THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Information design

Citation for published version:

Taneva, I 2019, 'Information design', *American Economic Journal: Microeconomics*, vol. 11, no. 4, pp. 151-185. <https://doi.org/10.1257/mic.20170351>

Digital Object Identifier (DOI):

[10.1257/mic.20170351](https://doi.org/10.1257/mic.20170351)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

American Economic Journal: Microeconomics

Publisher Rights Statement:

This is the accepted version of the following article: Taneva, Ina. 2019. "Information Design." *American Economic Journal: Microeconomics*, 11 (4): 151-85. DOI: 10.1257/mic.20170351, which has been published in final form at <https://www.aeaweb.org/articles?id=10.1257/mic.20170351>.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Information Design

INA TANEVA *

September 3, 2018

Abstract

A designer commits to a signal distribution that is informative about a payoff-relevant state. Conditional upon the privately observed signals, agents take actions that affect their payoffs as well as those of the designer. We show how to derive the (designer) optimal information structure in static finite environments. We fully characterize it in a symmetric binary setting for a parameterized game. In this environment, conditionally independent private signals are never strictly optimal.

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

JEL Classification: C72, D72, D82, D83.

*School of Economics, University of Edinburgh. E-mail Address: ina.taneva@ed.ac.uk
Financial support from ESRC grant ES/N00776X/1 is gratefully acknowledged. This paper is a revised version of the third chapter of my PhD thesis, UT Austin (2014). I am grateful to my advisor, Laurent Mathevet, for his continuous and invaluable help, support and encouragement. I am also indebted to Maxwell Stinchcombe and Thomas Wiseman for their thoughtful comments, advice and guidance. The paper has benefited substantially from the constructive and extensive comments of the Editor, Johannes Hörner. József Sákovics has provided helpful feedback on various versions of this paper. I also wish to thank numerous seminar and conference audiences. All remaining errors are my own.

1 Introduction

In many economic and social settings, one person or institution communicates with multiple agents, who are engaged in a strategic interaction (game). The kind of information supplied can significantly influence the actions taken and, thus, the outcome of the game. For example, in court, a prosecutor presents the results of her investigation to a jury, the members of which must vote and decide on a verdict. 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 lobby groups. In financial markets, firms disclose information about their profitability that is relevant to both shareholders and competitors. In economic policy, central banks release information about their stimulus campaigns, which affects the economic outlook and decisions of consumers, as well as of domestic and foreign investors.

These are but a few settings of economic importance that provide a context for the general questions that this paper strives to address: What is the optimal mode of information transmission between a self-interested designer (sender) and a group of strategically interacting agents (receivers), who form their beliefs and take actions based on the information provided? If agents are rational Bayesian players, how does the designer go about finding the information structure which induces the equilibrium most beneficial to her? These questions constitute the subject of *information design*.

In environments with incomplete information, the behavior of interacting agents is determined by their payoffs and by their beliefs about the payoff relevant states, as well as their beliefs about the information of their opponents. *Mechanism design* takes the informational environment as given and focuses on providing incentives for desired equilibrium behavior by committing to an extensive form of the strategic interaction, i.e. a mechanism. In contrast to this, *information design* studies the way a designer can manipulate the equilibrium behavior of agents by selecting the informational environment under which they operate while holding the mechanism fixed. Information design thus applies to situations where a designer is able to influence the optimal behavior of agents only through the information she provides about the state, without being able to change any aspects of the mechanism.

The main objective and contribution of this paper are simple: to outline the methodological approach to finding the information structure that maximizes the designer's objective in static finite environments. In doing so, we utilize an already existing method, which however has not been explicitly and systematically applied with this specific purpose in mind.¹ We first present the general approach to information design in static settings with finitely many agents, actions, and states, and subsequently apply it to a parameterized symmetric binary environment.

¹To the best of our knowledge, at the time this paper was originally posted (September 2013), the two-step approach presented in Bergemann and Morris (4) had not been articulated.

To fix ideas, consider a group of interacting agents in a static setting, who each choose from a finite set of possible actions. Their payoffs are determined by their action and the actions of their opponents, as well as the realization of a state, over which they hold a common prior belief. This constitutes the *basic game*. In order to analyze the strategic interactions in this setting, one needs 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 to achieve an equilibrium that maximizes the designer's objective. In contrast, information design takes the basic game as given and selects the information structure to create an equilibrium that maximizes the designer's objective.

This paper studies the general problem of a self-interested designer communicating with multiple agents engaged in a strategic interaction. The designer's objective is an arbitrary function of the state and the agents' actions. Without observing the state, the designer commits to an information structure: a mapping from states to joint distributions over signals. Once chosen, the information structure becomes common knowledge. Subsequently, agents observe the drawn signal realizations and form their beliefs about the state and about their opponents' beliefs. Finally, they take actions, which affect their own, their opponent's, and the designer's payoffs.

In this setting, the designer would like to choose an information structure under which, for the given basic game, the agents play a Bayes Nash equilibrium that maximizes the expected value of her objective. To this end, one would 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 fact that there are infinitely many information structures. We apply the concept of *Bayes correlated equilibrium* of Bergemann and Morris (3) for the particular case when the agents have no further information but their common prior. This allows us to characterize the set of all Bayes Nash equilibria associated with all information structures for a given basic game, while circumventing the explicit use of information structures. The designer then maximizes her objective over this set, identifies her preferred Bayes Nash equilibrium and backs out the information structure which supports it.

We illustrate the methodology outlined above in a class of symmetric problems with two agents, two actions and two states, for which it provides crisp results and conclusions. We work with a parameterized basic game, which is broad enough to capture various different strategic interactions. 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 allowing arbitrary designer objectives 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 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 (Farrell and Gibbons (7)). For the symmetric binary setting, we show that in almost all cases, the designer benefits from revealing some information rather than none, while conditionally independent private signals are never strictly optimal, irrespective of the designer’s objective function.

We make several important assumptions. First, the designer chooses the information structure without observing the state, i.e., at the *ex ante* stage, which removes any possibility for signalling through the choice of information structure. Second, the designer has full commitment power. This implies that the meaning of signals is established reliably: once the state is realized, the signals are drawn according to the previously announced probability distributions, and the agents observe undistorted signal realizations. Based on these, they can simply update their beliefs and take actions without questioning the interim incentives of the designer. Third, the designer can costlessly choose any information structure, regardless of its informativeness or correlations between signals.² Fourth, we abstract away from any communication between the agents. Fifth, we restrict attention to the best equilibrium from the designer’s perspective. This assumption is not innocuous, especially in the context of multiple strategic agents where equilibrium selection plays a non-trivial role. We provide some more detailed discussions regarding these assumptions, in general and in the context of our main application, after presenting the results.

The remainder of the paper is organized as follows. Section 2 provides an overview of the related literature, followed by a simple example presented in Section 3. Section 4 introduces the framework and outlines the general approach to information design in static finite settings. In Section 5 we apply the general approach to a symmetric binary environment, for which we provide a complete analysis and characterization of the optimal information structure. We also present some important extensions and discuss the ways in which they can be incorporated into the problem, as well as their impact on the optimal information structure. Section 6 concludes with some directions for future research. Proofs are relegated to the Appendix.

²In their Section I.C, Kamenica and Gentzkow (10) provide an excellent discussion on why these assumptions may, in fact, be not as restrictive as they appear at first glance.

2 Literature Review

This paper is related to the literature on cheap talk communication, which focuses on an informed sender that has observed the state and can send costless non-verifiable messages. Most related to our framework are the papers by Farrell and Gibbons (7) and Goltsman and Pavlov (8), which extend the cheap talk model of Crawford and Sobel (5) to an environment with two receivers. There are several consequential differences between these frameworks and ours. First, in our environment, the sender (designer) chooses the information structure before observing the state. Second, she has full commitment power. These two assumptions render the interim stage nonstrategic: the meaning of signals is established *ex ante* and the observed realizations are undistorted. This removes any of the strategic considerations present in cheap talk, as well as in signalling and informed principal frameworks. Third, the receivers in the above papers are decision makers whose payoffs depend only on the state and their own action. In our framework, in contrast, the receivers (agents) are involved in a strategic interaction, whereby their payoffs also depend on the actions of the other agents. Hence, the choice of information structure affects the equilibrium play by determining not only the first-order beliefs but also the higher order beliefs of the agents.

Closely linked to this paper is the literature on Bayesian persuasion, which is sometimes referred to as *cheap talk with commitment*: an uninformed sender costlessly commits to a signal. A central paper in this strand is Kamenica and Gentzkow (10), which is equivalent to single-agent information design. They characterize the optimal signal for any given set of preferences and initial beliefs using tools from convex analysis. However, the techniques and results of Kamenica and Gentzkow (10) 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.*”³

We use a definition of Bayes correlated equilibrium proposed by Bergemann and Morris (3) to answer this open question and show how things differ in the “multiple-interacting-receivers” environment. The proposed approach can be applied to obtain the optimal design of information in any static finite environment. We also provide insights regarding the characterization of the optimal information structure for different properties of the designer’s objective and of the underlying strategic interaction. While Bergemann and Morris (3) 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.

³See Kamenica and Gentzkow (10), p.2609.

cally. Our focus, in contrast, is on the designer, who maximizes her expected payoff by fully controlling the informational environment under which the agents interact.

Wang (16) also examines the question of “Bayesian persuasion with multiple receivers”. She restricts attention to a voting environment, in which the sender has a preference for the same alternative irrespective of the state. She further focuses on conditionally independent private signals or purely public signals and compares the two modes. In contrast, we impose no restriction on the designer’s objective, nor on the information structure. Public and conditionally independent private signals are special cases subsumed in our specification.

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 (13), 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 (6) 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 changing the number of elements in that set. They find that when the sender commits to a disclosure policy, 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 objective function they look at, on the strategic substitutability between actions and the availability of hard evidence. The main part of the analysis in Eliaz and Forges (6) focuses on the case of an informed principal who chooses the information disclosure policy 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 (3), Gossner (9), Lehrer, Rosenberg and Shmaya (11) and (12), Peski (14), etc. Closest to ours is Lehrer, Rosenberg and Shmaya (11) who restrict attention to symmetric games of common interest and rank information structures according to the magnitude of 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 payoff rather than the agents’ equilibrium payoffs.

3 A Simple Example

Consider a policy maker who would like to convince two of her peers to vote for a motion she has put forth. The motion could be thought of as a policy reform proposal, for example a reduction in the state fuel tax. Her objective is to convince her colleagues to vote in favor of the motion, as it will only be implemented if there is unanimity amongst the three parties. The suitability of the motion depends on an unknown policy-relevant state. Specifically, environmental concerns prevail and the tax cut would be considered appropriate only if the state of the local environment has not deteriorated recently. The motion's proponent may convince her colleagues by presenting to them the findings of a report on the status quo of the relevant policy variable (in this case the environment).

She can select the scope of the report, as well as the consultant she hires to produce it. For example, she could choose to include many items to be investigated in the report and hire a very thorough and diligent consultant. This would help her case if the realization of the policy-relevant variable is the kind that deems her proposal suitable, but would otherwise impede the passing of the motion. The motion proponent can commit to talking to her colleagues, simultaneously or individually, about different items in the report. However, once the findings of the report are produced, she must disclose those truthfully. Can this policy maker gain by choosing the scope and quality of the report optimally, in a way that maximizes the overall probability of achieving unanimity and the motion passing?

To formalize the example, suppose the two colleagues that need to be persuaded to vote in favor of the motion are indexed by i and j . There are two states of the world: the environment has deteriorated (θ_0) or not (θ_1). The proponent of the motion (designer) and her colleagues (agents) share a common prior belief, with $\Pr(\theta_0) = 0.7$.

The agents would like to make the right choice: vote against the motion (a_0) when the state is θ_0 and vote in favor of it (a_1) when the state is θ_1 . While unanimity is required to pass the reform, each agent derives utility independently when he makes the right choice, say, because if there is disagreement between the agents, the motion is sent to arbitration by an independent panel of experts. The payoffs of the two agents are summarized by the following matrix:

$\theta = \theta_0$	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

which is common knowledge to all parties involved. The motions's proponent obtains a positive payoff only if the reform is approved, regardless of the state. Therefore, she needs to convince both of her colleagues to vote in favor of the motion. Her payoff function is thus given by

$$V(a_i, a_j, \theta) = \begin{cases} 1 & \text{if } a_i = a_j = a_1 \\ 0 & \text{otherwise.} \end{cases}$$

Formally, the scope and quality of the report can be represented by conditional distributions $\pi(\cdot|\theta_0)$ and $\pi(\cdot|\theta_1)$ over a set of signals profiles $T = T_i \times T_j$. The number of possible signals per agent the proponent can choose is unrestricted; however, without loss, their number can be set equal to the number of possible actions — in this case two — as each signal can be viewed in terms of the action it induces in equilibrium. The proponent then chooses π , which becomes common knowledge, and her colleagues observe the undistorted findings (signal realizations) of the report. We also assume that the proponent can choose any distributions π , which implies that she is able to solicit a report with an arbitrary level of precision, including one that will perfectly reveal the true state (of the environment). This may seem unrealistic, as finding out the true state might be out of reach, irrespective of the thoroughness and diligence of the consultant. However, this is without loss of generality if we interpret the state of the world to be the most informative signal the report is able to produce. Choosing π would then determine how much and in what way the most informative signal gets obscured.

This specification also implies that the report can produce arbitrary correlations between the signals observed by the two agents. Since the proponent can choose to talk to each one of her colleagues individually, she can select which parts of the investigation to focus on when making her case. It could be that one of her colleagues cares more about the overall air quality, while the other one is more concerned with the impact of fracking. Therefore, by emphasizing the respective findings regarding these two items in her private meetings with each colleague, the proponent of the reform can calibrate the correlation between the agents' signals.

First, consider some baseline cases regarding the scope and precision of the report. If the designer chooses a completely uninformative report or equivalently, if she chooses not to solicit one, then both of the agents will vote against the reform. This is their default action profile since θ_0 is more likely than θ_1 . The designer 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 report, the agents will both vote in favor of the reform only when the state is indeed θ_1 . This happens 30 percent of the time and results in an expected payoff of 0.3 for the designer.

However, she can do better. The optimal choice for the report is achieved by the following signal structure, for some signals t_0, t_1 :

$\theta = \theta_0$	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

Under this information structure, it is incentive compatible for each agent to vote against the reform (a_0) when he observes t_0 and in favor of it (a_1) when he observes t_1 given that his opponent does the same. That is, the behavioral strategy profile $\beta^* = (\beta_1^*, \beta_2^*)$, where

$$\beta_i^*(a_k|t_n) = \begin{cases} 1 & \text{if } k = n \\ 0 & \text{otherwise} \end{cases}$$

is a BNE of the game. The BNE incentive compatibility constraints are as follows:

$$\mathbb{E}u_i(a_0|t_0, \beta_{-i}^*) = 1 \cdot 2 = 2 > \mathbb{E}u_i(a_1|t_0, \beta_{-i}^*) = 1 \cdot 0 = 0,$$

and

$$\mathbb{E}u_i(a_1|t_1, \beta_{-i}^*) = \frac{2}{3} \cdot 0 + \frac{1}{3} \cdot 2 = \frac{2}{3} = \mathbb{E}u_i(a_0|t_1, \beta_{-i}^*) = \frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0 = \frac{2}{3},$$

where $\mathbb{E}u_i(a|t, \beta_{-i}^*)$ is agent i 's expected utility from playing a after having observed signal t , assuming that the other agent follows the strategy β_{-i}^* . If the agents play this BNE, the proponent will succeed in passing the reform 90 percent of the time. Her colleagues are well aware that the report was chosen in a way to maximize the probability of passing the reform; yet they act in a rational way, given the signals they observe.

We emphasize two characteristics of the above information structure that are also shared by the optimal information structure in the single receiver case (Kamenica and Gentzkow (10)). First, when an agent votes against the implementation of the reform, the designer's least favorite option, she is certain that the state of the world is θ_0 . 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 designer 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 agent to vote in favor of the reform when observing t_1 . Both of these effects increase the expected payoff of the designer. Hence, the optimal information structure in this setting will always have $\pi(t_1, t_1|\theta_1) = 1$.

Second, when an agent votes in favor, he is indifferent between the two votes. If he were strictly in favor, then the designer could increase the probability of $\pi(t_1, t_1|\theta_0)$, to the point at which the agent becomes indifferent. That will not change the agent's optimal choice given t_1 — he will still choose to vote for the reform — but will increase the probability of (t_1, t_1) and hence, also the probability of a unanimous vote in favor. 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 agent is exactly indifferent between voting in favor and voting against the reform. 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 (10) and the current framework concerns the posterior beliefs. If there were only one agent, his posterior belief on θ_1 would have to be at least $\frac{1}{2}$ in order for him to vote in favor of the reform. Here, in contrast, each agent votes in favor as long as his posterior belief on θ_1 is at least $\frac{1}{3}$. This happens because of the complementarities in the players' actions. Since unanimity is needed for the motion to be passed, and since under the optimal information structure both agents always observe the same signal, each agent knows that voting in favor will only really make a difference if the other agent were to vote in the same way. Therefore, receiving a signal indicative of a state θ_1 , i.e., t_1 , makes an agent more willing to vote in favor for two reasons. First, if the state is indeed θ_1 , then his vote is needed for the motion to pass, which is the

right course of action in this case. Second, if the state is θ_0 , then voting against will not help achieve an outright rejection of the motion (a payoff of 2) and instead will only give him the payoff from unilaterally choosing the right action and sending the motion to a panel (a payoff of 1). This is because each agent knows that conditional on receiving a signal t_1 , the other agent has received the same signal and is, thus, voting in favor of the motion.⁴

4 The General Approach

This section describes the general approach. We first introduce key notation and definitions, after which we set up and solve the information design problem.

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 generally depends on both the agents' actions and the state. Note that, in principle, the designer can be one of the N players, rather than a completely external agent. Designer and agents share a common full support prior $\psi \in \text{int}(\Delta(\Theta))$ that is common knowledge. Let $G = ((A_i, u_i)_{i=1}^N, \psi)$ be 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 conditional signal distributions $\pi : \Theta \rightarrow \Delta(T)$, one for each possible state, with $T = T_1 \times \dots \times T_N$. We denote by t_i 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 A behavioral strategy for agent i in (G, S) is defined as $\beta_i : T_i \rightarrow \Delta(A_i)$.

4.2 Designer's Optimization Problem and Solution

Given a basic game G , the designer commits to and publicly announces an information structure S , which thus becomes common knowledge. Once the state has realized according to ψ ,

⁴If there is no benefit to choosing the right 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, in which both agent always vote in favor, to be achieved. This is because, in this case, under the null information structure, if the other agent were to always vote in favor, it is a best response to do the same irrespective of the beliefs about the state.

⁵This representation of an incomplete information game as a combination of a basic game and an information structure has been previously used in the literature; see for example Gossner (9) and Lehrer, Rosenberg and Shmaya (11).

the signals are drawn according to π and subsequently revealed to each agent privately. Notice that depending on S , these signals may in fact correspond to public information or be common knowledge among different subsets of agents. Upon observing a signal t_i , agent i formulates beliefs based on S about the state and about the beliefs of his opponents. Then, he selects an (mixed) action, $\beta_i(\cdot|t_i) \in \Delta(A_i)$, which maximizes his (interim) expected utility. The resulting distribution over action profiles conditional on states defines a Bayes Nash equilibrium (BNE) of the incomplete information game (G, S) at the interim level.

DEFINITION 1. (Bayes Nash Equilibrium)

For a given (G, S) , consider a strategy profile β such that 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 $a'_i \in A_i$. Then, the distribution $\nu : \Theta \rightarrow \Delta(A)$ given by

$$\nu(a | \theta) := \sum_{t \in T} \pi(t | \theta) \left(\prod_{j=1}^N \beta_j(a_j | t_j) \right), \quad (4.2)$$

is a BNE of (G, S) .

As there could be multiple BNE's of the game (G, S) , we denote the set of these as $BNE(G, S)$. The designer's problem is to choose an information structure which induces agents to play a BNE that maximizes her ex ante expected utility. Hence, for a given basic game G , the designer's problem can be approached as follows: 1) Characterize the set $\cup_S BNE(G, S)$, which comprises all BNEs of G that could emerge under any possible information structure. We refer to this as the *constraint set* of the optimization problem. 2) Maximize the *objective function* of the designer (her ex ante expected utility) over the set $\cup_S BNE(G, S)$. This gives the optimal distribution ν^* . 3) Find the information structure S^* which induces ν^* as a BNE of (G, S^*) .

The following subsections focus on each of the steps outlined above. We show that steps 1 and 2 reduce to a linear programming problem. We also argue that, without loss of generality, we can focus on a particular class of information structures when approaching step 3.

4.2.1 Constraint Set

To determine the constraint set, we need to characterize the set of all BNEs that could emerge under all possible information structures for the given basic game G , i.e. $\cup_S BNE(G, S)$.

Since there exist infinitely many information structures, this might seem like a daunting task. To accomplish it, we draw on a definition of correlated equilibrium introduced by Bergemann and Morris (3) called Bayes correlated equilibrium (BCE). We show below that applying this definition to the basic game G , under the assumption that agents have no further information but their prior beliefs, allows us to characterize the set $\cup_S BNE(G, S)$. With this purpose in mind, we introduce the following definition.

DEFINITION 2. (Bayes Correlated Equilibrium)

A distribution $\nu : \Theta \rightarrow \Delta(A)$ is a BCE of G if for each $i \in I$ and $a_i \in A_i$, we have

$$\sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a'_i, a_{-i}), \theta), \quad (4.3)$$

for all $a'_i \in A_i$.

Subsequently, we use $BCE(G)$ to denote the set of BCEs for the basic game G . We next establish an important equivalence.

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

In the designer's problem, the constraint set $\cup_S BNE(G, S)$ 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. The equivalence result above allows us instead to work with the set $BCE(G)$, which is easier to characterize.

Intuitively, the result can be interpreted as follows. In a BCE distribution, the correlation between the actions given the state is arbitrary, subjective to incentive constraints. In a BNE distribution, the correlation between the actions given the state is generated only through independent randomizations of individual actions conditional on signals. Therefore, in order to generate every possible element in $BCE(G)$ as a BNE under some information structure S , the additional correlation between the actions must come about through conditioning on the signals and their correlation with the common state. Every BCE distribution can thus be replicated as a BNE distribution for an appropriately chosen information structure S , which provides the necessary information about the state and about the information of the other players, so as to generate the required correlation between the equilibrium actions. To summarize, any distribution in the set $BCE(G)$ 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 rely on an information structure at all. In contrast, a distribution in the set $BNE(G, S)$ can be viewed as a combination of an (conditionally) independent stochastic device (the behavioral strategy profile) and the information structure S , where all the correlation between the actions is generated through S . The relationship between an element in $BCE(G)$ and the information structure which implements it as a BNE is established in section 4.2.3.

Using Proposition 1 we can characterize the set of all BNE by means of the BCE incentive constraints (4.3). The next lemma establishes the structure of the constraint set.

LEMMA 1. *The set $BCE(G)$ is a nonempty convex polygon (in $\Delta(A)^\Theta$).*

4.2.2 Designer's Objective

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

$$\max_{\nu \in \cup_S BNE(G, S)} \mathbb{E}_\nu[V] = \max_{\nu \in \cup_S BNE(G, S)} \sum_{a, \theta} V(a, \theta) \nu(a|\theta) \psi(\theta).$$

Since this objective is linear in $\nu(a|\theta)$, the designer's problem is to maximize a linear objective function over a non-empty convex polygon. By the fundamental theorem of linear programming, a solution exists and can be found at one of the corners of the constraint set. Denote by $\nu^* \in \cup_S BNE(G, S)$ the solution, that is, the BNE the designer would like to induce through her choice of information structure. We next characterise the information structure S^* that achieves this, 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 3. *Given a basic game G , an information structure $S = (T, \pi)$ is direct if $T_i = A_i$ for all $i \in I$.*

In other words, a direct information structure sends action recommendations as signals. The next proposition establishes the sufficiency of working with direct information structures, which is a "revelation principle" result.

PROPOSITION 2. *Given a basic game G , for every $\nu \in \cup_S BNE(G, S)$ there exist a direct information structure $S_\nu = (A, \nu)$ such that $\nu \in BNE(G, S_\nu)$.*

The proof 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. Thus, if the designer uses a direct information structure with the same probability distribution ν and actions as messages, it is interim incentive compatible for each agent to follow the action recommendation implied by the observed signal assuming that the other agents do so as well. This generates an BNE equilibrium distribution ν under that direct information structure, which in turn results in the same ex ante expected payoff for the designer. From here on, attention is restricted to direct information structures without loss of generality.

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

This corollary establishes the equivalence between optimal information structures and optimal BCE distributions over actions conditional on states. Given an optimal ν^* ⁶, the designer chooses a direct information structure where the signals are action recommendations and their distributions conditional on the state are given by ν^* . Under this structure, each agent has an incentive to follow the action recommendation given by the signal if all other agents do as well.

5 Application: Symmetric Binary Environments

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

5.1 Setup

There are $N = 2$ players, where i indexes the generic player, and j denotes 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 $u : A \times \Theta \rightarrow \mathbb{R}$ are symmetric. Further, we assume a common prior ψ , which is uniform on the two states, i.e. $\psi(\theta_0) = \psi(\theta_1) = \frac{1}{2}$. This specifies the basic game $G = (A^2, u, \psi)$ and we refer to this as a *symmetric binary environment*.

We parameterize the payoffs of the strategic interaction in the following way:

$\theta = \theta_0$	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

where $c \geq 0$ and $d \geq 0$. This non-negativity assumption 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 denote a basic game with parameters c and d by $G_{c,d}$.

The above payoff matrices assume that players have a preference for playing different actions in the different states of the world. If the same action were preferred in both states, there would be an equilibrium in strictly dominant strategies. In this case, the information

⁶Notice that ν^* need not be unique.

that players receive is irrelevant and there is no scope for information design. 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 refer to the agent that takes the action, i.e., a_k^i stands for agent i taking action a_k .

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 .

We refer to the preference of each player to coordinate his action with the state, for any given action of his opponent, as *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 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 his opponent is captured by the *strategic complementarity* (T), measured as:

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

for $k = 0, 1$. 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 a variety of strategic environments. For example, $c > d > 0$ represents an investment game where players want to coordinate on a project. The profitability of the project 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 one's opponent also invests in the same project. When $d > c > 0$, the payoffs capture a 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 the desired features, while the other one succeeds, then the second firm captures the whole market and obtains a higher payoff.

We consider a payoff function for the designer that is symmetric in the state and the agents' actions.

DEFINITION 4. A designer payoff function $V : A^2 \times \Theta \rightarrow \mathbb{R}$ is symmetric if:

$$(i) V(a_0, a_0, \theta_0) = V(a_1, a_1, \theta_1) = m,$$

$$(ii) V(a_0, a_1, \theta_0) = V(a_1, a_0, \theta_0) = V(a_0, a_1, \theta_1) = V(a_1, a_0, \theta_1) = l, \text{ and}$$

$$(iii) V(a_1, a_1, \theta_0) = V(a_0, a_0, \theta_1) = n,$$

for some $l, m, n \in \mathbb{R}$.

5.2 Designer's Optimization Problem and Solution

5.2.1 Constraint Set

To determine the constraint set, we first characterize the set of all possible BNE for the basic game $G_{c,d}$ under all possible information structures. By Proposition 1, we know that for a basic game G , the largest set of distributions over actions and states, which can be sustained as BNE under some information structure, is given by $BCE(G)$. In view of the symmetric and binary nature of the environment, it is without loss to restrict attention to distributions which are symmetric both with respect to the players' actions and the state. These can be fully described by two parameters, q and r . Hence, a symmetric distribution over action profiles conditional on states, $\nu(q, r)$, can be represented as follows:

$\nu(\cdot \theta_0)$	a_0	a_1	$\nu(\cdot \theta_1)$	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 $\Pr(a_0^i, a_0^j|\theta_0) + \Pr(a_0^i, a_1^j|\theta_0) = q$.

Our next result characterizes the set of symmetric $BCE(G_{c,d})$ distributions. We consider all possible values of the basic game parameters c and d which do not make the strategic interaction trivial. Given a set A , $Co(A)$ denotes the convex hull of A .

PROPOSITION 3. (BCE Distributions)

Consider the symmetric binary environment.

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

The proof of the proposition shows that the set of BCE distributions 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 prior.

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

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. The constraint set $BCE(G_{2,1})$ 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 agents have no information.

Figure 1: Constraint Set (Symmetric Example)

5.2.2 Objective function

If agents play a BNE distribution $\nu(q, r)$, the ex ante expected payoff of the designer can be written as

$$\mathbb{E}(V) = R \cdot r + Q \cdot q + const, \quad (5.3)$$

where

$$\begin{aligned}
R := & \psi(\theta_0) [V(a_1, a_1, \theta_0) - V(a_0, a_1, \theta_0) - V(a_1, a_0, \theta_0) + V(a_0, a_0, \theta_0)] \\
& + \psi(\theta_1) [V(a_1, a_1, \theta_1) - V(a_0, a_1, \theta_1) - V(a_1, a_0, \theta_1) + V(a_0, a_0, \theta_1)] \\
= & m + n - 2l,
\end{aligned} \tag{5.4}$$

and

$$\begin{aligned}
Q := & \psi(\theta_0) [V(a_0, a_1, \theta_0) - V(a_1, a_1, \theta_0) + V(a_1, a_0, \theta_0) - V(a_1, a_1, \theta_0)] \\
& + \psi(\theta_1) [V(a_0, a_1, \theta_1) - V(a_0, a_0, \theta_1) + V(a_1, a_0, \theta_1) - V(a_0, a_0, \theta_1)] \\
= & 2(l - n).
\end{aligned} \tag{5.5}$$

The coefficient R captures the “expected” preference for complementarity between the actions in the designer’s objective. It is a weighted average of the complementarities between the actions in each state, with the prior probabilities for each state as weights.

The coefficient 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 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 miscoordination 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 . Thus, Q measures the importance of unilateral coordination in the designer’s objective in expectation over the two states.

EXAMPLE. Suppose the designer benefits from mis-coordination between the agents’ actions irrespective of the state. That is, her payoffs are given by:

$$V(a^i, a^j, \theta) = \begin{cases} 1 & \text{if } a^i \neq a^j \\ 0 & \text{otherwise,} \end{cases} \tag{5.6}$$

which is symmetric in both the agents’ actions and the state. Substituting the values into (5.3), gives

$$\mathbb{E}(V) = -2r + 2q \tag{5.7}$$

as the designer’s *ex ante* expected payoff. 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 2).

5.2.3 Optimal Information Structure

We next maximize the designer’s objective function (5.3) over the constraint set. Let us denote by $\nu^*(q, r)$ a distribution which maximizes (5.3) over $BCE(G_{c,d})$ (notice that there might be multiple such distributions). 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 $\nu(q, r) \in BCE(G_{c,d})$ as BNE, can be parameterized in an analogous way, as given in Table 2. 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.

$\pi(\cdot \theta_0)$	a_0	a_1	$\pi(\cdot \theta_1)$	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 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 captured 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 $S_{1,1}$.

Before moving on, we showcase the optimal information structure in the symmetric binary example used throughout this section.

EXAMPLE. The symmetric BCE that maximizes the designer's objective 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 matrices of Table 3 and depicted in Figure 2. Under this information structure, the designer's ex ante expected payoff is $\frac{4}{5}$. Due to the symmetry of the binary environment and of the designer's payoff 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.

$\nu^*(\cdot \theta_0)$	a_0	a_1	$\nu^*(\cdot \theta_1)$	a_0	a_1
a_0	$\frac{1}{5}$	$\frac{2}{5}$	a_0	0	$\frac{2}{5}$
a_1	$\frac{2}{5}$	0	a_1	$\frac{2}{5}$	$\frac{1}{5}$

Table 3: Optimal Information Structure (Example)

Figure 2: Optimal Information Structure (Example)

Our next result is a complete characterization of the optimal information structure for all possible symmetric designer payoff functions and basic games $G_{c,d}$ in the symmetric binary environment. The different cases of the characterization theorem are summarized in Table 4 of Appendix B. 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 (5.4) and (5.5)). The information design problem reduces to a linear optimisation where the slope of the designer's level line, $-\frac{Q}{R}$, can be viewed as a marginal rate of substitution. It represents the tradeoff that the designer is willing to accept between the two parameters of the equilibrium distributions — the probability of state coordination (q) and the probability of action coordination (r).

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 $-\frac{Q}{R} > 2$.
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.

We next explain the intuition behind the general insights from this characterization. To begin with, notice that conditionally independent private signals are never strictly optimal. Indeed, conditionally independent private signals result in direct information structures characterized by $r = q^2$ for $q \in (0, 1)$ (excluding the purely public signal cases of $q \in \{0, 1\}$). Hence, these distributions are always in the interior of the constraint sets characterized in Proposition 3, and never at any of the vertices.

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

We next observe that full information is always optimal whenever the designer has a preference for both types of coordination — between the actions ($R > 0$) and between the state and the actions ($Q > 0$) — and that holds irrespective of whether the game is one of strategic complements or substitutes. Full information is also optimal whenever the designer's

preferences are “conflicting”, that is R and Q are of opposite signs, and the magnitude of the positive sign (coordination preference) is sufficiently strong.⁷ Therefore, whenever there is an overwhelmingly strong preference for coordination, be that for coordination between the agents’s actions or for coordination between their actions and the state, full information is optimal irrespective of the strategic forces present in the basic game.

Importantly, in contrast to the common perception that public signals are always optimal when the designer wants coordination between the players’ actions ($R > 0$), we show that this crucially depends on the sign and magnitude of the countervailing preference Q . Indeed, for sufficiently negative Q , the optimal information structure is private signals. We also show that the converse statement, that private signals are optimal when the designer wants the players’ actions to be uncorrelated, is also void: public signals are optimal in a lot of instances with $R < 0$, depending on the sign and magnitude of Q and on strategic substitutes/complements of the basic game. Therefore, the interplay between the designer’s preferences and the strategic forces inherent in the basic game determines the optimality of private versus public signals, as well as the precision of those signals.

For example, consider the case of $R > 0$ and $Q < 0$, i.e., 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 reveals the state fully and has the agents coordinate on both actions and the state. In contrast, if she really dislikes state coordination, she chooses 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.

There is no benefit from information design whenever revealing no information and letting the agents operate under their prior beliefs is optimal. As long as the optimal information structure differs from the null, the designer benefits from information design. In the symmetric binary setting, 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 $-\frac{Q}{R} < 2$. In this case, the designer has a prevailing preference for mis-coordination between the agents’ actions, and revealing no information would maximize the probability of mis-coordinated actions under the uniform prior. In all remaining cases, information design is beneficial.

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.*

To contrast our results with those from the literature on cheap talk without commitment

⁷For the case of $R > 0$ and $Q < 0$, R needs to be greater than Q in absolute value, while for the case of $Q > 0$ and $R < 0$, Q needs to be twice as large as R in absolute value.

(Farrell and Gibbons (7)), 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.⁸ 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, no information is optimal in this case for certain values of the parameters.

5.3 Mechanism and Information Design

In this section we compare the use of mechanism and information design in the context of the binary symmetric setting. Naturally, the designer can always do (weakly) better when she uses the tools of mechanism and information design combined. We first address the question of when the two can be viewed as substitutes in terms of achieving the same designer payoff. Then, we look at when the designer can strictly benefit from using ex post balanced transfers in addition to optimal information provision.

There are many ways in which the mechanism design problem can be formulated. In order to level the playing field and make the comparison between information and mechanism design meaningful, we allow for transfers contingent on both the realized state and action profiles. However, we also require ex post budget balance to ensure the mechanism is costless to the designer, as are all the information structures. We further restrict attention to symmetric transfers. Symmetry and ex post budget balance together are incorporated in the following parameterization of the transfers:

- (i) $t_i(a_0, a_0, \theta_k) = t_i(a_1, a_1, \theta_k) = 0$ for $i = 1, 2, k = 0, 1$,
- (ii) $t_1(a_0, a_1, \theta_0) = t_2(a_1, a_0, \theta_0) = t_1(a_1, a_0, \theta_1) = t_2(a_0, a_1, \theta_1) = s$, and
- (iii) $t_1(a_1, a_0, \theta_0) = t_2(a_0, a_1, \theta_0) = t_1(a_0, a_1, \theta_1) = t_2(a_1, a_0, \theta_1) = -s$,

for some $s \in \mathbb{R}$.

Writing down the BNE incentive constraints for the agents when they act under their prior beliefs, it becomes obvious that the above transfers are never useful in expanding the set of equilibria. That is, the only equilibria that mechanism design with ex post balanced

⁸In the binary framework of Farrell and Gibbons (7), making the preferences of the designer (sender) and agents (receivers) more aligned can be interpreted as assuming positive values for the designer's payoffs from each agent coordinating with the state ($v_i, w_i > 0$ for $i = 1, 2$). This ensures existence of a separating equilibrium with full information revelation (their Proposition 1), which is preferred by all parties. Translated into our framework, this corresponds to $R = 0$ and $Q > 0$. It is this absence of preference for complementarity between the actions of the agents ($R = 0$), stemming from the additive separability of the designer payoffs, that drives the different results. We have this added dimension in the preferences of both the designer and the agents, which needs to be taken into consideration.

Figure 3: Mechanism and Information Design (Example)

transfers can support are those that would naturally exist under the prior in the absence of transfers: both agents play a_0 , both agents play a_1 , and both agents mix with probability $1/2$. Therefore, mechanism design is never a substitute for information design as it can only achieve the outcomes that would prevail under no information provision, i.e. if the designer were to choose the null information structure.

Next, we look at how the use of transfers in conjunction with information design alters the constraint set. The BCE constraints boil down to the following inequality:

$$2(c - d)r \geq d + s + (c - 3d - 2s)q. \quad (5.8)$$

Indeed, manipulating this constraint through the choice of s might be beneficial to the designer, depending on the parameters of the problem. Notice also that for any $s \in \mathbb{R}$, the interim participation constraints for any symmetric BCE distribution are always satisfied.

While it is not possible to draw general and easily interpretable conclusions about when the use of transfers in addition to information is strictly beneficial for the designer, we demonstrate the complementarity between the two in the context of our running example.

EXAMPLE. *We next show that using state-and-action-contingent ex post balanced transfers leads to a higher designer payoff at the new optimal information structure. Setting $s = 1$, the BCE constraint pivots (from the dashed to the solid red line) as illustrated in Figure 3. Also depicted are the new optimal information structure ($q = 4/7$, $r = 1/7$) and level curve associated with an expected payoff of $6/7$ for the designer (higher than the previously optimal payoff of $4/5$). Indeed, further increasing s would lead to the BCE constraint converging to the vertical purple line, and the designer's expected payoff converging to 1.*

5.4 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.

Suppose that the designer can condition the new information she releases on both the state and the agents' signals, either because she observes those or has the technology to condition the new signals in such a way. Then, she can induce any distribution $\nu : \Theta \rightarrow \Delta(A)$ over action profiles conditional on states such that

$$\nu(a|\theta) = \sum_t \tilde{\pi}(t|\theta) \tilde{\nu}(a|t, \theta)$$

where for each $i \in I$, $t_i \in T_i$, and $a_i \in A_i$, $\tilde{\nu} : \Theta \times T \rightarrow \Delta(A)$ satisfies

$$\begin{aligned} \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \tilde{\pi}(t_i, t_{-i}|\theta) \tilde{\nu}(a_i, a_{-i}|t_i, t_{-i}, \theta) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \tilde{\pi}(t_i, t_{-i}|\theta) \tilde{\nu}(a_i, a_{-i}|t_i, t_{-i}, \theta) u_i((a'_i, a_{-i}), \theta), \end{aligned} \quad (5.9)$$

for all $a'_i \in A_i$. This corresponds to the BCE-constraint (4.3) adjusted for the fact that agents have observed signals drawn from $\tilde{\pi}$.

More specifically, consider the symmetric binary example from this section. Let us assume that the agents have already observed signals from the following information structure $\tilde{\pi}$:

$\tilde{\pi}(\cdot \theta_0)$	t_0	t_1	$\tilde{\pi}(\cdot \theta_1)$	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

which is common knowledge between all parties. It turns out that the best the designer can do in this situation is to not release any additional information and let the agents play according to the signals they have received through $\tilde{\pi}$. Indeed, for $(G_{2,1}, \tilde{\pi})$ the following behavioral strategy is a BNE:

$$\beta_i^*(a_i|t_i) = \begin{cases} 1, & \text{if } a_i = a_0 \text{ and } t_i = t_0, \text{ or } a_i = a_1 \text{ and } t_i = t_1 \\ 0, & \text{otherwise.} \end{cases} \quad (5.10)$$

This BNE would result in an expected payoff of $\frac{7}{10}$ for the designer, which is the maximum she

can achieve when the agents have prior information according to $\tilde{\pi}$. Indeed, this is lower than the maximum expected payoff of $\frac{4}{5}$ she could achieve when the agents were not exogenously informed.

Thus, the presence of prior information may significantly alter the optimal information structure or designer expected payoff. Similar considerations apply when the agents observe additional signals concurrently or subsequently to those from the designer's information structure. As long as the designer observes the distribution of the exogenous signals and can condition her information on them, as well as on the state, she can incorporate the constraints (5.9) into her optimization problem. These constraints may not always affect her ability to achieve the same maximum expected payoff as compared to when the agents have no exogenous information. Nevertheless, the form of the optimal information structure would change, as it would have to account for the informational content of the exogenous signals.

Notice that for the purposes of this discussion we have assumed that the designer has the largest possible amount of flexibility: she is able to send new signals conditional both on the exogenous signals of the agents, which are private information, and on the state. However, one could also imagine a situation where the designer is not able to do this, unless she first elicits the signals from the agents. This would add additional truth-telling constraints to the problem and would generally limit the range of distributions the designer can induce. An even more restrictive situation arises when the designer is not able to elicit the exogenous signals, for example because there is no way of communicating with the agents. In that case, she has to provide action recommendations for each possible type of each agent, which she can condition only on the state. These three scenarios are, in general, increasingly restrictive for the designer in terms of the distributions she can induce through the release of further information. In our particular example, however, she can do no better by knowing the exogenous signals of the agents than if she had to elicit them or if she had no way of doing so them. Indeed, her optimal choice would be to reveal no further information under any of these three regimes.

5.5 Additional Information

One implication of the commitment assumption is that the designer is not allowed to deviate and provide additional information upon the realization of a signal, which is unfavourable to her. In the single-agent case, Kamenica and Gentkow (10) (Lemma 2) show that this is not a restrictive assumption, as under the optimal information structure the designer would never find such a deviation beneficial. This result no longer holds in environments with multiple strategic agents where the designer observes the whole realized signal profile. In this case, unless the signal is public, the designer is always more informed than each individual agent and has a different posterior belief about the state. As the following example demonstrates, she can use this to her advantage and generate additional signals which would change the equilibrium play.

In the symmetric binary example, consider a designer who has chosen the optimal information structure in Table 3 and observes a realized signal (a_0, a_0) . In this case, each agent has privately observed a_0 , and attaches a posterior belief of $\frac{3}{5}$ to the state θ_0 , while the designer knows that the state is θ_0 with certainty. If she were to reveal no further information, her expected payoff would be $V(a_0, a_0, \theta_0) = 0$. However, she can do better by sending additional signals (denoted by tilde) in the following way:

θ_0	$a_0, \tilde{a}_0$	$a_0, \tilde{a}_1$	$a_1, \tilde{a}_0$	$a_1, \tilde{a}_1$	θ_1	$a_0, \tilde{a}_0$	$a_0, \tilde{a}_1$	$a_1, \tilde{a}_0$	$a_1, \tilde{a}_1$
$a_0, \tilde{a}_0$	0	1/5	2/5	0	$a_0, \tilde{a}_0$	0	0	2/5	0
$a_1, \tilde{a}_0$	2/5	0	0	0	$a_1, \tilde{a}_0$	1/5	0	0	0
$a_1, \tilde{a}_1$	0	0	0	0	$a_1, \tilde{a}_1$	0	1/5	0	1/5

This information structure is a combination of S^* and the new signals released by the designer. Under this combined information structure, which the designer makes common knowledge, it is incentive compatible for both agents to follow the action recommendation of the second signal irrespective of the initially observed action recommendation. In other words, the behavioral strategy

$$\beta_i^*(a_i | a'_i, \tilde{a}_i) = \begin{cases} 1, & \text{if } a_i = \tilde{a}_i \\ 0, & \text{if } a_i \neq \tilde{a}_i \end{cases} \quad (5.11)$$

for $i = 1, 2$ is a BNE under the above combined information structure. By doing this, the designer would increase her expected payoff to 1, since conditional upon (a_0, a_0, θ_0) the agents will receive the new signal profile $(\tilde{a}_0, \tilde{a}_1)$ with certainty, which under the above BNE leads to perfect miscoordination of their actions. Therefore, this is a profitable deviation for the designer, even after choosing the optimal information structure to begin with.⁹

6 Conclusion

In environments with incomplete information, the incentives of rational agents to behave in a certain way are determined by the extensive form of the underlying strategic interaction and by the informational environment in which they interact. A designer would like to maximize the value of her objective, which depends on the agents' equilibrium behavior. Mechanism design looks at the optimal choice of an extensive form to induce the desired equilibrium, while holding the informational environment fixed. This paper lays out the methodology of

⁹Notice that in this example, if the agents are aware of the designer's objective, they could infer that she has incentive to release additional information only in case signal profiles (a_0, a_0) or (a_1, a_1) have realized. Therefore, each agent, upon observing his initial action recommendation and the fact that the designer is releasing further information, would be able to infer which state of the world has realized. However, if we assume that agents are unaware of the designer's objective, then she does benefit from the release of additional signals. This is specific to the binary example considered here and is not necessary to assume in general, in order to demonstrate the designer's incentive to deviate from the optimal information structure upon realization of certain signal profiles.

information design in static and finite settings, which focuses on the analogous problem of choosing the informational environment optimally, while taking the extensive form as given.

There are a number of extensions and robustness issues that may constitute interesting directions for future research. An important assumption of the current framework is that agents are unable to engage in strategic communication and share their privately observed signals with each other. If the underlying basic game is such that communication is strategically beneficial under some information structures, this assumption becomes particularly pertinent. Another important assumption is the focus on designer preferred equilibria, which is implicitly subsumed in the linear programming approach described here. In the single-agent environment this assumption is innocuous, as it boils down to indifferences. In the multiple-strategic-agent environment, however, it is salient. Indeed, equilibrium multiplicity can be detrimental to the designer in many settings of interest, which suggests the need for different selection criteria to be incorporated into the information design framework. For example, a pessimistic designer would like to maximize her expected payoff under adversarial equilibrium selection, that is, assuming agents will coordinate on the worst equilibrium from the designer’s perspective. A further restriction of the methodology outlined in this paper is its focus on Bayes Nash equilibrium as a solution concept, which relies on common knowledge of rationality. Relaxing this and assuming agents have bounded depths of reasoning may provide interesting insights into behavioral implications for optimal information design. The common prior assumption is yet another dimension on which the current framework can be extended and modified to incorporate heterogeneous prior beliefs.

Appendix A

Proof of Proposition 1

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

$$\sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a'_i, a_{-i}), \theta) \quad (6.1)$$

for each $i \in I$, $a_i \in A_i$ and $a'_i \in A_i$. Consider the information structure $S^* = (A, \pi^*)$ with

$$\pi^*(a | \theta) = \nu(a | \theta) \quad (6.2)$$

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

$$\beta_i^*(a_i | a'_i) = \begin{cases} 1, & \text{if } a_i = a'_i \\ 0, & \text{if } a_i \neq a'_i \end{cases} \quad (6.3)$$

for all $a_i, a'_i \in A_i$. The interim payoff to agent i observing signal a_i and choosing action a'_i when his opponents follow β_{-i}^* is

$$\begin{aligned} \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi^*(a_i, a'_{-i} | \theta) \left(\prod_{j \neq i} \beta_j^*(a_j | a'_j) \right) u_i((a'_i, a_{-i}), \theta) \\ = \sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (6.4)$$

where we use (6.2) and (6.3). Therefore, the BNE interim incentive compatibility constraint

$$\begin{aligned} \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi^*(a_i, a'_{-i} | \theta) \left(\prod_{j \neq i} \beta_j^*(a_j | a'_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi^*(a_i, a'_{-i} | \theta) \left(\prod_{j \neq i} \beta_j^*(a_j | a'_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (6.5)$$

is equivalent to and implied by the BCE constraint (6.1). Hence, given the strategy profile β^* and by Definition 1, the distribution over actions conditional on states given by

$$\sum_{a' \in A} \pi^*(a' | \theta) \left(\prod_{j=1}^N \beta_j(a_j | a'_j) \right) = \nu(a | \theta). \quad (6.6)$$

is a BNE of (G, S^*) , i.e., $\nu \in BNE(G, S^*)$. This implies $BCE(G) \subseteq \cup_S BNE(G, S)$.

Next we prove that $BCE(G) \supseteq \cup_S BNE(G, S)$. 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 β such that

$$\tilde{\nu}(a|\theta) = \sum_{\tilde{t} \in \tilde{T}} \tilde{\pi}(\tilde{t}|\theta) \left(\prod_{j=1}^N \beta_j(a_j|\tilde{t}_j) \right). \quad (6.7)$$

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

$$\begin{aligned} \sum_{a_{-i}, \tilde{t}_{-i}, \theta} \psi(\theta) \tilde{\pi}(\tilde{t}_i, \tilde{t}_{-i}|\theta) \left(\prod_{j \neq i} \beta_j(a_j|\tilde{t}_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, \tilde{t}_{-i}, \theta} \psi(\theta) \tilde{\pi}(\tilde{t}_i, \tilde{t}_{-i}|\theta) \left(\prod_{j \neq i} \beta_j(a_j|\tilde{t}_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (6.8)$$

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

$$\begin{aligned} \sum_{a_{-i}, \tilde{t}_{-i}, \theta} \psi(\theta) \tilde{\pi}(\tilde{t}_i, \tilde{t}_{-i}|\theta) \left(\prod_{j=1}^N \beta_j(a_j|\tilde{t}_j) \right) u_i((a_i, a_{-i}), \theta) \\ \geq \sum_{a_{-i}, \tilde{t}_{-i}, \theta} \psi(\theta) \tilde{\pi}(\tilde{t}_i, \tilde{t}_{-i}|\theta) \left(\prod_{j=1}^N \beta_j(a_j|\tilde{t}_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (6.9)$$

which by (6.7) is equivalent to

$$\sum_{a_{-i}, \theta} \psi(\theta) \tilde{\nu}(a|\theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \tilde{\nu}(a|\theta) u_i((a'_i, a_{-i}), \theta). \quad (6.10)$$

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

Proof of Lemma 1

The set $BCE(G)$ is the collection of distributions $\nu : \Theta \rightarrow \Delta(A)$ 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 $a'_i \in A_i$.

Constraints (i) and (ii) ensure each ν is a proper probability distribution. Constraints (iii) are the BCE incentive constraints (4.3). All constraints are linear in $\nu(a|\theta)$. By Theorem A

of Stinchcombe (15), the set of $BCE(G)$ is non-empty. Therefore, $BCE(G)$ is a non-empty convex polygon. $\square$

Proof of Proposition 2

Take a basic game G and a distribution $\nu \in \cup_S BNE(G, S)$. By Proposition 1 we know that $\nu \in BCE(G)$ and hence,

$$\sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a_i, a_{-i}), \theta) \geq \sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a'_i, a_{-i}), \theta) \quad (6.11)$$

for each $i \in I$, $a_i \in A_i$ and $a'_i \in A_i$. Consider the direct information structure $S_\nu = (A, \pi_\nu)$ with $\pi_\nu(a | \theta) = \nu(a | \theta)$ for all $a \in A$ and $\theta \in \Theta$. In the game (G, S_ν) consider the following behavioral strategy β_i for agent i :

$$\beta_i(a_i | a'_i) = \begin{cases} 1, & \text{if } a_i = a'_i \\ 0, & \text{if } a_i \neq a'_i \end{cases} \quad (6.12)$$

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

$$\begin{aligned} & \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi_\nu(a_i, a'_{-i} | \theta) \left(\prod_{j \neq i} \beta_j(a_j | a'_j) \right) u_i((a'_i, a_{-i}), \theta) \\ &= \sum_{a_{-i}, \theta} \psi(\theta) \pi_\nu(a_i, a_{-i} | \theta) u_i((a'_i, a_{-i}), \theta) = \sum_{a_{-i}, \theta} \psi(\theta) \nu(a_i, a_{-i} | \theta) u_i((a'_i, a_{-i}), \theta), \end{aligned} \quad (6.13)$$

where the first equality follow by (6.12) and the second equality follows from $\pi_\nu(a | \theta) = \nu(a | \theta)$ for all $a \in A$ and $\theta \in \Theta$. By (6.11) and (6.13) we obtain

$$\begin{aligned} & \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi_\nu(a_i, a'_{-i} | \theta) \left(\prod_{j \neq i} \beta_j(a_j | a'_j) \right) u_i((a_i, a_{-i}), \theta) \\ & \geq \sum_{a_{-i}, a'_{-i}, \theta} \psi(\theta) \pi_\nu(a_i, a'_{-i} | \theta) \left(\prod_{j \neq i} \beta_j(a_j | a'_j) \right) u_i((a'_i, a_{-i}), \theta) \end{aligned} \quad (6.14)$$

for all i , which gives the interim incentive compatibility conditions for strategy profile β to constitute a BNE of (G, S_ν) . The distribution of actions conditional on states under β and S_ν is

$$\sum_{a' \in A} \pi_\nu(a' | \theta) \left(\prod_{i=1}^N \beta_i(a_i | a'_i) \right) = \pi_\nu(a | \theta) = \nu(a | \theta) \quad (6.15)$$

for all $a \in A$ and $\theta \in \Theta$. Hence, $\nu \in BNE(G, S_\nu)$. $\square$

Proof of Proposition 3

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

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

for all i . These two constraints are equivalent and reduce to only one inequality:

$$2(c-d)r \geq d + (c-3d)q. \quad (6.16)$$

Additionally, the parameters need to satisfy:

$$r \leq q \quad (6.17)$$

$$r \geq \max\{2q-1, 0\} \quad (6.18)$$

and

$$q \in [0, 1]. \quad (6.19)$$

Therefore, the set $BCE(G_{c,d})$ is equivalent to the set of (q, r) -pairs which satisfy constraints (6.16)–(6.19). In the following figures, constraint (6.17) is represented by the 45-degree dashed diagonal, while constraints (6.16) and (6.18) are depicted in red and blue respectively.

Case 1: Assume $c > d \geq 0$ (strategic complements). The BCE constraint (6.16) can thus be written as:

$$r \geq \frac{d}{2(c-d)} + \frac{c-3d}{2(c-d)}q \quad (6.20)$$

In this case, constraint (6.19), which essentially coincides with the x -axis of the graph, is never binding. The reason behind this is the following. The intercept of constraint (6.20) (depicted in red on Figure 4) is always positive. When in addition $c \geq 3d$, the slope is also positive. Hence, this constraint is always more binding than (6.19), as it always lies above the x -axis. On the other hand, when $c < 3d$, the slope of (6.20) is negative. However, it is easy to show that (6.20) intersects (6.18) (depicted in blue) before it intersects the x -axis. Therefore, for the relevant range of values, (6.20) lies above the x -axis also in this case. Hence, (6.19) is never binding.

The set of distributions which satisfy (6.20), (6.17) and (6.18) 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 (6.20) and (6.17)),

Figure 4: Strategic Complements ($c > d \geq 0$)

$(q_2, r_2) = \left(\frac{2c-d}{3c-d}, \frac{c-d}{3c-d}\right)$ (of (6.20) and (6.18)) and $(q_3, r_3) = (1, 1)$ (of (6.17) and (6.18)).

Case 2: Assume $d > c \geq 0$ (strategic substitutes). The obedience constraint (6.16) can thus be written as:

$$r \leq \frac{d}{2(c-d)} + \frac{c-3d}{2(c-d)}q \quad (6.21)$$

This constraint has a negative intercept and a positive slope (depicted in red on Figure 5). 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 (6.21) intersects only constraints (6.19) and (6.17). In this case, all four constraints (6.21), (6.17), (6.18) and (6.19) 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 (6.21) and (6.17)), $(q_3, r_3) = (1, 1)$ (of (6.17) and (6.18)), $(q_5, r_5) = \left(\frac{1}{2}, 0\right)$ (of (6.18) and (6.19)) and $(q_4, r_4) = \left(\frac{d}{3d-c}, 0\right)$ (of (6.19) and (6.21)).

When $c = 0$, the slope of (6.21) is exactly equal to $\frac{3}{2}$. In this case (6.21), (6.17) and (6.18) all intersect at one point — $(q_3, r_3) = (1, 1)$ — and (6.17) is never binding. The set of BCE distributions, in this case, 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 (6.16) becomes $q \geq \frac{1}{2}$ (depicted in red on Figure 6). The set of BCE distributions 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 5: Strategic Substitutes ($d > c \geq 0$)

Figure 6: $c = d > 0$

Appendix B

Table 4: 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(2016), “Persuading Voters,” *American Economic Review*, 106(11), 3590-3605. 6
- [2] Angeletos, G.-M., A. Pavan(2009), “Policy with Dispersed Information,” *Journal of the European Economic Association*, 7(1), 11-60.
- [3] Bergemann, D., S. Morris (2016a), “Bayes Correlated Equilibrium and the Comparison of Information Structures,” *Theoretical Economics*, 11, 487-522. 3, 5, 6, 12
- [4] Bergemann, D., S. Morris (2016b), “Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium,” *American Economic Review: Papers and Proceedings*, 106(5), 586-591. 2
- [5] Crawford, V., J. Sobel (1982), “Strategic Information Transmission,” *Econometrica*, 50(6), 1431-1451. 5
- [6] Eliaz, K., F. Forges (2012), “Disclosing Information to Interacting Agents,” *mimeo*. 6
- [7] Farrell, J., R. Gibbons (1989), “Cheap Talk with Two Audiences,” *American Economic Review*, 79(5), 1214-1223. 4, 5, 23
- [8] Goltsman, M., G. Pavlov (2011), “How to Talk to Multiple Audiences,” *Games and Economic Behavior*, 72, 100-122. 5
- [9] Gossner, O. (2000), “Comparison of Information Structures,” *Games and Economic Behavior*, 30, 44-63. 6, 10
- [10] Kamenica, E., M. Gentzkow (2011), “Bayesian Persuasion,” *American Economic Review*, 101, 2590-2615. 4, 5, 9, 26
- [11] Lehrer, E., D. Rosenberg, E. Shmaya (2010), “Signaling and Mediation in Games with Common Interest,” *Games and Economic Behavior*, 68, 670-682. 6, 10
- [12] 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
- [13] 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
- [14] Peski, M. (2008), “Comparison of Information Structures in Zero-Sum Games,” *Games and Economic Behavior*, 62, 732-735. 6
- [15] Stinchcombe, M. B. (2011), “Correlated Equilibrium Existence for Infinite Games with Type-Dependent Strategies,” *Journal of Economic Theory*, 146(2), 638-655. 31
- [16] Wang, Y. (2012), “Bayesian Persuasion with Multiple Receivers,” *WISE Working Papers*, WISE, Xiamen University. 6