

INFORMATION DESIGN, BAYESIAN PERSUASION
AND BAYES CORRELATED EQUILIBRIUM

By

Dirk Bergemann and Stephen Morris

January 2016

COWLES FOUNDATION DISCUSSION PAPER NO. 2027



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Information Design, Bayesian Persuasion and Bayes Correlated Equilibrium*

Dirk Bergemann[†]

Stephen Morris[‡]

January 12, 2016

Abstract

A set of players have preferences over a set of outcomes. We consider the problem of an "information designer" who can choose an information structure for the players to serve his ends, but has no ability to change the mechanism (or force the players to make particular action choices). We describe a unifying perspective for information design. We consider a simple example of Bayesian persuasion with both an uninformed and informed receiver. We extend information design to many players and relate it to the literature on incomplete information correlated equilibrium.

KEYWORDS: Information Design, Bayesian Persuasion, Bayes Correlated Equilibrium, Information Structure.

JEL CLASSIFICATION: C72, D82, D83.

*We acknowledge financial support from NSF SES 1459899. We would like to thank our discussant, Drew Fudenberg, and Ben Brooks, Jeff Ely, Emir Kamenica, Laurent Mathevet and Tymofiy Mylanov for informative comments, and Xinyang Wang for valuable research assistance.

[†]Department of Economics, Yale University, New Haven, U.S.A., dirk.bergemann@yale.edu.

[‡]Department of Economics, Princeton University, Princeton, U.S.A., smorris@princeton.edu.

1 Introduction

A set of players have preferences over a set of outcomes. Consider the problem of an "information designer" who can choose an information structure for the players to serve his ends, but has no ability to change the mechanism (or force the players to make particular action choices). A mechanism here describes the set of players, their available actions and a mapping from action profiles to outcomes. Contrast this "information design" problem with the "mechanism design" problem, where a "mechanism designer" can choose a mechanism for the players to serve his ends, but has no ability to choose the information structure (or force the players to make particular action choices).¹ In each case, the problem is sometimes studied with a restricted choice set. In the information design problem, we could restrict the designer to choose whether the players are given no information or complete information about the environment. In the mechanism design problem, we could restrict the designer to choose between a first price and a second price auction. However, in each case, there is a revelation principle argument that allows for the analysis of all information structures or all mechanisms respectively. For the mechanism design problem, we can restrict attention to direct mechanisms where the players' action sets are equal their possible types. Conversely, for the information design problem, we can restrict attention to information structures where the players' type sets are equal to their action sets. In this note, using this observation, we consider information design problems when all information structures are available to the designer.

When there are many players, but the information designer (or "mediator") has *no informational advantage* over the players, this problem reduces to the analysis of communication in games (Myerson (1991)). If there is only one player (or "receiver") but the information designer (or "sender") has an informational advantage over the player, the problem reduces to the "Bayesian persuasion" problem (Kamenica and Gentzkow (2011)). Some of our recent work corresponds to the information design problem when there are *both* many players *and* the information designer has an informational advantage over the players (Bergemann and Morris (2013), (2016)).

This note explores this unifying perspective on information design. In the next section, we discuss the simplest example of Bayesian persuasion, with both an uninformed and informed receiver. We provide a couple of novel perspectives with this treatment. First, we consider "omniscient persuasion" in the informed receiver case, where the sender knows the receiver's signal, contrasting this with the more usual assumptions that the receiver either cannot condition on the receiver's signal or can only do so if he can elicit this information from the receiver. Second, we use a two step procedure to solve the problem, by first characterizing the set of outcomes that could be attained by the sender and then analyzing the sender's

¹We follow Taneva (2015) in our use of the term "information design".

choice of outcome among those that are attainable.² These novel perspectives are of independent interest for the Bayesian persuasion literature. But more importantly, they also clarify how the analysis extends to the many player case. In the final section, we report the extension to the many player case, discuss the connection to the incomplete information correlated equilibrium literature and survey our own applied and theoretical work in this area. While the discussion in this note is informal, a companion piece (Bergemann and Morris (2016b)) discusses these connections formally.

2 A Bayesian Persuasion Example

A bank is solvent in a good state (G) and insolvent in a bad state (B). A depositor can either run (r) or stay (s) with the bank. Each state of the world is equally likely a priori. A regulator can design the depositor's information in order to influence his behavior. For the depositor, the payoff from staying with the bank is -1 if the bank is insolvent, and y if solvent, with $0 < y < 1$; the payoff for running is normalized to 0 in either state. The regulator seeks to minimize the probability of the depositor running. This is the leading example of Kamenica and Gentzkow (2011), with the information designer being a regulator (instead of a prosecutor) and the single player being a depositor (instead of a judge). We rephrase this example in order to tie the analysis with the many player generalization discussed in Section 3.

2.1 The Uninformed Depositor Case

We briefly review the analysis of Bayesian persuasion with an uninformed depositor (receiver). We describe the receiver's behavior by a decision rule specifying his behavior given the true state of the world, writing ρ_θ for the probability that he will run in each state $\theta \in \{B, G\}$; thus a decision rule is a pair (ρ_G, ρ_B) . We can think of a decision rule as being a (stochastic) action recommendation from an informed mediator. A decision rule is *obedient* if the depositor always has an incentive to follow the action recommendation. The depositor will then have an incentive to stay if

$$(1/2)(1 - \rho_G)y - (1/2)(1 - \rho_B) \geq 0, \tag{1}$$

and an incentive to run if

$$0 \geq (1/2)\rho_G y + (1/2)\rho_B(-1). \tag{2}$$

Obedience conditions reflect the fact that the agent may have information (that we do not need to describe explicitly) that leads him to act differently across the two states, hence ρ_G and ρ_B may differ in value.

²We thus do not appeal to a concavification argument (Aumann and Maschler (1995)) which is useful to solve this problem at least in specific settings (Kamenica and Gentzkow (2011)).

Since $y < 1$, (1) is always the binding constraint and we can rewrite (1) as

$$\rho_B \geq 1 - y + y\rho_G. \quad (3)$$

Thus in any obedient decision rule, the probability of running in the bad state has to exceed the probability of running in the good state, with the difference being at least $1 - y$. In particular, staying with the bank for sure can never be an equilibrium. The regulator's preferred outcome, with the lowest probability of running across states, has $\rho_G = 0$ and $\rho_B = 1 - y$.

Now any obedient decision rule corresponds to optimal behavior under some information structure. For the regulator's preferred outcome, it suffices to give the depositor a bad signal with probability $1 - y$ if the state is bad, and otherwise always give the depositor a good signal. From the point of view of the motivating example, this corresponds to a regulator running stress tests to obfuscate the state of a bank in order to prevent a run. By pooling good and bad states in this way, the depositor is made indifferent between running and staying. More generally, any obedient decision rule can always arise when the depositor is given a signal equal to his action recommendation.

2.2 The Informed Depositor Case

Now suppose that the depositor receives information, independent of the regulator. We assume that the depositor receives a signal which is "correct" with probability $q > 1/2$. Formally, the depositor observes a signal g or b , with signals g and b being observed with conditionally independent probability q when the true state is G or B respectively.

The depositor's information will act as a constraint on the ability of the regulator to influence the depositor's decision, since he now has less control over the depositor's information. In this enriched setting, a decision rule specifies the probability that the depositor receives a recommendation to run, as a function of both the state and the signal. We will write $\rho_{\theta t}$ for the probability of running in state $\theta \in \{G, B\}$ if the signal is $t \in \{g, b\}$. A decision rule is now described by the quadruple $(\rho_{Bb}, \rho_{Gb}, \rho_{Bg}, \rho_{Gg})$.

Omniscient Persuasion. The analysis of the informed depositor case will depend on what the regulator knows about the depositor's information. We first consider the case where the regulator knows the information of the depositor *as well as* the state and, in this sense, is omniscient. This case has not been the focus of the Bayesian persuasion literature with an informed receiver. However, it is a natural case and, in some cases, the most natural case to consider. For example, a regulator may know the information available to the depositor (in the form of newspapers, public reports, etc.). And while the regulator may be unable to suppress that information, he may be able to condition the information he releases on the depositor's initial information.

The obedience constraints now reflect the conditional belief of the agent about the state given the realization of the signal. The obedience condition for depositor who observes a good signal and an action recommendation to stay is

$$q(1 - \rho_{Gg})y - (1 - q)(1 - \rho_{Bg}) \geq 0. \quad (4)$$

The obedience condition for a depositor who observes a good signal but an action recommendation to run is:

$$0 \geq q\rho_{Gg}y - (1 - q)\rho_{Bg}. \quad (5)$$

As long as the private information of the agent is sufficiently noisy, or $q \leq 1/(1 + y)$, the binding constraint is (4), otherwise it is the inequality (5) that determines the conditional probabilities. The obedience conditions for the bad type are derived in an analogous manner. The obedience conditions are defined type by type and we compute the restrictions on the conditional probabilities averaged across types. Now the decision rule $(\rho_{Bb}, \rho_{Gb}, \rho_{Bg}, \rho_{Gg})$ will induce behavior (ρ_B, ρ_G) integrating over signals $t \in \{g, b\}$.

Proposition 1 (Omniscient Persuasion) *The probabilities (ρ_B, ρ_G) form an equilibrium outcome for some information structure if:*

$$\rho_B \geq \max\{q(1 + y), 1\} - y + y\rho_G. \quad (6)$$

The behavior of the equilibrium set as a function of the precision q of the private information is illustrated in Figure 2.2. As the private information becomes precise and q approaches one, we converge to the complete information outcome with $\rho_B = 1, \rho_G = 0$. The depositor's private information thus limits the regulator's ability to influence the depositor's decision as the private information tightens obedience constraints.

But if the regulator had to elicit the depositor's information, the constraints imposed by the private information would become even more severe. We briefly contrast omniscient persuasion with this case.

Private Persuasion. We now consider a screening problem where the regulator offers a recommendation, a probability of running as a function of the reported type and the true state. Kolotilin, Li, Mylovanov, and Zapechelnuyk (2015) refer to this informational environment as "private persuasion".³ We now have three set of constraints. First, each type has to truthfully report his type; second, each type has to be willing to follow the recommendation, the obedience constraints; and third, double deviations, by means of misreporting and disobeying at the same time are not profitable. With these additional constraints, the set of outcomes that can arise in equilibrium with private persuasion is strictly smaller than under omniscient persuasion as can be seen by comparison between the blue and red areas in Figure 2.2. We

³Bergemann, Bonatti, and Smolin (2015) analyze this environment with monetary transfers.

can verify that the truthtelling constraints impose restrictions on how the differences in the conditional probabilities across signals can vary across states. This imposes additional restrictions on the ability of the designer when he seeks to attain either very low and very high running probabilities in both states.




Figure 2.2. Bayesian persuasion with an informed receiver (blue)
and private persuasion (red)
for $y = 9/10$ and $q = 0.575, 0.7, 0.825$.

Public Persuasion. A yet more restrictive model of persuasion with an informed receiver is to assume that the sender not only does not know the receiver's private information but cannot elicit it. Kolotilin, Li, Mylovanov, and Zapechelnuyk (2015) call this scenario "public persuasion." Such public persuasion has been the focus of the recent literature. In our example, one can show that any outcome that is attainable by private persuasion is attainable by public persuasion. Kolotilin, Li, Mylovanov, and Zapechelnuyk (2015) identify sufficient conditions for an equivalence.

3 The Many Player Case

Omniscient persuasion is the one player version of an approach to information design that we have been pursuing in recent work (Bergemann and Morris (2016a)). As in the analysis of omniscient persuasion presented here, our work suggests a two step procedure to information design. First, what are the set of outcomes that could arise for all extra information with which the players could be endowed? Second, which of those outcomes would be chosen by some interested party (and what is the information structure giving rise to those outcomes)?

We have analyzed the first step with many players by giving a general characterization of the set of outcomes that could arise in some equilibrium (Bergemann and Morris (2016a)). We consider a joint distribution of states, initial information signals and action profiles that satisfy the relevant obedience constraints: that is, the constraint that each player is always choosing a best response given that he knows his initial signal and the action that he is going to play. We call such distributions "Bayes correlated equilibria," since they correspond to one (very permissive) version of incomplete information correlated equilibrium. This set clearly reduces to the set of outcomes attainable by omniscient persuasion in the one player case. We have characterized the set of Bayes correlated equilibria in the context of price discrimination and auctions (Bergemann, Brooks, and Morris (2015a), (2015b)) and symmetric linear best response games with normal signals (Bergemann and Morris (2013), Bergemann, Heumann, and Morris (2015) and others). In these applications, we consider both the case where players have initial information (as in omniscient persuasion) and the case where players have no initial information (as in Bayesian persuasion with an uninformed receiver).

The second step of an information design procedure is to identify the information structure that would be chosen by an interested designer. This step has not been the focus of our earlier analysis. However, our results can and have been used to make "information design" observations. One novel question that arises in the many player information design is whether the information designer wants to convey information to the players in the form of private or public signals. If the information designer would like the players' actions to be correlated, then he will choose to give them public signals, whereas if the designer wants the players' actions to be uncorrelated, then he will choose private signals. If he wants the actions to be as uncorrelated as possible, then (conditional on the amount of payoff relevant information conveyed to the players), he would like signals to be as uncorrelated as possible. We have studied the case where an information designer acts in the joint interest of the players. In the appendix of Bergemann and Morris (2016a), we analyze a two player, two action, two state example that is the generalization of the one player example above. In this case, public signals are best for the information designer under strategic complementarities (e.g., with the usual payoffs for a many player bank run) while private signals are best under strategic

substitutes. Our analysis of information sharing in oligopoly in Bergemann and Morris (2013) followed the same logic. Under oligopoly (where actions are strategic substitutes), firms would like to have accurate information about demand. But they would also like their actions to be as uncorrelated as possible (since each firm would produce more when other firms are producing less). There is a conflict between these two objectives. We show that the optimal information structure trades off these two competing objectives by having firms observe conditionally independent noisy signals of demand. In the elegant dynamic two player analysis of Ely (2015), the information designer wants to minimize the probability of both players running, and thus wants uncorrelated actions, and thus chooses signals which are as uncorrelated as possible.

A striking feature of the one player example is that the set of attainable outcomes shrinks as the receiver becomes more informed. This monotonicity arises because (i) the sender has the option of giving the receiver any information that he wants, and thus less information does not limit the feasible outcomes for the receiver, while (ii) more information implies tighter obedience constraints for the sender. In a general omniscient persuasion analysis, there is a tight generalization of this observation: fixing a state space, the set of attainable outcomes shrinks in all decision problems if and only if there is more information, in the sense of the Blackwell's sufficiency order. This result also extends to the many player case: more information reduces the set of Bayes correlated equilibria in all games, and vice versa. The subtlety in this statement is formalizing what is meant by "more information" in the many player case. In Bergemann and Morris (2016a), we identify the relevant many player generalization of Blackwell's order under which this result is tight.

We can make a tight connection between the three versions of persuasion discussed in this note and the incomplete information correlated equilibrium literature. As discussed in the introduction, persuasion problems have an informed information designer choosing information for a single player while the communication in games literature has an information designer with no informational advantage choosing information for many players. The latter scenario corresponds to the classic incomplete information correlated equilibrium literature. Omniscient persuasion corresponds to the *Bayesian solution* of Forges (1993): the mediator is able to condition his recommendation on any information possessed by the players jointly; private persuasion corresponds to the *communication equilibria* of Forges (1993), where the mediator is able only to condition on information that players can be induced to report truthfully; and public persuasion corresponds to the *strategic form correlated equilibrium* of Forges (1993), where a mediator cannot condition on players' information at all. Our work on *Bayes correlated equilibria* generalizes omniscient persuasion by allowing many players and generalizes the Bayesian solution by allowing the information designer to have information that the players do not collectively possess.

References

- AUMANN, R., AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*. MIT Press.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2015): “Selling Experiments: Menu Pricing of Information,” .
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015a): “First Price Auctions with General Information Structures: Implications for Bidding and Revenue,” Discussion paper.
- (2015b): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–957.
- BERGEMANN, D., T. HEUMANN, AND S. MORRIS (2015): “Information and Volatility,” *Journal of Economic Theory*, 158, 427–465.
- BERGEMANN, D., AND S. MORRIS (2013): “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81, 1251–1308.
- (2016a): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, forthcoming.
- (2016b): “Information Design and Incomplete Information Correlated Equilibrium,” Discussion paper.
- ELY, J. (2015): “Beeps,” Discussion paper, Northwestern University.
- FORGES, F. (1993): “Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information,” *Theory and Decision*, 35, 277–310.
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KOLOTLIN, A., M. LI, T. MYLOVANOV, AND A. ZAPECHELNYUK (2015): “Persuasion of Privately Informed Receiver,” Discussion paper.
- MYERSON, R. (1991): *Game theory*. Harvard University Press, Cambridge.
- TANEVA, I. (2015): “Information Design,” Discussion paper, University of Edinburgh.