q has to be weakly larger than $\frac{1}{2}$. Thus, the set of indirect information structures is smaller than the set of direct information structures due to this additional constraint $q \geq \frac{1}{2}$. Figure 4 depicts the set of indirect information structures.

Every $\nu^*(q, r) \in BCE(G_{c,d}, \underline{S})$ can be thus supported as a BNE by an indirect information structure $S_{q,r}$ as long as $q \geq \frac{1}{2}$. In this case, following the signal and playing a_0 when t_0 is received, and a_1 when t_1 is received, is a BNE. This can be viewed as “truthtelling” under indirect information structures.

However, for $\nu^*(q, r) \in BCE(G_{c,d}, \underline{S})$ with $q < \frac{1}{2}$ we can no longer use an indirect information structure $S_{q,r}$, as it is not defined for precision values less than a half. We need to find an indirect information structure, which will decentralize $\nu^*(q, r)$ as a BNE. The next proposition establishes the information structure that accomplishes this.

PROPOSITION 5. *For $q < \frac{1}{2}$, $\nu^*(q, r) \in BCE(G_{c,d}, \underline{S})$ only if $\nu^*(q, r) \in BNE(G_{c,d}, S_{1-q, 1-2q+r})$.*

The indirect information structure which decentralizes a BCE $\nu^*(q, r)$ with $q < \frac{1}{2}$ as a BNE is given by $S_{1-q, 1-2q+r}$. The intuition is that the designer creates an indirect structure that is a “mirror image” of the direct information structure she would have used. This is necessary

⁶The chairman of the Fed does not typically talk about how economic agents should be behaving. Rather, his statements include signals regarding the Fed’s stimulus policy and the economic outlook.

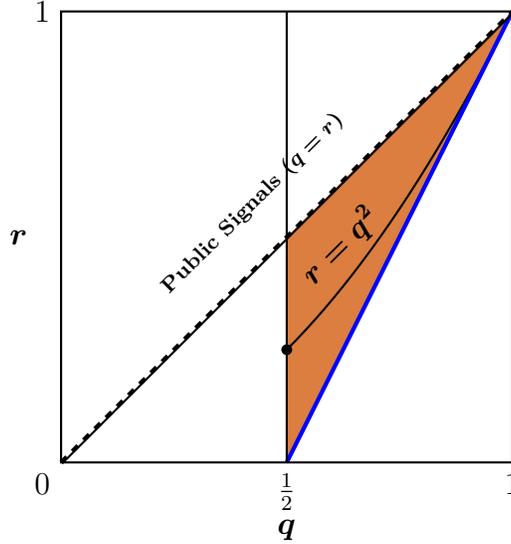


Figure 4: Set of Indirect Information Structures

so that the precision of the indirect structure is greater than $\frac{1}{2}$. Under $S_{1-q,1-2q+r}$, it is a BNE for both players to play the opposite action of what the signal suggests; that is, play a_1 if t_0 is received, and play a_0 if t_1 is received. This BNE results in a conditional probability distribution which is exactly equivalent to $\nu^*(q, r)$. Since the information structure $S_{1-q,1-2q+r}$ is a mirror image of $\nu^*(q, r)$ and the BNE strategy we consider is in turn a mirror image of truthtelling, the resulting distribution over actions conditional on states of the world is exactly the desired distribution $\nu^*(q, r)$.

5.3 Mechanism and Information Design

Thus far we maintained the assumption that the payoffs of the underlying game — the parameters c and d — were constant. In this section we offer insights into how changes in these parameters affect the maximal payoff of the information designer. This change in parameter values may come about due to exogenous factors. Another possibility is if the designer uses state and action contingent transfers to modify the payoffs of the agents. This is feasible if the state of the world becomes observable after the game has been played and if the actions of the agents are verifiable. Angeletos and Pavan (3) use taxes contingent on ex post public information about the realized state and aggregate activity. They show that such policies can improve the equilibrium use of information.

The main idea of this section is that in some cases the designer can induce an even better equilibrium outcome by combining the tools of mechanism design (payoff modification) and information design (belief modification), than when she uses only one or the other. Changes in the payoff parameters through state and payoff contingent transfers can affect the maximal utility of the information designer. An increase in c , while d is held constant, increases both

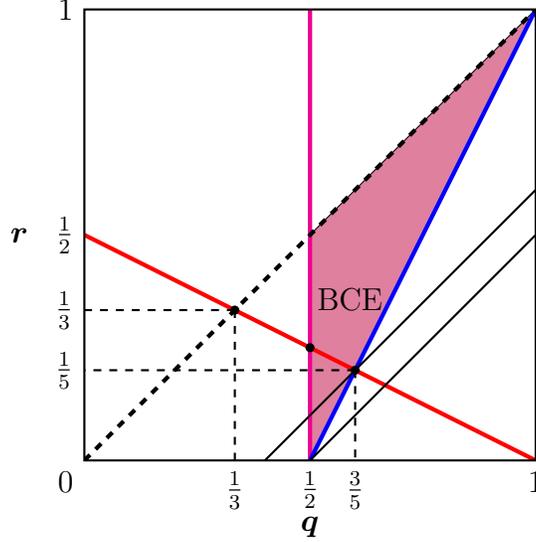


Figure 5: Changing Payoffs (Symmetric Example)

the complementarities with the state and between the agents. In contrast, an increase in d , while c is held constant, increases the complementarities with the state, while decreasing the complementarities between the agents. These changes have different effects on the maximum utility of the designer, depending on her preferences and also on the extent of the changes.

SYMMETRIC EXAMPLE. *We next show that by reducing c from 2 to 1, while holding $d = 1$ (for example, by introducing a state-contingent tax whenever both players coordinate), increases the maximal expected payoff for the designer. The new constraint set and the new optimal information structure ($q = 1/2$, $r = 0$) are illustrated in Figure 5, together with the higher level curve achievable for the designer, who now has an expected payoff of 1.*

It is not always possible to draw general conclusions from the comparative statics analysis. Frequently, if the change in the parameter is substantial enough, it may cause a shift from the case of strategic complements to the case of strategic substitutes. When this happens, the comparison is very sensitive to the size of the shift and we cannot evaluate the changes based only on its direction. In some instances, however, we are able to make clear-cut conclusions regarding the direction of the effects. The next proposition established the cases for which this is possible and for which mechanism design can improve on the outcome achieved with information design.

PROPOSITION 6. *Holding all else equal,*

1. *and starting with $c > d$, an increase in c or a decrease in d is always beneficial to a utility maximizing designer with $R < 0$, $Q < 0$, who uses public signals.*
2. *and starting with $c < d$, a decrease in c or an increase in d is always beneficial to a utility*

maximizing designer with $R > 0$, $Q < 0$ who uses private signals and to a designer with $R < 0$, $Q < 0$.

We next explain the intuition behind the last case: $R < 0$, $Q < 0$ and $c < d$. If the designer wants players to mis-coordinate both with the state and with each other, and the game exhibits strategic substitutes, the optimal information structure is private signals with low precision. This allows the players to use the signals in order to mis-match their actions, while obtaining very little information about the state. The signals need to have some level of precision in order for the players to pay attention to them and utilize them when formulating their strategies. This comes to the designer at the cost of the agents being able to predict the state better. When c decreases, or d increases, the level of strategic substitutability decreases. Therefore, the incentives of the agents for mis-coordinating their actions become stronger. They will now need less precise signals about the state to achieve this. This is beneficial to the designer as the less precise signals also result in less coordination with the state.

6 Discussion and Extensions

6.1 Communication

In certain environments it is unreasonable to assume that agents will not share the information they observe with each other, if that is beneficial to them. When the underlying game exhibits strategic complements, this issue is of particular relevance. The agents want to match the state and each other with their actions. Therefore, each agent has an incentive to disclose the private signal he observes to his opponent. The benefit from doing so is twofold. First, the agents can coordinate their actions perfectly once they have the same information. Second, by sharing the signal realizations they have observed, they can improve the precision of their information and update their beliefs accordingly.

A designer who faces a situation where the agents have incentives to communicate with each other, needs to take that into account when designing the information structure. This implies making the information design robust to communication. In this case, the designer is restricted to public signal information structures, as every signal realization will be ultimately observed by both agents. Hence, she needs to include the constraint $q = r$ in her linear optimization program. This constraint imposes the restriction of public signals on the choice of information structure. From Table 3 it becomes clear that in some instances this constraint is binding and leads to lower optimal value of the designer's objective function. Looking at the case of strategic complements ($c > d$), whenever the optimal information structure consists of private signals, the communication constraint is binding. In these cases, the designer needs to choose a constrained optimum information structure consisting of a public signal, which will be communication robust.

6.2 Multiplicity of Equilibria and Other Solution Concepts

Information design is about finding the information structure under which the most beneficial BNE is played. However, this does not exclude the possibility of there being multiple BNE under the optimal information structure. Information design is subject to the same criticisms as mechanism design with regards to the multiple equilibrium problem. For example, consider the symmetric $2 \times 2 \times 2$ environment when there are strategic complements ($c > d$). Under the null information structure, there are three (agent) symmetric BNE: always play a_0 , always play a_1 , and mix with equal probabilities between the two actions. Only the last equilibrium is also symmetric in the state, which is what we focus on in our two-dimensional representation. Further restrictions need to be imposed on the information structure through the incentive constraints in order to ensure the uniqueness of equilibrium.

In our framework, we exclusively focus on BNE as a solution concept. Working with different solution concepts may change the environment, by modifying the incentive constraints or adding new ones. For example, Lehrer, Rosenberg and Shmaya (13) consider solutions concepts which allow for degrees of communication and correlation between the agents that are different from Nash equilibrium. The optimal information structure, which supports the most desirable equilibrium under a different solution concept, can be derived using the approach described here with appropriately modified incentive constraints. An interesting question to investigate would be how the optimal information structures compare, in terms of their complexity, across the different solution concepts.

6.3 Exogenous Information

Our analysis and results are based on the assumption that the designer is in complete control of the informational environment. In particular, we assumed away any signals observed by the agents prior to the ones sent by the designer. In some instances, however, this assumption is unrealistic as the agents may already have some information about the state. Depending on the nature of this information, the designer's ability to achieve the highest possible objective may be impeded. In either case, the designer needs to take into account the prior signals of the agents and incorporate that as an additional constraint into her information design problem.

Consider the motivating example of Section 3. Let us assume that in the course of the trial, the jury already observed some evidence. In particular, she knows that the jury members have observed signal realizations from the following information structure:

θ_0	t_0	t_1	θ_1	t_0	t_1
t_0	6/20	7/20	t_0	0	7/20
t_1	7/20	0	t_1	7/20	6/20

The precision of these signals is $\frac{13}{20}$ which is slightly higher than the precision of $\frac{3}{5}$ that the optimal information structure was characterized by. Given this prior information structure

observed by the jurors, the prosecutor can no longer achieve the unconstrained maximum objective value of $\frac{4}{5}$. In fact, the best she could do is choose a completely uninformative investigation and obtain a value of $\frac{7}{10}$.

Therefore, in the presence of prior information, the optimal information structure may significantly change as compared to the case of no prior information. Similar considerations apply when the agents observe additional signals beyond what the designer reveals. As long as the designer knows the structure of the exogenous information, she can incorporate that as additional constraints into her optimization problem. These constraints may not affect her ability to achieve the same maximum value of the objective as when the agents have no exogenous information. Nevertheless, she needs to modify the optimal information structure, since it is augmented by the exogenous signals.

7 Conclusion

The incentives of rational agents to behave in a certain way are determined by their payoffs and by their beliefs. Mechanism design concerns the modification of payoffs so that people have incentives to behave in desired ways. This paper lays out the methodology of information design. Information design operates on the beliefs of the agents through the choice of information structure. It thus focuses on the choice and creation of information structures under which agents achieve the most favorable outcomes.

As suggested in the previous section, there are many important extensions and robustness issues that can be studied with the proposed method. All of these constitute interesting directions for future research.

Appendix A

Proof of Proposition 1

First we prove that $BCE(G, \underline{S}) \subseteq \cup_S BNE(G, S)$. Choose $\nu \in BCE(G, \underline{S})$. Hence, it must hold that

$$\sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a'_i, a_{-i}), \theta) \quad (7.1)$$

for each $i \in I$, $a_i \in A_i$ and $a'_i \in A_i$. Consider the information structure $S^* = (T^*, \pi^*)$ with $T_i^* = A_i^* \times \underline{t}_i$ and

$$\pi^*((a_i, \underline{t}_i)_{i=1}^N | \theta) = \pi^*(a, \underline{t} | \theta) = \nu(a | \theta) \quad (7.2)$$

for each $a \in A$ and $\theta \in \Theta$. In the game (G, S^*) consider the “truthful” behavioral strategy β_i^* for agent i with

$$\beta_i^*(a_i | a'_i, \underline{t}_i) = \begin{cases} 1, & \text{if } a_i = a'_i \\ 0, & \text{if } a_i \neq a'_i \end{cases} \quad (7.3)$$

for all $a_i, a'_i \in A_i$. The interim payoff to agent i observing signal (a_i, \underline{t}_i) and choosing action a'_i when his opponents follow β_{-i}^* is

$$\begin{aligned} \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi^*((a_i, a'_{-i}), \underline{t} | \theta) \left(\prod_{j \neq i} \beta_j^*(a_j | a'_j, \underline{t}_j) \right) u_i((a'_i, a_{-i}), \theta) \\ = \sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (7.4)$$

where we use (7.2) and (7.3). Therefore, the BNE interim incentive compatibility condition

$$\begin{aligned} \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi^*((a_i, a'_{-i}), \underline{t} | \theta) \left(\prod_{j \neq i} \beta_j^*(a_j | a'_j, \underline{t}_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi^*((a_i, a'_{-i}), \underline{t} | \theta) \left(\prod_{j \neq i} \beta_j^*(a_j | a'_j, \underline{t}_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (7.5)$$

is equivalent to and implied by the BCE obedience constraint (7.1). Hence, β^* is a BNE of (G, S^*) . The distribution over actions conditional on states generated from this equilibrium strategy is

$$\sum_{a' \in A} \pi^*(a', \underline{t} | \theta) \left(\prod_{j=1}^N \beta_j(a_j | a'_j, \underline{t}_j) \right) = \nu(a | \theta). \quad (7.6)$$

Thus, ν is a BNE of the game (G, S^*) , i.e. $\nu \in BNE(G, S^*)$. This implies $BCE(G, \underline{S}) \subseteq \cup_S BNE(G, S)$.

Next we prove that $BCE(G, \underline{S}) \supseteq \cup_S BNE(G, S)$. Choose $\tilde{\nu} \in \cup_S BNE(G, S)$. Hence,

there exist an information structure $\tilde{S} = (\tilde{T}, \tilde{\pi})$ and a BNE behavioral strategy $\beta(a|\tilde{t})$ of (G, \tilde{S}) such that

$$\tilde{\nu}(a|\theta) = \sum_{\tilde{t} \in \tilde{T}} \tilde{\pi}(\tilde{t}|\theta) \left(\prod_{j=1}^N \beta_j(a_j|\tilde{t}_j) \right). \quad (7.7)$$

We write $\tilde{\pi}(\tilde{t}, \underline{t}|\theta) = \tilde{\pi}(\tilde{t}|\theta)$, $\beta(a|\tilde{t}, \underline{t}) = \beta(a|\tilde{t})$ and $\tilde{\nu}(a|\underline{t}, \theta) = \tilde{\nu}(a|\theta)$ for $\tilde{t} \in \tilde{T}$ and $\{\underline{t}\} = \underline{T}$, which trivially holds.

For each a_i such that $\beta_i(a_i|\tilde{t}_i, \underline{t}_i) > 0$, by the BNE incentive compatibility condition it must hold that

$$\begin{aligned} \sum_{a_{-i}, \tilde{t}_{-i}, \theta} \psi(\theta) \tilde{\pi}((\tilde{t}_i, \tilde{t}_{-i}), \underline{t}|\theta) \left(\prod_{j \neq i} \beta_j(a_j|\tilde{t}_j, \underline{t}_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, \tilde{t}_{-i}, \theta} \psi(\theta) \tilde{\pi}((\tilde{t}'_i, \tilde{t}_{-i}), \underline{t}|\theta) \left(\prod_{j \neq i} \beta_j(a_j|\tilde{t}_j, \underline{t}_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (7.8)$$

for each $i \in I$, $\tilde{t}_i \in \tilde{T}_i$, and $a'_i \in A_i$. Multiplying both sides by $\sum_{\tilde{t}_i} \beta_i(a_i|\tilde{t}_i, \underline{t}_i)$ gives

$$\begin{aligned} \sum_{a_{-i}, \tilde{t}, \theta} \psi(\theta) \tilde{\pi}((\tilde{t}_i, \tilde{t}_{-i}), \underline{t}|\theta) \left(\prod_{j=1}^N \beta_j(a_j|\tilde{t}_j, \underline{t}_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, \tilde{t}, \theta} \psi(\theta) \tilde{\pi}((\tilde{t}'_i, \tilde{t}_{-i}), \underline{t}|\theta) \left(\prod_{j=1}^N \beta_j(a_j|\tilde{t}_j, \underline{t}_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (7.9)$$

which by (7.7) is equivalent to

$$\sum_{a_{-i}, \theta} \psi(\theta) \tilde{\nu}(a|\underline{t}, \theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \tilde{\nu}(a|\underline{t}, \theta) u_i((a'_i, a_{-i}), \theta). \quad (7.10)$$

Thus, $\tilde{\nu} \in BCE(G, \underline{S})$, which implies $BCE(G, \underline{S}) \supseteq \cup_S BNE(G, S)$. \square

Proof of Proposition 2

By definition, (ii) implies (i) and (iii). Let us first show that (i) implies (ii). Take a basic game G and an information structure $S = (T, \pi)$ with value V^* . Suppose that β is the BNE of (G, S) which generates that value, that is

$$\sum_{a, t, \theta} V(a, \theta) \pi(t|\theta) \left(\prod_{i=1}^N \beta_i(a_i|t_i) \right) \psi(\theta) = V^*. \quad (7.11)$$

Let $T^a = \{t | \beta_i(a_i | t_i) > 0 \forall i \in I\}$. Consider the direct information structure $S' = (A, \pi')$ with

$$\pi'(a|\theta) = \sum_{t \in T^a} \pi(t|\theta) \left(\prod_{i=1}^N \beta_i(a_i | t_i) \right). \quad (7.12)$$

We will show that the truthful strategy of playing the action implied by the signal realization is a BNE of (G, S') .

Since $\beta_i(a_i | t_i) > 0$ we have the BNE incentive compatibility condition

$$\begin{aligned} \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i}) | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i}) | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (7.13)$$

for each $i \in I$, $t_i \in T_i$ and $a'_i \in A_i$. Multiplying both sides of the above inequality by $\beta_i(a_i | t_i)$ and summing across all t_i we get

$$\begin{aligned} \sum_{t_i} \beta_i(a_i | t_i) \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i}) | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{t_i} \beta_i(a_i | t_i) \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i}) | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (7.14)$$

which can be rewritten as

$$\begin{aligned} \sum_{a_{-i}, \theta} \psi(\theta) \sum_t \pi((t_i, t_{-i}) | \theta) \left(\prod_{i=1}^N \beta_i(a_i | t_i) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, \theta} \psi(\theta) \sum_t \pi((t_i, t_{-i}) | \theta) \left(\prod_{i=1}^N \beta_i(a_i | t_i) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (7.15)$$

for each $i \in I$ and $a'_i \in A_i$. Substituting in with (7.12) we obtain

$$\sum_{a_{-i}, \theta} \psi(\theta) \pi'((a_i, a_{-i}) | \theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \pi'((a'_i, a_{-i}) | \theta) u_i((a'_i, a_{-i}), \theta). \quad (7.16)$$

In the game (G, S') consider the behavioral strategy β'_i for agent i with

$$\beta'_i(a_i | a'_i) = \begin{cases} 1, & \text{if } a_i = a'_i \\ 0, & \text{if } a_i \neq a'_i \end{cases} \quad (7.17)$$

for all $a_i, a'_i \in A_i$. The interim payoff to agent i observing signal a_i and choosing action a'_i in

(G, S') when each opponent j follows strategy β'_j is

$$\begin{aligned} \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi'((a_i, a'_{-i})|\theta) \left(\prod_{j \neq i} \beta'_j(a_j|a'_j) \right) u_i((a'_i, a_{-i}), \theta) \\ = \sum_{a_{-i}, \theta} \psi(\theta) \pi'((a_i, a_{-i})|\theta) u_i((a'_i, a_{-i}), \theta). \end{aligned} \quad (7.18)$$

Hence, (7.16) implies the BNE incentive compatibility conditions for strategy profile β' . Under β' the expected payoff to the designer is given by

$$\mathbb{E}[V] = \sum_{a, \theta} V(a, \theta) \pi'(a|\theta) \psi(\theta) = V^*. \quad (7.19)$$

where we use (7.11) and (7.12). Hence, the direct information structure S' also has value V^* .

Next we show that (iii) implies (ii). For basic game G , consider a BNE distribution ν such that

$$\mathbb{E}_\nu[V] = \sum_{a, \theta} V(a, \theta) \nu(a|\theta) \psi(\theta) = V^*.$$

By Proposition 1 we know that $\nu \in BCE(G, \underline{S})$. Hence, it holds

$$\sum_{a_{-i}, \theta} \psi(\theta) \nu((a_i, a_{-i})|\theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \nu((a_i, a_{-i})|\theta) u_i((a'_i, a_{-i}), \theta). \quad (7.20)$$

Consider the direct information structure $S = (A, \pi)$ with $\pi(a|\theta) = \nu(a|\theta)$ for all $a \in A$ and $\theta \in \Theta$. In the game (G, S) consider the behavioral strategy β_i for agent i with

$$\beta_i(a_i|a'_i) = \begin{cases} 1, & \text{if } a_i = a'_i \\ 0, & \text{if } a_i \neq a'_i \end{cases} \quad (7.21)$$

for all $a_i, a'_i \in A_i$. The interim payoff to agent i observing signal a_i and choosing action a'_i in (G, S) when each opponent j follows strategy β_j is

$$\begin{aligned} \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi((a_i, a'_{-i})|\theta) \left(\prod_{j \neq i} \beta_j(a_j|a'_j) \right) u_i((a'_i, a_{-i}), \theta) \\ = \sum_{a_{-i}, \theta} \psi(\theta) \pi((a_i, a_{-i})|\theta) u_i((a'_i, a_{-i}), \theta) = \sum_{a_{-i}, \theta} \psi(\theta) \nu((a_i, a_{-i})|\theta) u_i((a'_i, a_{-i}), \theta). \end{aligned} \quad (7.22)$$

where the first equality follow by (7.21) and the second equality follows from $\pi(a|\theta) = \nu(a|\theta)$. Hence, (7.22) implies the BNE incentive compatibility conditions for strategy profile β . The

distribution of actions conditional on states of the world under β is

$$\sum_{a' \in A} \pi(a'|\theta) \left(\prod_{i=1}^N \beta_i(a_i|a'_i) \right) = \pi(a|\theta) = \nu(a|\theta) \quad (7.23)$$

and thus, the expected payoff of the designer under strategy profile β is given by

$$\mathbb{E}[V] = \sum_{a, \theta} V(a, \theta) \nu(a|\theta) \psi(\theta) = V^*. \quad (7.24)$$

Hence, the direct information structure S has value V^* . \square

Proof of Proposition 3

Let ν^* be an optimal direct signal structure that generates an action profile \underline{a} , i.e. $\nu^*(\underline{a}|\theta) > 0$ for some $\theta \in \Theta$. The expected payoff for the designer under realization \underline{a} is given by $\mathbb{E}[V(\underline{a})] = \sum_{\theta} V(\underline{a}, \theta) \nu^*(\theta|\underline{a})$. Suppose that upon observing the signal realization \underline{a} , the designer decides to release a new signal structure $\hat{\nu}$. She will only strictly prefer to do that if this gives her a higher expected payoff, that is if:

$$\sum_{\theta} V(\underline{a}, \theta) \nu^*(\theta|\underline{a}) < \sum_{\theta} \nu^*(\theta|\underline{a}) \sum_a V(a, \theta) \hat{\nu}(a|\theta). \quad (7.25)$$

Multiplying both sides of the above inequality by $\nu^*(\underline{a}) = \sum_{\theta \in \Theta} \nu^*(\underline{a}|\theta)$ we obtain:

$$\sum_{\theta} V(\underline{a}, \theta) \nu^*(\underline{a}, \theta) < \sum_{\theta} \nu^*(\underline{a}, \theta) \sum_a V(a, \theta) \hat{\nu}(a|\theta). \quad (7.26)$$

Suppose that the designer released the following information structure to start with:

$$\nu(a|\theta) = \begin{cases} \nu^*(\underline{a}|\theta) \hat{\nu}(a|\theta) & \text{if } \hat{\nu}(a|\theta) > 0 \text{ and } \nu^*(a|\theta) = 0 \\ \nu^*(a|\theta) + \nu^*(\underline{a}|\theta) \hat{\nu}(a|\theta) & \text{if } \hat{\nu}(a|\theta) > 0 \text{ and } \nu^*(a|\theta) > 0 \\ \nu^*(a|\theta) & \text{if } \hat{\nu}(a|\theta) = 0, \nu^*(a|\theta) > 0 \text{ and } a \neq \underline{a} \\ 0 & \text{otherwise} \end{cases} \quad (7.27)$$

for each $a \in A$ and $\theta \in \Theta$. It is straightforward to see that $\nu(a|\theta) \in [0, 1]$ and $\sum_a \nu(a|\theta) = 1$ and hence $\nu(a|\theta)$ is a proper probability distribution. Obedience of $\nu(a|\theta)$ requires:

$$\sum_{a_{-i}, \theta} \psi(\theta) \nu((a_i, a_{-i})|\theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \nu((a_i, a_{-i})|\theta) u_i((a'_i, a_{-i}), \theta) \quad (7.28)$$

which in turn is equivalent to

$$\begin{aligned} \sum_{a_{-i}, \theta} \psi(\theta) [\nu^*((a_i, a_{-i})|\theta) + \nu^*(\underline{a}|\theta)\hat{\nu}(a_i, a_{-i}|\theta)] u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, \theta} \psi(\theta) [\nu^*((a_i, a_{-i})|\theta) + \nu^*(\underline{a}|\theta)\hat{\nu}(a_i, a_{-i}|\theta)] u_i((a'_i, a_{-i}), \theta). \end{aligned} \quad (7.29)$$

This condition is satisfied due to the obedience of ν^* under the prior $\psi(\theta)$ given by

$$\sum_{a_{-i}, \theta} \psi(\theta) \nu^*((a_i, a_{-i})|\theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \nu^*((a_i, a_{-i})|\theta) u_i((a'_i, a_{-i}), \theta) \quad (7.30)$$

and the obedience of $\hat{\nu}$ under the prior $\nu^*(\theta|\underline{a})$ given by

$$\sum_{a_{-i}, \theta} \frac{\psi(\theta) \nu^*(\underline{a}|\theta)}{\nu^*(\underline{a})} \hat{\nu}((a_i, a_{-i})|\theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \frac{\psi(\theta) \nu^*(\underline{a}|\theta)}{\nu^*(\underline{a})} \hat{\nu}((a_i, a_{-i})|\theta) u_i((a'_i, a_{-i}), \theta). \quad (7.31)$$

The signal structure $\nu(a|\theta)$ would thus result in an expected payoff of

$$\begin{aligned} \sum_{\theta} \nu^*(\underline{a}|\theta) \psi(\theta) \sum_a V(a, \theta) \hat{\nu}(a|\theta) + \sum_{a \setminus \underline{a}, \theta} V(a, \theta) \nu^*(a|\theta) \psi(\theta) \\ = \sum_{\theta} \nu^*(\underline{a}, \theta) \sum_a V(a, \theta) \hat{\nu}(a|\theta) + \sum_{a \setminus \underline{a}, \theta} V(a, \theta) \nu^*(a, \theta). \end{aligned} \quad (7.32)$$

By (7.26) the above expression is strictly larger than the expected payoff under ν^* , which can be written as

$$\sum_{\theta} V(\underline{a}, \theta) \nu^*(\underline{a}, \theta) + \sum_{a \setminus \underline{a}, \theta} V(a, \theta) \nu^*(a, \theta). \quad (7.33)$$

This is a contradiction to ν^* being optimal. □

Proof of Proposition 4

For the null information structure \underline{S} and a basic game $G_{c,d}$ the general BCE constraints given in Definition 2 become:

for $a_i = a_0, a'_i = a_1$:

$$\frac{1}{2}rc + \frac{1}{2}(q-r)d \geq \frac{1}{2}(q-r)c + \frac{1}{2}(1-2q+r)d$$

and

for $a_i = a_1, a'_i = a_0$:

$$\frac{1}{2}rc + \frac{1}{2}(q-r)d \geq \frac{1}{2}(q-r)c + \frac{1}{2}(1-2q+r)d.$$

These two constraints are equivalent and reduce to only one inequality:

$$2(c-d)r \geq d + (c-3d)q. \quad (7.34)$$

Additionally, the parameters need to satisfy:

$$r \leq q \quad (7.35)$$

$$r \geq \max\{2q-1, 0\} \quad (7.36)$$

and

$$q \in [0, 1]. \quad (7.37)$$

Therefore, the set of BCE distributions of $(G_{c,d}, \underline{S})$ is equivalent to the set of (q, r) -pairs which satisfy constraints (7.34)–(7.37).

Case 1: Assume $c > d \geq 0$ (strategic complements). The obedience constraint (7.34) can thus be written as:

$$r \geq \frac{d}{2(c-d)} + \frac{c-3d}{2(c-d)}q \quad (7.38)$$

In this case, constraint (7.37), which essentially coincides with the x -axis of the graph, is never binding. The reason behind this is the following. The intercept of constraint (7.38) is always positive. When in addition $c \geq 3d$, the slope is also positive. Hence, this constraint is always more binding than (7.37), as it always lies above the x -axis. On the other hand, when $c < 3d$, the slope of (7.38) is negative. However, it is easy to show that (7.38) intersects (7.36) before it intersects the x -axis. Therefore, for the relevant range of values, (7.38) lies above the x -axis also in this case. Hence, (7.37) is never binding.

The set of distributions which satisfy (7.38), (7.35) and (7.36) is thus equivalent to the convex hull formed by the intersection points $(q_1, r_1) = \left(\frac{d}{c+d}, \frac{d}{c+d}\right)$ (of (7.38) and (7.35)), $(q_2, r_2) = \left(\frac{2c-d}{3c-d}, \frac{c-d}{3c-d}\right)$ (of (7.38) and (7.36)) and $(q_3, r_3) = (1, 1)$ (of (7.35) and (7.36)).

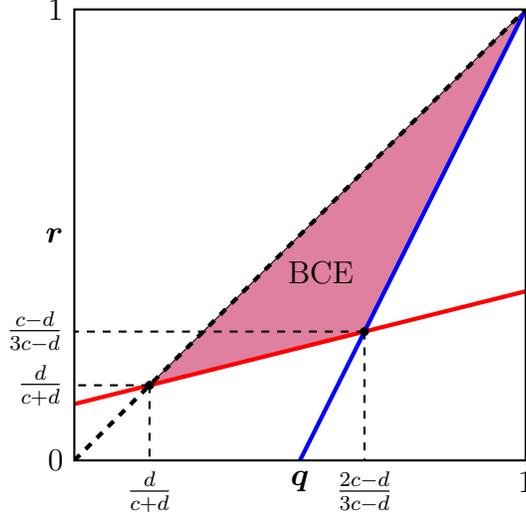


Figure 6: Strategic Complements ($c > d \geq 0$)

Case 2: Assume $d > c \geq 0$ (strategic substitutes). The obedience constraint (7.34) can thus be written as:

$$r \leq \frac{d}{2(c-d)} + \frac{c-3d}{2(c-d)}q \quad (7.39)$$

This constraint has a negative intercept and a positive slope. In fact, it always holds that the slope $\frac{c-3d}{2(c-d)} \geq \frac{3}{2}$. When $c > 0$ the slope is strictly greater than $\frac{3}{2}$ and (7.39) intersects only constraints (7.37) and (7.35). In this case, all four constraints (7.39), (7.35), (7.36) and (7.37) are binding. The set of BCE distributions which satisfy all of them is equivalent to the hull formed by the intersection points $(q_1, r_1) = \left(\frac{d}{c+d}, \frac{d}{c+d}\right)$ (of (7.39) and (7.35)), $(q_3, r_3) = (1, 1)$ (of (7.35) and (7.36)), $(q_5, r_5) = \left(\frac{1}{2}, 0\right)$ (of (7.36) and (7.37)) and $(q_4, r_4) = \left(\frac{d}{3d-c}, 0\right)$ (of (7.37) and (7.39)).

When $c = 0$, the slope of (7.39) is exactly equal to $\frac{3}{2}$. In this case (7.39), (7.35) and (7.36) all intersect at one point — $(q_3, r_3) = (1, 1)$ — and (7.35) is never binding. The set of BCE distributions is equivalent to the hull formed by the intersection points $(q_3, r_3) = (1, 1)$, $(q_5, r_5) = \left(\frac{1}{2}, 0\right)$ and $(q_4, r_4) = \left(\frac{1}{3}, 0\right)$.

Case 3: In the special case of $c = d > 0$, the obedience constraint (7.34) becomes $q \geq \frac{1}{2}$. The set of BCE is then equivalent to the convex hull of $(q_1, r_1) = \left(\frac{1}{2}, \frac{1}{2}\right)$, $(q_2, r_2) = \left(\frac{1}{2}, 0\right)$, and $(q_3, r_3) = (1, 1)$. \square

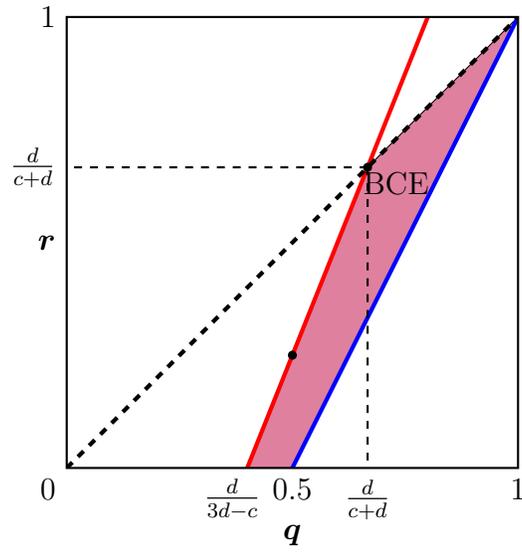


Figure 7: Strategic Substitutes ($d > c \geq 0$)

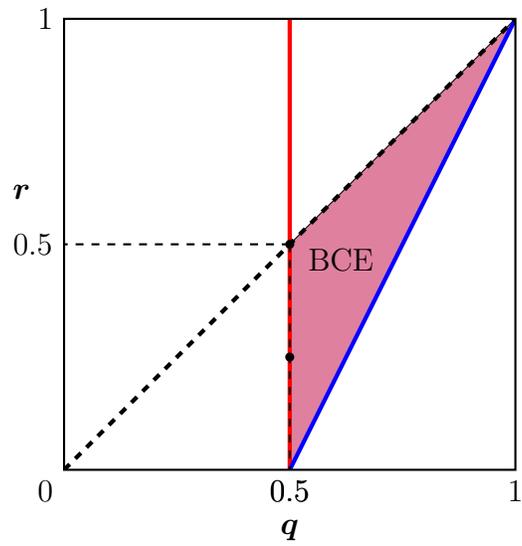


Figure 8: $c = d > 0$

Proof of Proposition 5

Since $\nu^*(q, r) \in BCE(G_{c,d}, \underline{S})$, we know that (7.34) holds. We will show that this condition implies the behavioral strategy

$$\beta_i(a_k|t_n) = \begin{cases} 1 & \text{if } k \neq n \\ 0 & \text{if } k = n \end{cases} \quad (7.40)$$

for $k, n = 0, 1$ and $i = 1, 2$ is a BNE in the incomplete information game $(G_{c,d}, S_{1-q,1-2q+r})$. For (7.40) to be an equilibrium, it needs to hold that the BNE incentive compatibility conditions given by (4.1) in Definition 1 are satisfied. Due to the symmetry of the players, we need to consider only player i . The incentive constraint for $\beta_i(a_1|t_0) = 1$ is given by

$$\begin{aligned} & \frac{1}{2}\pi(t_0, t_0|\theta_0)\beta_j(a_1|t_0)u(a_1, a_1, \theta_0) + \frac{1}{2}\pi(t_0, t_0|\theta_1)\beta_j(a_1|t_0)u(a_1, a_1, \theta_1) \\ & + \frac{1}{2}\pi(t_0, t_1|\theta_0)\beta_j(a_0|t_1)u(a_1, a_0, \theta_0) + \frac{1}{2}\pi(t_0, t_1|\theta_1)\beta_j(a_0|t_1)u(a_1, a_0, \theta_1) \\ & \geq \frac{1}{2}\pi(t_0, t_0|\theta_0)\beta_j(a_1|t_0)u(a_0, a_1, \theta_0) + \frac{1}{2}\pi(t_0, t_0|\theta_1)\beta_j(a_1|t_0)u(a_0, a_1, \theta_1) \\ & + \frac{1}{2}\pi(t_0, t_1|\theta_0)\beta_j(a_0|t_1)u(a_0, a_0, \theta_0) + \frac{1}{2}\pi(t_0, t_1|\theta_1)\beta_j(a_0|t_1)u(a_0, a_0, \theta_1). \end{aligned} \quad (7.41)$$

When we substitute in the probabilities $\pi(\cdot|\theta)$ of the information structure $S_{1-q,1-2q+r}$, the equilibrium strategy probabilities β_j and the basic game payoffs, the above condition reduces to

$$rc + (q - r)d \geq (1 - 2q + r)d + (q - r)c. \quad (7.42)$$

The incentive constraint for $\beta_i(a_0|t_1) = 1$ is given by

$$\begin{aligned} & \frac{1}{2}\pi(t_1, t_0|\theta_0)\beta_j(a_1|t_0)u(a_0, a_1, \theta_0) + \frac{1}{2}\pi(t_1, t_0|\theta_1)\beta_j(a_1|t_0)u(a_0, a_1, \theta_1) \\ & + \frac{1}{2}\pi(t_1, t_1|\theta_0)\beta_j(a_0|t_1)u(a_0, a_0, \theta_0) + \frac{1}{2}\pi(t_1, t_1|\theta_1)\beta_j(a_0|t_1)u(a_0, a_0, \theta_1) \\ & \geq \frac{1}{2}\pi(t_1, t_0|\theta_0)\beta_j(a_1|t_0)u(a_1, a_1, \theta_0) + \frac{1}{2}\pi(t_1, t_0|\theta_1)\beta_j(a_1|t_0)u(a_1, a_1, \theta_1) \\ & + \frac{1}{2}\pi(t_1, t_1|\theta_0)\beta_j(a_0|t_1)u(a_1, a_0, \theta_0) + \frac{1}{2}\pi(t_1, t_1|\theta_1)\beta_j(a_0|t_1)u(a_1, a_0, \theta_1). \end{aligned} \quad (7.43)$$

After substituting in we obtain:

$$(q - r)d + rc \geq (q - r)c + (1 - 2q + r)d. \quad (7.44)$$

Notice that (7.42) and (7.44) are equivalent and, moreover, implied by the BCE condition (7.34). Hence, the behavioral strategy (7.40) is an equilibrium in the incomplete information game $(G_{c,d}, S_{1-q,1-2q+r})$.

We now need to show that this BNE strategy generates the distribution $\nu^*(q, r)$. We will show that for the distribution conditional on θ_0 , as the rest follows by analogy. The decision rule induced by strategy (7.40) is given by $\sigma(a_0, a_0|t_1, t_1, \theta_0) = \sigma(a_0, a_1|t_1, t_0, \theta_0) = \sigma(a_1, a_0|t_0, t_1, \theta_0) = \sigma(a_1, a_1|t_0, t_0, \theta_0) = 1$ and zero otherwise. Hence, we obtain:

$$\begin{aligned}\nu(a_0, a_0|\theta_0) &= \sigma(a_0, a_0|t_1, t_1, \theta_0)\pi(t_1, t_1|\theta_0) = r \\ \nu(a_0, a_1|\theta_0) &= \sigma(a_0, a_1|t_1, t_0, \theta_0)\pi(t_1, t_0|\theta_0) = q - r \\ \nu(a_1, a_0|\theta_0) &= \sigma(a_1, a_0|t_0, t_1, \theta_0)\pi(t_0, t_1|\theta_0) = q - r \\ \nu(a_1, a_1|\theta_0) &= \sigma(a_1, a_1|t_0, t_0, \theta_0)\pi(t_0, t_0|\theta_0) = 1 - 2q + r.\end{aligned}$$

By analogy, we obtain the corresponding values for $\nu(\cdot|\theta_1)$. This is exactly equivalent to the distribution $\nu^*(q, r)$. Thus, (7.40) is a BNE strategy which induces $\nu^*(q, r)$ in $(G_{c,d}, S_{1-q,1-2q+r})$. \square

Proof of Proposition 6

Case 1: Consider a designer with $R < 0$, $Q < 0$, and $-\frac{Q}{R} < \frac{c-3d}{2(c-d)}$ in the case of strategic complements. By Theorem 3 we know that the optimal information structure is a public signal with precision $q = \frac{d}{c+d}$. Notice that $\frac{c-3d}{2(c-d)}$ is increasing in c and decreasing in d . Therefore an increase in c and a decrease in d will both increase its value. This implies that after the change, the optimal information structure will be a public signal with a new precision level. Therefore, when we consider such changes in the parameters, we can write the designer's utility under the optimal information structure as:

$$V^* = (R + Q)\frac{d}{c + d} + C$$

where C is a constant. Notice that V^* is increasing in c :

$$\frac{\partial V^*}{\partial c} = -(R + Q)\frac{d}{(c + d)^2} > 0$$

and decreasing in d :

$$\frac{\partial V^*}{\partial d} = (R + Q)\frac{c}{(c + d)^2} < 0.$$

Case 2: Consider a designer with $R > 0$, $Q < 0$, and $-\frac{Q}{R} > \frac{c-3d}{2(c-d)}$ in the case of strategic substitutes. Since $\frac{c-3d}{2(c-d)}$ is increasing in c and decreasing in d , a decrease in c and a increase

in d will both decrease its value. This implies that after the change, the optimal information structure will be a private signal with a new precision level and same correlation $r = 0$. Therefore, when we consider such changes in the parameters, we can write the designer's utility under the optimal information structure as:

$$V^* = Rr + Q\frac{d}{3d-c} + C = Q\frac{d}{3d-c} + C$$

where C is a constant. Notice that V^* is decreasing in c :

$$\frac{\partial V^*}{\partial c} = Q\frac{d}{(3d-c)^2} < 0$$

and increasing in d :

$$\frac{\partial V^*}{\partial d} = -Q\frac{c}{(3d-c)^2} > 0$$

a decrease in c or an increase in d always increases the maximal utility of the designer. The same argument applies to the case of $R < 0$, $Q < 0$ and strategic substitutes. \square

Appendix B

Table 3: Characterization of Optimal Information Structure

	complements [$c > d$]	substitutes [$c < d$]
$R > 0, Q > 0$	full information	full information
$R < 0, Q > 0$	public signal ($q = \frac{d}{c+d}$) if $-\frac{Q}{R} < \frac{c-3d}{2(c-d)}$	—————
	private signals ($q = \frac{2c-d}{3c-d}, r = \frac{c-d}{3c-d}$) if $\frac{c-3d}{2(c-d)} < -\frac{Q}{R} < 2$	null information if $-\frac{Q}{R} < 2$
	full information if $-\frac{Q}{R} > 2$	full information if $-\frac{Q}{R} > 2$
$R > 0, Q < 0$	—————	private signals ($q = \frac{d}{3d-c}, r = 0$) if $-\frac{Q}{R} > \frac{c-3d}{2(c-d)}$
	private signals ($q = \frac{2c-d}{3c-d}, r = \frac{c-d}{3c-d}$) if $-\frac{Q}{R} > 1$	public signal ($q = \frac{d}{c+d}$) if $1 < -\frac{Q}{R} < \frac{c-3d}{2(c-d)}$
	full information if $-\frac{Q}{R} < 1$	full information if $-\frac{Q}{R} < 1$
$R < 0, Q < 0$	public signal ($q = \frac{d}{c+d}$) if $-\frac{Q}{R} < \frac{c-3d}{2(c-d)}$	private signals ($q = \frac{d}{3d-c}, r = 0$) always
	private signals ($q = \frac{2c-d}{3c-d}, r = \frac{c-d}{3c-d}$) if $-\frac{Q}{R} > \frac{c-3d}{2(c-d)}$	—————

Full information: ($q = 1, r = 1$); null information: ($q = \frac{1}{2}, r = 0$); public signals: $q = r$.

References

- [1] Alonso, R., O. Câmara(2014), “Persuading Voters,” *mimeo*. 6
- [2] Angeletos, G.-M., A. Pavan(2007), “Efficient Use of Information and Social Value of Information,” *Econometrica*, 75(4), 1101-1142. 6
- [3] Angeletos, G.-M., A. Pavan(2009), “Policy with Dispersed Information,” *Journal of the European Economic Association*, 7(1), 11-60. 28
- [4] Bergemann, D., T. Heumann, S. Morris (2013), “Information, Interdependence, and Interaction: Where Does the Volatility Come From?,” *mimeo*. 7
- [5] Bergemann, D., S. Morris (2013), “Bayes Correlated Equilibrium and the Comparison of Information Structures,” *mimeo*. 3, 5, 6, 11, 12, 13
- [6] Bergemann, D., S. Morris (2013), “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81(4), 1251-1308. 6
- [7] Crawford, V., J. Sobel (1982), “Strategic Information Transmission,” *Econometrica*, 50(6), 1431-1451. 4
- [8] Eliaz, K., F. Forges (2012), “Disclosing Information to Interacting Agents,” *mimeo*. 6
- [9] Farrell, J., R. Gibbons (1989), “Cheap Talk with Two Audiences,” *American Economic Review*, 79(5), 1214-1223. 4
- [10] Goltsman, M., G. Pavlov (2011), “How to Talk to Multiple Audiences,” *Games and Economic Behavior*, 72, 100-122. 4
- [11] Gossner, O. (2000), “Comparison of Information Structures,” *Games and Economic Behavior*, 30, 44-63. 6
- [12] Kamenica, E., M. Gentzkow (2011), “Bayesian Persuasion,” *American Economic Review*, 101, 2590-2615. 4, 5, 7, 9, 11
- [13] Lehrer, E., D. Rosenberg, E. Shmaya (2010), “Signaling and Mediation in Games with Common Interest,” *Games and Economic Behavior*, 68, 670-682. 6, 11, 31
- [14] Lehrer, E., D. Rosenberg, E. Shmaya (2011), “Garbling of Signals and Outcome Equivalence,” *Discussion Paper*, Tel Aviv University, University of Paris and Northwestern University. 6
- [15] Lester, B., N. Persico, L. Visschers (2012), “Information Acquisition and the Exclusion of Evidence in Trials,” *Journal of Law, Economics, and Organization*, 28, 163-182. 6
- [16] Myerson, R. (1979), “Incentive Compatibility and the Bargaining Problem,” *Econometrica*, 47, 61-73. 15
- [17] Peski, M. (2008), “Comparison of Information Structures in Zero-Sum Games,” *Games and Economic Behavior*, 62, 732-735. 6

- [18] Stinchcombe, M. B. (2011), “Correlated Equilibrium Existence for Infinite Games with Type-Dependent Strategies,” *Journal of Economic Theory*, 146(2), 638-655. 14
- [19] Wang, Y. (2012), “Bayesian Persuasion with Multiple Receivers,” *Job Market Paper*, Department of Economics, University of Pittsburgh, 1-45. 5