

Information Design: A Unified Perspective[†]

DIRK BERGEMANN AND STEPHEN MORRIS*

Given a game with uncertain payoffs, information design analyzes the extent to which the provision of information alone can influence the behavior of the players. Information design has a literal interpretation, under which there is a real information designer who can commit to the choice of the best information structure (from her perspective) for a set of participants in a game. We emphasize a metaphorical interpretation, under which the information design problem is used by the analyst to characterize play in the game under many different information structures. We provide an introduction to the basic issues and insights of a rapidly growing literature in information design. We show how the literal and metaphorical interpretations of information design unify a large body of existing work, including that on communication in games (Myerson 1991), Bayesian persuasion (Kamenica and Gentzkow 2011), and some of our own recent work on robust predictions in games of incomplete information. (JEL C70, D82, D83)

1. Introduction

Players' payoffs in a game depend on their actions and also on the realization of a payoff-relevant state. An "information designer" can commit how to provide information about the states to the players. "Information design"

studies how the information designer, through the choice of the information provided, can influence the individually optimal behavior of the players to achieve her objective. She can achieve this objective even though she has no ability to change outcomes or force the

*Bergemann: Department of Economics, Yale University. Morris: Department of Economics, Princeton University. We acknowledge financial support from NSF Grants SES 0851200 and 1459899. We are grateful for productive suggestions by the Editor, Steven Durlauf, and five anonymous referees. This material has been presented in lectures at the 2015 Istanbul meetings of the Society of Economic Design, the 2015 Delhi School of Economics Winter School, the 2016 AEA meetings, the North American Summer Meeting of the Econometric Society, the Gerzensee Summer Symposium in Economic Theory, the Becker Friedman Institute conference on Frontiers and Economic Theory and Computer Science, the Harris

lecture at Harvard University, and the 2016 fall Finance Theory group conference at Princeton. Some of the material in this paper was previewed in Bergemann and Morris (2016b). We have benefited from discussions and correspondence about the material in this paper with Ben Brooks, Laura Doval, Jeff Ely, Françoise Forges, Drew Fudenberg, Olivier Gossner, Tibor Heumann, Atsushi Kajii, Emir Kamenica, Rohit Lamba, Laurent Mathevet, Roger Myerson, Tymofiy Mylanov, Alessandro Pavan, Eran Shmaya, Jonathan Weinstein, and Juan Pablo Xandri; and from valuable research assistance from Ian Ball, Leon Músolff, Denis Shishkin, Áron Tóbiás and Xinyang Wang.

[†] Go to <https://doi.org/10.1257/jel.20181489> to visit the article page and view author disclosure statement(s).

players to choose particular actions that determine outcomes.¹

The past decade has seen a rapidly growing body of literature in information design. An influential paper by Kamenica and Gentzkow (2011) phrased the optimal design of information as a “Bayesian persuasion” problem between a sender and single receiver. A large body of work fits this rubric, including important contributions of Brocas and Carrillo (2007) and Rayo and Segal (2010). The economic applications of information design have been investigated in areas as far apart as grade disclosure and matching markets (Ostrovsky and Schwarz 2010), voter mobilization (Alonso and Câmara 2016), traffic routing (Das, Kamenica, and Mirka 2017), rating systems (Duffie, Dworzak, and Zhu 2017), and transparency regulation (Asquith, Covert, and Pathak 2013) in financial markets, price discrimination (Bergemann, Brooks, and Morris 2015), and stress tests in banking regulation (Inostroza and Pavan 2017).

One purpose of the paper is to provide an overview of information design that unifies this recent work with a number of literatures sometimes treated as distinct. If we assume that 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, section 6.3) and, more generally, the literature on correlated equilibrium in incomplete information games (Forges 1993). 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 of Kamenica and Gentzkow (2011). Information design concerns the general case where there are *both* many players *and*

the information designer has an informational advantage over the players. This case has been the focus of some of our own work (Bergemann and Morris 2013b, 2016a), where we show that the set of outcomes that can arise in this setting corresponds to a version of incomplete information equilibrium (Bayes-correlated equilibrium, or BCE) that allows outcomes to be conditioned on states that the players do not know.

A second purpose of the paper is to highlight a distinction between *literal information design* and *metaphorical information design*. The information design problem has a literal interpretation (given above): there really is an information designer (or mediator, or sender) who can commit to provide extra information to players to serve her own interests. While the commitment assumption may be problematic in many settings, it provides a useful benchmark. But the information design formulation might also be a metaphor that the analyst uses as a tool. For example, we might be interested in finding an upper bound (across information structures) on the aggregate variance of output in a given economy with idiosyncratic and common shocks to agents’ productivity (Bergemann, Heumann, and Morris 2015). We can understand this as an information design problem, where the information designer is interested in choosing an information structure to maximize aggregate variance in output. But in this case, we do not have in mind that there is an actual information designer maximizing aggregate variance. We will discuss this application, and other applications where information design is metaphorical, below.

This survey reviews the *pure information design* problem where a designer can commit to a certain information structure for the players but has no control over outcomes. This problem is a special case of the more *general mechanism design* problem where a mechanism designer can control outcomes but may also be able to manipulate information

¹We follow Taneva (2015) in our use of the term “information design” in this context.

in the course of doing so.² We study the case where all information structures are available to the designer. It is thus possible to appeal to the revelation principle from the general mechanism design problem, and without loss of generality restrict attention to information structures where the signals that the information designer sends to a player can be identified with action recommendations. This revelation principle/mechanism design approach to information design thus contrasts with work where there is no commitment to the information structure or attention is restricted to a parameterized class of information structures.

We use a family of two-player, two-action, and two-state examples to survey the literature, and to provide some graphical illustrations. We start with the leading example of Bayesian persuasion (with a single player/receiver with no prior information) from the work of Kamenica and Gentzkow (2011). We can use extensions of this example—with many players and prior information—to illustrate many of the key ideas in the survey. Three key substantive general insights are illustrated in these examples.

First, it is often optimal for the information designer to selectively obfuscate information. This insight is familiar from the case of one player without prior information.

Second, the information designer has less ability to manipulate outcomes in his favor if players have more prior information: if the players are endowed with their own information, the designer has less influence over the information structure that they end up with. This insight can already be illustrated in the

one-player case. But we will also describe a general partial order on information structures—generalizing the Blackwell order for the one-player case—which characterizes the right definition of “more information” in this context (Bergemann and Morris 2016a).

Third, we can ask whether the information designer prefers to give the information to players in a public or in a private message. Of course, this last question only arises once we have multiple players. Public information is optimal if the information designer wants perfect correlation between players’ actions; otherwise private information will be optimal. While the information designer may have intrinsic preferences over whether players’ actions are correlated (or not), the designer may care about correlation for purely instrumental reasons: if there are strategic complementarities between the players’ actions, she may want to correlate players’ actions to relax the obedience constraints on her ability to attain specific outcomes. The converse holds for strategic substitutability. We will illustrate the case when there are only instrumental preferences over correlation.

The examples also illustrate a methodological point. The information design problem can be solved in two steps. First, we can identify the set of outcomes that could be induced by the information designer. Second, we can identify which of these outcomes would be preferred by the information designer. This, too, parallels the mechanism design literature: we can first identify which outcomes are implementable and then identify the one most preferred by the designer. As noted above, in the information design problem, the set of implementable outcomes corresponds to the set of Bayes-correlated equilibria. This approach reduces the problem to a linear program.

The information designer is assumed to be able to commit to an information structure that maps payoff-relevant states of the world and prior information of the agents

²Myerson (1982, 1991) describes the problem of “communication in games” where the designer cannot control outcomes but can elicit information from players and pass it to other players. Thus, what we are defining as the “information design” problem can be viewed as Myerson’s “communication in games” problem with the important new feature that the designer may have access to information that is not available to the players.

(types) into possibly stochastic signals to the players.³ As the mapping essentially recommends actions to the players, we refer to it as a *decision rule*. We initially focus on what we sometimes call the *omniscient* case: here, the information designer faces no constraints on her ability to condition the signals on the payoff-relevant states of the world *and* all the players' prior information (i.e., their types). But we also consider information design constrained by private information, where the prior information of the players is not accessible to the information designer, even though she can condition on the payoff-relevant states. There are two cases to consider here: an information designer may be able to condition on the reported realizations of the players' signals even if she does not observe them (information design with elicitation) or she may be unable to do so (information design without elicitation). If the information designer cannot condition on the payoff state at all, and has to rely entirely on the private information of the players, then these three scenarios (omniscient, private information with elicitation, and private information without elicitation) correspond to versions of incomplete information correlated equilibrium: in the terminology of Forges (1993), the Bayesian solution, communication equilibrium, and strategic form correlated equilibrium, respectively.

Once the information designer has picked the information structure, the players decide how to play the resulting game of incomplete

information. There may be multiple Bayes–Nash equilibria of the resulting game. In our treatment of the information design problem, we have been implicitly assuming that the designer can pick which equilibrium is played. Under this maintained assumption, we can appeal to the revelation principle and focus attention on information structures where the signal space is set equal to the action space, and the signals have the interpretation that they are action recommendations. In the single-player case, this maintained equilibrium selection assumption is without loss of generality. But just as the revelation principle breaks down in mechanism design if the designer does not get to pick the best equilibrium (as in Maskin 1999), it similarly breaks down for information design.⁴ We follow Mathevet, Perego, and Taneva (2017) in formally describing a notion of maxmin information design, where an information designer gets to pick an information structure but the selected equilibrium is the worst one for the designer. We note how some existing work can be seen as an application of maxmin information design, in particular, an extensive literature on “robustness to incomplete information” (Kajii and Morris 1997).

Other assumptions underlying the revelation principle—and maintained throughout this paper—are that (i) all information structures are feasible, (ii) there is zero (marginal) cost of using information, (iii) there is a single information designer, and (iv) the setting is static. Of course, there are many (static) settings where the impact of different information structures has been studied, without allowing all information structures. Two classic examples would be information sharing in oligopoly (a literature beginning with Novshek and Sonnenschein 1982) and the revenue comparison across different auction formats

³Whether or not the information designer will observe the payoff relevant state is irrelevant—what is important is whether she can condition the signals she sends on the realization of the state and the players' prior signals. For example, a prosecutor might never know whether the defendant is guilty or innocent, but can nevertheless set up an investigation process that would provide different evidence depending on the actual guilt or innocence of the subject and the information of the judge. We thank an anonymous referee who stressed the distinction between conditioning on a state, and actually knowing the realization of the state.

⁴This point has been highlighted by Carroll (2016) and Mathevet, Perego, and Taneva (2017).

in auction theory (Milgrom and Weber 1982). In the former, there is a restriction to normally distributed signals, and in the latter there is the restriction to affiliated signals.

Optimal information design in dynamic settings has been studied recently in Ely, Frankel, and Kamenica (2015); Passadore and Xandri (2014); Doval and Ely (2016); Ely (2017); Ely and Szydlowski (2017); Ball (2018); and Makris and Renou (2018). A new aspect to information design that appears in these dynamic settings is that information can be used as an incentive device to reward behavior over time. Horner and Skrzypacz (2016) survey work on information design in dynamic settings. Gentzkow and Kamenica (2014) consider the case of costly information and Gentzkow and Kamenica (2017) allow for multiple information designers.

This paper provides a conceptual synthesized guide to the literature; we discuss applications when they are relevant for this purpose, but make no attempt to provide a comprehensive survey of the many applications of information design. A recent survey by Kamenica (forthcoming) reviews the literature of Bayesian persuasion and discusses some leading economic applications.

We describe the basic information design problem in section 2. We illustrate the main notions and ideas with an investment example in section 3. We discuss key ideas from information design in the investment example in section 4. Here, we discuss private versus public signals, intrinsic versus instrumental preferences over correlation, the two-step procedure for solving information design problems, ordering information, and the use of concavification in information design (instead of pure linear programming methods). Section 5 describes, in more detail, two applications of information design with a microeconomic and macroeconomic perspective, respectively: limits of price discrimination, and the link between information and volatility. These

two applications emphasize the relevance of the metaphorical interpretation of information design. In section 6, we describe what happens when players' prior information is not known by the information designer; this discussion allows us to locate the information design problem within mechanism design and within a larger literature on incomplete information correlated equilibrium reviewed by Forges (1993). In section 7, we discuss the role of equilibrium selection.

Given the synthetic treatment of the literature, there is much terminology that has been introduced and used in different contexts (including by us in prior work), and which is at times inconsistent or redundant. To give one example, what we are calling an "information designer" has in previous work been called a sender, mediator, principal, or mechanism designer. We are attempting throughout to use a unified and consistent language, but compromising at times between the use of a consistent terminology and precedents set by earlier work.

2. *The Information Design Problem*

We begin by describing the general setting and notation that we maintain throughout the paper. We will fix a finite set of players and a finite set of payoff states of the world. There are I players, $1, 2, \dots, I$, and we write i for a typical player. We write Θ for the set of payoff states of the world and θ for a typical element of set Θ .

A "basic game" G consists of (1) for each player i , a finite set of actions A_i , and an (ex post) utility function

$$u_i : A \times \Theta \rightarrow \mathbb{R},$$

where we write $A = A_1 \times \dots \times A_I$, and $a = (a_1, \dots, a_I)$ for a typical element of A ; and (2) a full support prior $\psi \in \Delta(\Theta)$ that is shared by all players and the information designer. Thus $G = ((A_i, u_i)_{i=1}^I, \Theta, \psi)$. We

define the ex post objective of the information designer by:

$$v : A \times \Theta \rightarrow \mathbb{R}.$$

An “information structure” S consists of (i) for each player i , a finite set of types $t_i \in T_i$; and (ii) a type distribution $\pi : \Theta \rightarrow \Delta(T)$, where we write $T = T_1 \times \cdots \times T_I$. Thus $S = ((T_i)_{i=1}^I, \pi)$.

Together, the “payoff environment” or “basic game” G and the “belief environment” or “information structure” S define a standard “incomplete information game” (G, S) . While we use different notation, this division of an incomplete information game into the “basic game” and the “information structure” is a common one in the literature, see, for example, Gossner (2000).

We are interested in the problem of an information designer who has the ability to commit to provide the players with *additional* information, in order to induce them to make particular action choices. In this section, we will consider the leading case where the designer can condition on the state *and* on all the players’ types—their prior information—if they have any. We will sometimes refer to this setting as *omniscient* information design. In section 6, we will consider the case where prior information of the players is truly private to them, and hence the information designer cannot condition on their prior information unless she is able to induce them to reveal it.⁵

If viewed as an extensive form game between the information designer and the players, the timing is as follows:

- (i) the information designer picks and commits to a rule for providing the players with extra messages;

- (ii) the true state θ is realized;
- (iii) each agent’s type t_i is privately realized;
- (iv) the players receive extra messages according to the information designer’s rule;
- (v) the players pick their actions based on their prior information and the messages provided by the information designer;
- (vi) payoffs are realized.

We emphasize that the information designer commits to a decision rule before the realization of the state and type profile. This structure in the timing sets the information design problem apart from informed principal, cheap talk, or signaling environments where the informed party chooses a message only after the state has been revealed.

In general, the information designer could follow any rule for generating messages. However, a “revelation principle” argument implies that it is without loss of generality to assume that the information designer sends only “action recommendations” that are obeyed. The argument is that any message will give rise to an action in equilibrium and we might as well label messages by the actions to which they give rise. We discuss the revelation principle in more detail below. Given this restriction, the information designer is choosing among *decision rules*

$$(1) \quad \sigma : T \times \Theta \rightarrow \Delta(A).$$

The information designer can condition the recommended action profile on the true state $\theta \in \Theta$ and the type vector $t \in T$. We stress that the designer does not need to know the true realization of the state or the type profile—it is sufficient that the decision

⁵For this reason, we refer to the types of a player as his prior information, and only in section 6 does this information become truly private.

rule can condition on these. For example, in a medical test, the information designer, the doctor, may not know the true condition of the patient, but can choose a diagnostic test that reveals the condition of the patient with the desired precision and accuracy.

The decision rule encodes the information that the players receive about the realized state of the world, the types and actions of the other players. The conditional dependence of the recommended action a on state of the world θ and type profile t represents the information conveyed to the players.

The key restriction on the decision rule is a notion of obedience that we now define. Obedience is the requirement that the information privately communicated to player i in the form of an action recommendation a_i according to σ is such that each player i would want to follow his recommended action a_i rather than choose any other available action a'_i .

DEFINITION 1 (Obedience): *Decision rule $\sigma : T \times \Theta \rightarrow \Delta(A)$ is obedient for (G, S) if, for each i , $t_i \in T_i$ and $a_i \in A_i$, we have*

(2)

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

for all $a'_i \in A_i$.

For notational compactness, we only list the elements, but not the sets, over which we sum under the summation operator, thus, e.g., \sum_a rather than $\sum_{a \in A}$. The obedience inequality requires that each player i , after receiving his recommendation a_i , finds that no other action a'_i could yield him a strictly higher utility. We emphasize that each player, when computing his expected utility, is indeed using the information contained in the action recommendation a_i , and

thus the above inequality is written from an interim perspective (conditioning on t_i and a_i). We can state the above inequality explicitly in terms of the interim beliefs that agent i holds, given his type t_i and his action recommendation a_i , thus using Bayes's rule:

$$\begin{aligned} & \sum_{a_i, t_i, \theta} u_i((a_i, a_{-i}), \theta) \\ & \times \frac{\sigma((a_i, a_{-i})|(t_i, t_{-i}), \theta) \pi((t_i, t_{-i})|\theta) \psi(\theta)}{\sum_{a''_i, t''_i, \theta''} \sigma((a_i, a''_i)|(t_i, t''_i), \theta'') \pi((t_i, t''_i)|\theta'') \psi(\theta'')} \\ & \geq \sum_{a'_i, t_i, \theta} u_i((a'_i, a_{-i}), \theta) \\ & \times \frac{\sigma((a_i, a_{-i})|(t_i, t_{-i}), \theta) \pi((t_i, t_{-i})|\theta) \psi(\theta)}{\sum_{a''_i, t''_i, \theta''} \sigma((a_i, a''_i)|(t_i, t''_i), \theta'') \pi((t_i, t''_i)|\theta'') \psi(\theta'')}. \end{aligned}$$

We observe that the belief of agent i updates independently of whether he is following the recommendation a_i or deviating from it to a'_i . Moreover, the denominator in Bayes's rule sums up over all possible profiles $a''_i \in A_{-i}$, $t''_i \in T_{-i}$, $\theta'' \in \Theta$, and hence is constant across all possible realizations of $a_{-i} \in A_{-i}$, $t_{-i} \in T_{-i}$, $\theta \in \Theta$. Hence, we can multiply through to obtain the earlier inequality (2), provided that the denominator is strictly positive.

Bergemann and Morris (2016a) define a *Bayes-correlated equilibrium* (BCE) to be any decision rule σ satisfying obedience. An important aspect of this solution concept is that the decision rule σ enters the obedience constraints, as stated above in (2), in a linear manner as a probability. Thus, an obedient decision rule can be computed as the solution to a linear program.

PROPOSITION 1 (Revelation Principle): *An omniscient information designer can attain decision rule σ if and only if it is a BCE, i.e., if it satisfies obedience.*

By “can attain decision rule” we mean that there exists a (perhaps indirect)

communication rule that gives rise to this decision rule in Bayes–Nash equilibrium.⁶ We refer to the resulting (ex ante) joint distribution of payoff state θ and action profile a as the *outcome* of the information design, thus integrating out the prior information t :

$$\sum_{t \in T} \sigma(a|t, \theta) \pi(t|\theta) \psi(\theta).$$

In this (and later) propositions, we omit formal statements and proofs that correspond to revelation principle arguments. Bergemann and Morris (2016a) give a formal statement⁷ and proof of this proposition as theorem 1. The argument is a straightforward adaptation of standard characterizations of complete and incomplete information correlated equilibrium.

If there is no payoff uncertainty—the set Θ is a singleton—then the notion of BCE exactly coincides with the complete information correlated equilibrium as introduced in the seminal contributions of Aumann (1974, 1987). In the absence of payoff uncertainty, we can simply suppress the dependence of the payoff function on the state of the world θ . Thus, the decision rule σ does not vary with θ , nor is there any private information t about the state of the world θ . A decision rule σ is then simply a joint distribution over actions, or $\sigma \in \Delta(A)$. Now, a distribution $\sigma \in \Delta(A)$ is defined to be a correlated equilibrium if for each i and $a_i \in A_i$, we have:

$$\begin{aligned} (3) \quad & \sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \sigma(a_i, a_{-i}) \\ & \geq \sum_{a'_{-i} \in A_{-i}} u_i(a'_i, a_{-i}) \sigma(a_i, a_{-i}), \quad \forall a'_i \in A_i. \end{aligned}$$

⁶We do not discuss information design under solution concepts other than Bayes–Nash equilibrium in this paper. Mathevet, Peregó, and Taneva (2017) study information design under bounded level rationalizability and Inostroza and Pavan (2017) under full rationalizability.

⁷A formal statement also appears in section 7.

The obedience condition (2) thus collapses to the best response property (3) in the absence of payoff uncertainty. Aumann (1987) argued that correlated equilibrium captured the implications of common knowledge of rationality in a complete information game (under the common prior assumption). An alternative interpretation is that the set of correlated equilibria is the set of outcomes attainable by an information designer in the absence of payoff uncertainty. We discuss these interpretational issues and the literature on incomplete information correlated equilibrium more broadly in section 6.4.

The proof of proposition 1 in Bergemann and Morris (2016a), like the proof of Aumann (1987), is a “revelation principle” argument, establishing that it is without loss of generality to focus on a set of signals that equals the set of actions to be taken by the agents—so that there is “direct communication”—and to recommend actions in such way that they will be obeyed—so that “incentive compatibility” gives rise to “obedience” conditions. In the case of complete information, Myerson (1991, section 6.2) describes this as the “revelation principle for strategic form games.” Note that while the expression “revelation principle” is sometimes limited to the case where agents are sending messages rather than receiving them (e.g., Fudenberg and Tirole 1991 and Mas-Collel, Whinston, and Green 1995), we follow Myerson in using the broader meaning throughout the paper. In the basic information design described in this section, the only incentive constraints are obedience conditions, but we discuss the extension to the case where the information designer must elicit players’ private information in section 6, where truth-telling incentive constraints also arise. We postpone a discussion of how to place information design in the broader context of the mechanism design literature until then.

Proposition 1 characterizes the set of outcomes that an information designer could

attain, i.e., a feasible set. To complete the description of the basic information design problem, we need an objective. Given the information designer's ex post utility $v(a, \theta)$, her ex ante utility from the decision rule σ for a given game of incomplete information (G, S) is:

$$(4) V(\sigma) = \sum_{a, \theta} v(a, \theta) \sigma(a | \theta) \pi(\theta) \psi(\theta).$$

The (omniscient) information design problem is then to pick a BCE σ to maximize $V(\sigma)$. When there is a single player with no prior information, the information design problem reduces to the benchmark Bayesian persuasion problem described by Kamenica and Gentzkow (2011). In this case, the single player is called the “receiver” and the information designer is called the “sender.”

3. An Investment Example

We now apply this framework to an investment game and discuss the main themes of information design through the lens of this example.

We first consider the following benchmark setting. There is a bad state (B) and a good state (G). The two states are equally likely:

$$\psi(G) = \psi(B) = \frac{1}{2}.$$

There is one player (the “firm”). The firm can decide to *invest* or *not invest*. The payoff from not investing is normalized to 0. The payoff to investing is -1 in the bad state and x in the good state, with $0 < x < 1$. These payoffs, $u(a, \theta)$, are summarized in the following matrix:

(5)

$u(a, \theta)$	bad state B	good state G
invest	-1	x
not invest	0	0

3.1 Single Player without Prior Information

We begin the analysis when the firm has no prior information about the state (beyond the uniform prior). Together with the above assumptions about the payoff matrix, the firm would therefore choose to *not invest* if it had no additional information.

We will assume that an information designer (the “government”) is interested in maximizing the probability of investment independent of the state, or

$$1 = v(\text{invest}, \theta) > v(\text{not invest}, \theta) = 0,$$

$$\theta = B, G.$$

This example is (modulo some changes in labelling) the leading example in Kamenica and Gentzkow (2011). We will describe this example first, but then use variations to illustrate more general points. The decision rule is now simply:

$$\sigma : \Theta \rightarrow \Delta(A),$$

where we can omit the type space T due to the absence of prior information. In this binary decision environment, the decision rule $\sigma(\theta)$ specifies the probability of investment, denoted by p_θ , conditional on the true state $\theta \in \{B, G\}$. Thus, a decision rule is a pair (p_B, p_G) of investment probabilities. We can think of a decision rule as a (stochastic) action recommendation from the government. If the recommendations are obeyed, the *outcome*—the ex ante distribution over states and actions—is given by:

$\sigma(a \theta)\psi(\theta)$	bad state B	good state G
invest	$\frac{1}{2}p_B$	$\frac{1}{2}p_G$
not invest	$\frac{1}{2}(1 - p_B)$	$\frac{1}{2}(1 - p_G)$

If the firm receives a recommendation to invest, it will update its beliefs about the state by Bayes's rule. The firm's interim expected utility from following the recommendation represents the left-hand side of the inequality below. If the firm were to disobey the recommendation and chose not to invest, then its payoff would be zero. This gives rise to the following obedience constraint:

$$(6) \quad \frac{\frac{1}{2}p_B}{\frac{1}{2}p_B + \frac{1}{2}p_G}(-1) + \frac{\frac{1}{2}p_G}{\frac{1}{2}p_B + \frac{1}{2}p_G}x \geq 0.$$

As we discussed following definition 1, we can simplify the inequality by multiplying through with the conditioning probability $\frac{1}{2}p_B + \frac{1}{2}p_G$, and thus write the obedience condition equivalently in terms of the interim probabilities:

$$(7) \quad \frac{1}{2}p_B(-1) + \frac{1}{2}p_Gx \geq 0 \Leftrightarrow p_G \geq \frac{p_B}{x}.$$

There is an analogous obedience constraint corresponding to the recommendation not to invest, namely:

$$0 \geq \frac{1}{2}(1 - p_B)(-1) + \frac{1}{2}(1 - p_G)x.$$

Because the firm would not invest absent information from the designer—by our maintained assumption that $x < 1$ —the binding obedience constraint will be the one corresponding to investment, i.e., inequality (7). We see that the highest probability of investment corresponds to the decision rule with $p_G = 1$ and $p_B = x$.

We illustrate the set of BCE decision rules for the case where $x = 55/100$ in figure 1. Any decision rule (p_B, p_G) in the blue shaded area can arise as some BCE. We observe that the feasible set of BCE does not depend on the government's preference $v(a, \theta)$.

Now every BCE decision rule corresponds to optimal behavior under some information structure S . By the *revelation principle* for

the BCE, it suffices to give the firm a *binary* information structure S to implement any BCE decision rule in the *binary* action environment. For the outcome that maximizes the probability of investment, it suffices to generate a no-investment recommendation with probability $1 - x$ if the state is bad, and otherwise give the firm an investment recommendation. The resulting *outcome*—the ex ante distribution over states and actions—is given by:

$$(8) \quad \begin{array}{|c|c|c|} \hline \sigma(a|\theta)\psi(\theta) & \text{bad state } B & \text{good state } G \\ \hline \text{invest} & \frac{1}{2}x & \frac{1}{2} \\ \hline \text{not invest} & \frac{1}{2}(1-x) & 0 \\ \hline \end{array}.$$

Thus, a government trying to encourage investment will *obfuscate* the states of the world in order to maximize investment. By pooling realizations of the bad and good states in the recommendation to invest, the firm is made exactly indifferent between investing or not when recommended to invest. The bad state is completely isolated in the recommendation not to invest. Finally, we observe that under complete information the firm would always invest in the good state and never invest in the bad state. We thus have described three different information structures—zero information, partial information, and complete information—that support the three vertices of the above investment triangle. Thus, the set of all investment probabilities that satisfy the obedience constraints can be described by a set of linear inequalities that jointly form a polyhedron of implementable outcomes.

3.2 Single Player with Prior Information

We remain with the investment example where there is still only one firm, but now the firm has some prior information about the true state that it receives independently




Figure 1. Investment Probability with Uninformed Player: $x = 55/100$

of the government.⁸ In particular, the firm has a type (or receives a signal) that is “correct” with probability $q > 1/2$. Formally, the firm observes its type $t \in \{b, g\}$ with probability q conditional on the true state being B or G , respectively:

$\pi(t \theta)$	bad state B	good state G
bad signal b	q	$1 - q$
good signal g	$1 - q$	q

Here, signals refer to the prior information that firms are endowed with. Conditional on the type of the firm, the analysis of the obedience constraints reduces immediately to the analysis of the previous section, but where the firm has an updated belief, q or $1 - q$, depending on the type. We nonetheless analyze this problem because we want to trace the ex ante implications of a player’s prior information for information design. A decision rule σ now specifies the probability of investment $p_{\theta t}$ conditional on the true state $\theta \in \{B, G\}$ and the type $t \in \{b, g\}$. Thus a decision rule is now a quadruple:

$$(9) \quad p = (p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg}).$$

⁸Some detailed calculations for this example appear in the appendix.

We can solve the problem—conditional on state *and* type—as before. For example, the obedience constraint for the recommendation to invest after receiving a good type g now becomes:

$$(10) \quad (1 - q)p_{Bg}(-1) + qp_{Cg}x \geq 0.$$

However, we are interested in what we can say about the joint distribution of states and actions *ex ante*, integrating out the types, say

$$p_G = qp_{Cg} + (1 - q)p_{Cb}.$$

One can show that there is a lower bound on investment in the good state given by:

$$(11) \quad p_G \geq q - \frac{1 - q}{x},$$

which approaches 1 as q approaches 1. As before, the bound for p_G depends on p_B and the lowest bound is obtained by taking $p_B = 0$ in that expression.

The set of BCE is illustrated in figure 2. More prior information shrinks the set of BCE, since the obedience constraints become tighter. Once q reaches 1, the firm knows the state and the information designer has no ability to influence the outcome. The firm is simply pursuing the complete information optimal decision, which is to invest in the good state, $p_G = 1$, and not to invest in the bad state, $p_B = 0$.

The set of BCE across different q has two notable features. First, we notice that the set of BCE happens to be constant across some information structures near precision $q = 0.5$, for example at $q = 0.5$ and $q = 0.6$. With low precision in the signals, such as $q = 0.6$, the firm would pursue the same action for either type realization, absent any additional recommendation by the government. Thus, the weak prior information by the agent does not constrain the government in its recommendation policy. In

consequence, prior information of the player only affects the set of BCE if the prior information by itself already generates a differential response by the player. Second, the slope of the boundary is constant across different levels of precision q . This occurs because the rate at which the optimal decision of the player can be reversed by additional information of the designer (and hence indifference is attained) is constant across q by Bayes's law.

3.3 Many Players without Prior Information

We can now generalize the analysis to two firms and return to the assumption that the firms have no prior information.⁹ We assume, for now, that the government wants to maximize the sum over each individual firm's probability of investment. If there is no strategic interaction between firms, the previous analysis can be carried out firm by firm and will thus be unchanged.

But we now perturb the problem to make it strategic, assuming that each firm gets an extra payoff ε if both invest, where ε may be positive or negative. If ε is positive, we have a game of strategic complementarities; if $\varepsilon < 0$, we have a game of strategic substitutes. We can write firm 1's state dependent payoffs for the game as follows (and symmetrically for firm 2):

		firm 2	
(12)	firm 1		

⁹Other two player, two action, and two state examples appear in Bergemann and Morris (2013a, 2016a) and Taneva (2015).




Figure 2. Investment Probability with Informed Player: $x = 55/100$

		firm 2	
		$\theta = G$	
		invest	not invest
firm 1	invest	$x + \varepsilon$	x
	not invest	0	0

We can focus on symmetric decision rules, given the symmetry of the basic game, for any symmetric objective of the information designer. To see why, note that if we found an asymmetric maximizing decision rule, the decision rule changing the names of the firms would also be optimal, and so would the (symmetric) average of the two decision rules. Therefore, we will continue to write p_θ for the probability that each firm will

invest in state $\theta \in \{G, B\}$; but we will now write r_θ for the probability that both invest. Thus a decision rule is a vector (p_B, r_B, p_G, r_G) . A decision rule can now be represented in a table as

(13)

$\theta = B$	invest	not invest
invest	r_B	$p_B - r_B$
not invest	$p_B - r_B$	$1 + r_B - 2p_B$

$\theta = G$	invest	not invest
invest	r_G	$p_G - r_G$
not invest	$p_G - r_G$	$1 + r_G - 2p_G$

To ensure that all probabilities are nonnegative, we require that for all $\theta \in \{B, G\}$:

$$\max\{0, 2p_\theta - 1\} \leq r_\theta \leq p_\theta.$$

The firm has an incentive to invest when told to do so if

$$(14) \quad -\frac{1}{2}p_B + \frac{1}{2}p_G x + \frac{1}{2}(r_B + r_G)\varepsilon \geq 0,$$

and an incentive to not invest when told to not do so if

$$\begin{aligned} &-\frac{1}{2}(1 - p_B) + \frac{1}{2}(1 - p_G)x \\ &+ \frac{1}{2}(p_B - r_B + p_G - r_G)\varepsilon \leq 0. \end{aligned}$$

Since $x < 1$, (14) is always the binding constraint and—for $|\varepsilon|$ sufficiently close to 0—we can rewrite it by the same reasoning as in section 3.1 as

$$(15) \quad p_G \geq \frac{p_B}{x} - (r_B + r_G)\frac{\varepsilon}{x}.$$

Now maximizing the sum of the probabilities of each firm investing corresponds to maximizing p_B , (or $p_B + p_G$, but we will have $p_G = 1$ always) subject to (15). For fixed $x < 1$ and $|\varepsilon| \approx 0$, it is clearly optimal to have firms always invest when the state is good (so $p_G = 1$ and $r_G = 1$) and it is not possible to get both firms to always invest when the signal is bad.

If $\varepsilon > 0$, (15) implies that it is optimal to choose r_B as large as possible given p_B . Thus we will set $r_B = p_B$. Substituting these variables into expression (15), we have

$$1 \geq \frac{p_B}{x} - (p_B + 1)\frac{\varepsilon}{x},$$

and so it is optimal to set

$$p_B = r_B = \frac{x + \varepsilon}{1 - \varepsilon},$$

and we can summarize the optimal decision rule in the following table:

$\theta = B$	invest	not invest
invest	$\frac{x + \varepsilon}{1 - \varepsilon}$	0
not invest	0	$\frac{1 - x - 2\varepsilon}{1 - \varepsilon}$

$\theta = G$	invest	not invest
invest	1	0
not invest	0	0

This decision rule entails a public signal: there is common certainty among the firms that they always observe the same signal.

If $\varepsilon < 0$, it remains optimal to have both firms always invest when the state is good ($p_G = r_G = 1$). But now we want to minimize r_B given (p_B, p_G, r_G) . To reduce cases, let us assume that $x > 1/2$ and restrict attention to $|\varepsilon| \leq x - 1/2$. In this case, it will be optimal to set $r_B = 0$. Substituting these expressions into (15), we have

$$1 \geq \frac{p_B}{x} - \frac{\varepsilon}{x}.$$

Thus we will now have $p_B = x + \varepsilon$, and we can summarize the optimal decision rule in the following table:

$\theta = B$	invest	not invest
invest	0	$x + \varepsilon$
not invest	$x + \varepsilon$	$1 - 2x - 2\varepsilon$

$\theta = G$	invest	not invest
invest	1	0
not invest	0	0

Under this decision rule, firms told to invest know neither whether the state is good or bad, nor if the other firm is investing or not. Thus, signals are *private* to each firm. Given that, in the bad state, each firm will invest with (roughly) probability x and will not with (roughly) probability $1 - x$, the above information structure minimizes the unconditional correlation of the signals across firms (or equivalently minimizes the negative correlation conditional on the bad state.)

Strategic complementarities increase the private return from investing if the other player invests as well. Below, we display the set of investment probabilities that can be attained by the government while varying the size of the strategic effect ε . As the strategic effect ε increases, the boundaries of the investment probabilities attainable by the government shift outwards as illustrated in figure 3. As the strategic complementarity increases (or strategic substitutability decreases), the government can support a larger probability of investment in both states. The intermediate case of $\varepsilon = 0$ reduces to the case of a single player, and hence reduces to the area depicted earlier in figure 1.

3.4 Many Players with Prior Information

We analyzed the case of two players *and* prior information in Bergemann and Morris (2016a). Here, we illustrate this case without formally describing it. As in the single-player case, an increase in players' prior information limits the ability of the designer to influence the players' choices. Consequently, the impact of prior information on the set of attainable investment probabilities with many players is similar to the one-player case. In figure 4 we illustrate the set of attainable investment probabilities under increasing prior information with strategic complementarities. The strategic complementarities give rise to a kink in the set of attainable probabilities (p_B, p_C), unlike in the single-player case depicted earlier in figure 2.

4. Issues in Information Design Illustrated by the Examples

Let us draw out the significance of these examples. One basic point that has been extensively highlighted (e.g., by Kamenica and Gentzkow 2011 and the following Bayesian persuasion literature) is that when there is a conflict between the designer and the player(s), it will, in general, be optimal for the designer to obfuscate, that is, hide information from the player(s) in order to induce him to make choices that are in the designer's interests. And conditional on obfuscation being optimal, it may not be optimal to hide all information, but will, in general, be optimal to partially reveal information. This issue already arises in the case of one player with no prior information.

In this section, we draw out a number of additional insights about information design that emerged from the examples. First, we observe that information will be supplied to players publicly or privately depending on whether the designer would like to induce positive or negative correlation in players' actions; we also discuss designers' possible intrinsic or instrumental reasons for wanting positive or negative correlation. Second, we note that in the case of one player with prior information, more prior information constrains the ability of the designer to control outcomes; we discuss the many-player generalization of this observation. Third, we discuss the elegant "concavification" approach as an alternative to the linear programming representation used above to characterize and provide insights into the information designer's problem. We also discuss an extension of the concavification approach to the many-player case but note limitations of the concavification approach, both in the one-player case and (even more) in the many-player case.




Figure 3. Investment Probability with Negative or Positive Strategic Term ε

4.1 Public versus Private Signals and Instrumental versus Intrinsic Motivation for Preferences over Correlation

An information designer will often have preferences over whether players' actions are correlated with each other or not. The case of many players without prior information illustrates the point that if the designer wants players' actions to be correlated, it will be optimal to give them public signals, and if he wants players' actions to be uncorrelated he will give them private signals. However, there are different reasons why the designer might want to induce positive or negative correlation in actions.

In our analysis of the case of two players without prior information, we made the assumption that the information designer wanted to maximize the sum of the probabilities that each player invests. Thus, we assumed that the information designer did not care whether players' actions were correlated or not. Put differently, we assumed that the information designer had no *intrinsic* preferences over correlation. Yet, despite this assumption, we observed that the information designer wants—for *instrumental* reasons—to induce correlated behavior when players' actions are strategic complements, and to induce negative correlation when there were strategic substitutes among the players. This in turn gen-




Figure 4. Investment Probability with Two Players with Prior Information, with Strategic Term $\varepsilon = 3/10$

erated the insight that the designer would like to generate public signals when there are strategic complementarities and private signals when there are strategic substitutes. The reason for this instrumental objective is that under strategic complements, the designer can slacken obedience constraints by correlating play, with the opposite mechanism holding under strategic substitutes.

We now describe three environments where there will be only instrumental concerns about correlation. First, Mathevet, Perego, and Taneva (2017) consider an environment with one-sided strategic complementarities. The designer cares about the action of a first player who cares about

the action of a second player who has no strategic concerns, i.e., does not care about the first player's action. In this case, the information designer does not have intrinsic preferences over correlation (because she only cares about the first player's action), but has an instrumental incentive to correlate actions because she can use information design to influence the action of the second player and correlate behavior in order to slacken the first player's obedience constraint. In the formulation of Mathevet, Perego, and Taneva (2017), the information designer is a manager, the first player is a worker, and the second player is a supervisor.

Second, Bergemann and Morris (2016a) consider an environment with two-sided strategic complementarities but where a nonstrategic payoff externality removes intrinsic preferences over correlation. To illustrate this, suppose that we take our many player with no prior information example from section 3.3, but now suppose that—in addition to the existing payoffs—each firm would like the other firm to invest, and thus there are spillovers. In the following payoff table, we are assuming that each firm gets an extra payoff of $z > 0$ if the other firm invests:

$\theta = B$	invest	not invest
invest	$-1 + \varepsilon + z$	-1
not invest	z	0

$\theta = G$	invest	not invest
invest	$x + \varepsilon + z$	x
not invest	z	0

Observe that this change in payoffs has no impact on the firms' best responses: neither firm can influence whether the other firm invests. But now suppose that the government is interested in maximizing the sum of the firms' payoffs. Consider the case that z is very large. As z becomes larger and larger, the government's objective will approach maximizing the sum of the probabilities that each firm invests. In this sense, the government's instrumental preference for correlation is micro-founded in the benevolent government's desire to make each firm invest in the interests of the other firm. This example illustrates a distinctive point about strategic information design. Recall that in the one-player case where the designer and the player have common interests, it is always optimal for the designer to fully reveal all information in order to allow the

player to take an action that is optimal given their shared preferences. In the many-player case, however, the players themselves may not act in their joint interest for the usual (noncooperative strategic) reasons. In this case—as in the above example—a benevolent information designer might want to obfuscate information.

For a last case with only instrumental concerns over correlation, Bergemann and Morris (2013b) considered quantity (Cournot) competition in a market, where the information designer wants to maximize the sum of the firms' payoffs, i.e., the industry profits.¹⁰ A continuum of firms choose output where there is uncertainty about the intercept of the demand curve, i.e., the level of demand. In this case, the information designer would like the firms' total output to be correlated with the level of demand, but total profits do not depend on the correlation of firms' output conditional on the level of aggregate output. However, firms would like their actions to be negatively correlated (because the game is one of strategic substitutes); but they would also like output to be correlated with the state. The information designer can induce players to make total output choices that are closer to the optimal level, but allow them to negatively correlate their output. In the optimal outcome (for some parameters), firms observe conditionally independent private signals about the state of demand, trading off these two objectives.¹¹

Having considered the case where the information designer cares about correlation for instrumental but not intrinsic reasons, we can also consider the opposite case where the information designer cares about correlation

¹⁰This corresponds to a large literature on information sharing in oligopoly following Novshek and Sonnenschein (1982).

¹¹In this setting, the information designer would like to induce firms to lower output on average, but cannot do so. The designer can only influence correlation.

for intrinsic but not instrumental reasons. We can illustrate this case with the example of section 3.3, also. Suppose that the payoffs remain the same, but now the government would like to maximize the probability that at least one firm invests, so that the government has intrinsic preferences over correlation. But in this case—under our maintained assumption that $x > 1/2$ —it is possible to ensure that one firm always invests. Consider the following decision rule:

$\theta = B$	invest	not invest
invest	0	$\frac{1}{2}$
not invest	$\frac{1}{2}$	0

$\theta = G$	invest	not invest
invest	1	0
not invest	0	0

If ε were equal to 0, this decision rule would be obedient, with all constraints holding strictly: a firm told to not invest would have a strict incentive to obey, since it would know that the state was bad; a firm told to invest would have a strict incentive to obey, since its expected payoff will be $\frac{2}{3}(x - \frac{1}{2}) > 0$. Because the obedience constraints hold strictly, this decision rule will continue to be obedient, for positive or negative ε , as long as $|\varepsilon|$ is sufficiently small. Note that the government's objective of maximizing the probability that at least one firm invests necessitates private signals.

In a dynamic setting, Ely (2017) shows how an information designer with intrinsic preferences for negative correlation will optimally use private signals to induce it (he also shows that this is consistent with players' strategic objectives). Arieli and Babichenko (2016) and Meyer (2017) provide a general analysis of optimal information design when

players have binary actions and the information designer has an *intrinsic* motive for correlation, but there is no strategic interaction—and thus no instrumental motive for caring about correlation. With supermodular payoffs, public signals are optimal, whereas with submodular payoffs, private signals are optimal and it is optimal to minimize correlation.¹²

4.2 *Tightening Obedience Constraints and BCE Outcomes*

There is never any reason for an information and/or mechanism designer to provide players with more information than they will use in making their choices. Giving more information will impose more incentive constraints on players' choices, and thus reduce the ability of an information designer to attain outcomes that are desirable for him. In dynamic mechanism design, giving players information about others' past reports will tighten truth-telling constraints. Myerson (1986) emphasizes that a similar observation is true in dynamic problems of communication in games; in these games the extra information imposes more obedience constraints as well. Recall that in our language, communication in games corresponds to information design when the information designer has no information of her own.

Our examples have illustrated this general observation: giving players more information will impose more obedience constraints and thus reduce the set of (BCE) outcomes that can occur. However, the examples illustrate a more subtle point that is the focus of Bergemann and Morris (2016a): it is not only true that sending additional signals reduces the set of outcomes that can occur; it is

¹²In a recent survey on the algorithmic aspects of information design, Dughmi (2017) emphasizes how the structure of the payoff environment impacts the algorithms to compute the optimal information structure and their computational complexity.

also possible to construct a partial order on arbitrary information structures that exactly characterizes the notion of “more informative,” and one that corresponds exactly to adding more obedience constraints.

This was illustrated in our one-player example with prior information. In that example, the set of implementable BCE outcomes shrunk in size as the accuracy q of the prior information increased (as illustrated in figure 2). As q increases, we are intuitively giving the player more information. The additional information is given to the player not in the form of more signals, but rather more precise signals.

We will now informally describe how this observation can be generalized in many directions. In the one-player case, an information structure reduces to an experiment in the sense of Blackwell (1951, 1953). He defines the “more informative” ordering in terms of a *feasibility ordering*. An experiment is said to be more informative than another experiment if the set of outcomes (joint distributions over *actions* and *states*) that can be induced by decision rules mapping *signals* into *actions* is larger—in any decision problem—under the first experiment.¹³

Blackwell (1951, 1953) offered an alternative, entirely *statistical* ordering on experiments, without reference to actions or payoffs: one experiment is *sufficient* for another if the latter can be attained by adding noise to the former.

We have described an *incentive* ordering on experiments: one experiment is more

incentive constrained than another if the set of BCE outcomes under the former experiment is smaller (reflecting the tighter obedience constraints) in every decision problem (or one-player basic game).

Taken together, there is now an elegant set of connections between Blackwell’s theorem and the information design problem. One can show that there is a three-way equivalence between (i) the “more informative” ordering; (ii) the “sufficiency” ordering; and (iii) the “incentive-constrained” ordering. Blackwell’s theorem shows an equivalence between (i) and (ii). Thus, if an experiment is “more informative” in the sense of Blackwell, then—in any decision problem—the set of BCE outcomes for a given experiment is *smaller* under the more informative experiment. There is naturally also a converse. If an experiment is not more informative than another, then one can find a decision problem and an outcome that is a BCE for the first experiment but not for the second.

This equivalence result holds in the one-player case for general games, i.e., decision problems with finitely many states and actions. This result is the one-player special case of the main result (for many-player information structures) from Bergemann and Morris (2016a).

The definition of the incentive ordering, as well as the definition of the feasibility ordering, generalizes naturally to the many-player case. Bergemann and Morris (2016a) offer a new statistical ordering, termed *individual sufficiency*, for many players, and show that is equivalent to the incentive ordering in the many-player case. Individual sufficiency is defined as follows. Fix two information structures. A combined information structure is one where players observe a pair of signals, corresponding to the two information structures, with the marginal on signal profiles of each information structure corresponding to the original information structures. Thus, there are many combinations of any two

¹³Blackwell (1951, 1953) defines “more informative” in terms of “risk vectors” rather than joint distributions between states and action. These two feasibility conditions are equivalent. Blackwell defines as risk vector the vector of expected payoffs that can be sustained by a decision rule measurable with respect to the information structure alone. The resulting payoff vector is, however, computed conditional (on the vector) of the true state. It then follows easily that a larger set of risk vectors is sustained if and only if a larger set of joint distribution of actions and states is sustained.

information structures corresponding to different ways of correlating signals across the two information structures. One information structure is now individually sufficient for another if there is a combined information structure such that each player's signal in the former information structure is a sufficient statistic for his beliefs about the state of the world and others' signals in the latter information structure. A subtle feature of this ordering is that one information structure being individually sufficient for another neither implies nor is implied by the property that players' *joint information* in the former case is sufficient (in the statistical sense) for their joint information in the latter case. We should add that in the special case of a single player individual sufficiency and sufficiency naturally coincide. The ordering of *individual sufficiency* has a number of natural properties. Two information structures are individually sufficient for each other if and only if they correspond to the same beliefs and higher-order beliefs about states, and differ only in the redundancies of the type identified in Mertens and Zamir (1985). One information structure is individually sufficient for another only if we can get from the latter to the former by providing additional information and removing redundancies.

In the one-player setting, there is an alternative but equivalent ordering to the feasibility ordering. It is phrased in terms of optimality and appears more frequently in the economics, rather than statistic, literature. For example, Laffont (1989) defines one experiment to be "more valuable" than another if it leads to (weakly) higher maximal expected utility for every decision problem. In a single-player problem, a larger set of feasible joint distributions clearly implies a larger maximal expected utility. Less obvious perhaps, the converse also holds, as the ranking has to hold across all decision problems.

A common observation is that in strategic situations, there is no many-player analogue

to the "more valuable" ordering; see, for example, Hirshleifer (1971); Neyman (1991); Gossner (2000); and Bassan et al. (2003). The above discussion provides a novel perspective. Intuitively, there are two effects of giving players more information in a strategic setting. First, it allows players to condition on more informative signals, and thus—in the absence of incentive constraints—attain more outcomes. Second, more information can reduce the set of attainable outcomes by imposing more incentive constraints on players' behavior. The value of information in strategic situations is ambiguous in general because both effects are at work. Following Lehrer, Rosenberg, and Shmaya (2010), we can abstract from the second (incentive) effect by focusing on common-interest games. Here, more information in the sense of individual sufficiency translates into more attainable outcomes. But looking at BCE abstracts from the first (feasibility) effect by allowing the information designer to supply any information to the players. Now, more information in the sense of individual sufficiency translates into less attainable outcomes.

4.3 *BCE Outcomes and Information Design without Concavification*

We have described a "two-step" approach to solving information design problems. First, provide a linear algebraic characterization of *implementable outcomes*, meaning the set of joint distributions over actions and states that can be induced by some information structure that the information designer might choose to give the players. The set of implementable outcomes is exactly the set of BCE. Second, we select among the BCE the one that is optimal for the information designer. This second step implicitly identifies the optimal information structure. The first problem is solved by finding the set of outcomes that satisfy a set of linear (obedience) constraints. The second

problem corresponds to maximizing a linear objective subject to linear constraints. Both steps of this problem are well behaved. There is a separate reason why we might pursue this two-step procedure: for many questions of interest, it is critical to first understand the set of BCE outcomes. The next two subsections describe two contexts where the structure of the BCE outcomes is the focus of the analysis.

However, there is a different approach to information design: concavification. The “concavification” approach is based on a geometric analysis of the function mapping receiver posterior beliefs to sender payoffs. Concavification focuses more on the “experiments” or the distribution of posteriors that are induced for the receivers, rather than on the joint distribution between actions and states. In the one-person problem, we can identify the payoff that the information designer receives for any given probability distribution over states, subject to the fact that the player will make an optimal choice. But the information designer has the ability to *split* the player’s beliefs about the state, i.e., supply the player with information that will induce any distribution over posteriors with the constraint that the prior equals the expected posterior. This implies that the set of attainable payoffs for the information designer, as a function of prior distributions of states, is the concavification of the set of payoffs of the designer in the absence of information design. This concavification argument (building on Aumann and Maschler 1995) is the focus of both Kamenica and Gentzkow (2011) and the large and important literature inspired by their work. The many-players case is significantly harder than the single-player case, as it is no longer the set of probability distributions over states that matters, but rather the set of (common prior) subsets of the universal type space of Mertens and Zamir (1985) that are relevant for strategic

analysis. Mathevet, Perego, and Taneva (2017) describe this generalization of concavification for the many-player case.

Concavification and its many-player analogue are important for two reasons. First, they offer structural insights into the information design problem. Second, they provide a method for solving information design problems. As a solution method, the concavification approach and its generalization do not always help without some special structure. Our own work (Bergemann, Brooks, and Morris 2015) on (one player) price discrimination, discussed in the next section, relies heavily on linear programming; but while the solution must correspond to the concavification of an objective function, it is very difficult to visualize the concavification or provide a proof using it. However, linear programming methods do not always help, either: in our work on (many player) auctions (Bergemann, Brooks, and Morris 2017a), neither generalized concavification nor linear programming methods are used in stating or proving our results (although linear programming played an important role in supplying conjectures for the results).

5. *Metaphorical Information Design and Applications*

Mechanism design sometimes has a literal interpretation. For example—in some settings—a seller may be able to commit to an auction for selling an object. In other settings, the mechanism design problem is studied even though there does not exist a mechanism designer able to commit. For example, suppose that we are interested in a buyer and seller bargaining over an object. There may be no rules for how the players bargain and no one who could enforce such rules. Nonetheless, Myerson and Satterthwaite (1983) studied what would be the optimal mechanism for realizing gains

from trade because it bounds what could happen under any bargaining protocol that ends up being used. In this sense, there is not a literal mechanism designer, but we are rather using the language of mechanism design for another purpose.

Similarly for information design, the most literal interpretation of the information design problem is that there is an actual information designer who can commit to choosing the players' information structure in order to achieve a particular objective. In many contexts, this commitment assumption may not be plausible.¹⁴ Yet, the information design perspective can be used to address many important questions even where there is not a literal information designer. In particular, understanding the set of outcomes that an information designer can induce corresponds to identifying the set of all outcomes that could arise from some information structure.

In our own applications of information design, we have mostly been interested in metaphorical interpretations. The set of BCE is precisely the set of outcomes that can arise with additional information for a given basic game and prior information structure. If there are properties that hold for all BCE, we have identified predictions that are robust to the exact information structure. Identifying the best or worst outcome that can arise under some information structure according to some objective function as criterion is the same as solving an information design problem where the designer is maximizing or minimizing that criterion. In this section, we will review two such economic applications of information design. We will highlight the implications of this approach in the context of third-degree price discrimination (Bergemann, Brooks, and Morris 2015)

and a linear interaction game with a focus on aggregate variance, and macroeconomic implications (Bergemann, Heumann, and Morris 2015).

Caplin and Martin (2015) adopt a similar, metaphorical, approach to the recovery of preference orderings and utility from choice data. They allow for the possibility that the decision maker has imperfect information while satisfying Bayes's law and iterated expectation. They ask what they can learn from the observed choice data about the underlying preference profile without making strong assumptions on the information available to the decision maker at the moment of choice. In related work, Caplin and Dean (2015) develop a revealed preference test giving conditions under which apparent choice "mistakes" can be explained in terms of optimal costly information acquisition by the player in the presence of imperfect information.

5.1 *The Limits of Price Discrimination*

A classic issue in the economic analysis of monopoly is the impact of discriminatory pricing on consumer and producer surplus. A monopolist engages in *third-degree price discrimination* if he uses additional information about consumer characteristics to offer different prices to different segments of the aggregate market. Bergemann, Brooks, and Morris (2015) characterize what could happen to consumer and producer surplus for all possible segmentations of the market.

One can provide some elementary bounds on consumer and producer surplus in any market segmentation. First, consumer surplus must be nonnegative as a consequence of the participation constraint: a consumer will not buy the good at a price above his valuation. Second, the producer must get at least the surplus that he could get if there was no segmentation and he had no additional information beyond the prior distribution. In this case, *an* optimal policy is

¹⁴Forges and Koessler (2005) observe that conditioning on players' exogenous information makes sense if players' types are ex post verifiable.

always to offer the product with probability one at a given price to all buyers. We therefore refer to it as the *uniform* monopoly price, and correspondingly, uniform monopoly profit. Third, the sum of consumer and producer surplus cannot exceed the total social value that is generated by the good, which is willingness-to-pay minus unit cost of production. The shaded right-angled triangle in figure 5 illustrates these three bounds.

The main result in Bergemann, Brooks, and Morris (2015) is that *every* welfare outcome satisfying these constraints is attainable by some market segmentation. This is the entire shaded triangle in figure 5. If the monopolist has no information beyond the prior distribution of valuations, there will be no segmentation. The producer charges the optimal monopoly price and gets the associated monopoly profit, and consumers receive a positive surplus; this is marked by point A in figure 5. If the monopolist has complete information, then he can charge each buyer his true valuation, i.e., engage in perfect or *first-degree price discrimination*; this is marked by point B. The point marked C is where consumer surplus is maximized; the outcome is efficient and the consumer gets all the surplus gains over the uniform monopoly profit. At the point marked D, social surplus is minimized by holding producer surplus down to uniform monopoly profits and holding consumer surplus down to zero.

The main result states that we can make only very weak predictions about producer and consumer surplus. It can be understood as the outcome of a set of metaphorical information design problems. If an information designer wanted to maximize consumer surplus, she would choose point C. If she wanted to minimize consumer surplus, or producer surplus, or any weighted combination of the two, she could choose point D. Any other point on the boundary of the

triangle is the solution to some maximization problem of the information designer defined by some preferences over producer and consumer surplus.

The information design problem has a very clear literal interpretation in the case where the monopolist knows the consumer's valuation. She can then achieve perfect price discrimination at point B. However, giving a literal information design interpretation of point C is more subtle. We would need to identify an information designer who knew consumers' valuations and committed to give partial information to the monopolist in order to maximize the sum of consumers' welfare. Importantly, even though the disclosure rule is optimal for consumers as a group, individual consumers would not have an incentive to truthfully report their valuations to the information designer, given the designer's disclosure rule, since they would want to report themselves to have low values.

As discussed in the previous section, it seems hard to explain the main result using concavification, but there is an elementary geometric argument. One can show that any point where the monopolist is held down to his uniform monopoly profits with no information beyond the prior distribution—including outcomes A, C, and D in figure 5—can be achieved *with the same segmentation*. In this segmentation, consumer surplus varies because the monopolist is indifferent between charging different prices.

We can use a simple example to illustrate these results. There are three valuations, $v \in V = \{1, 2, 3\}$, which arise in equal proportions, and there is zero marginal cost of production. The feasible social surplus is $w^* = (1/3)(1 + 2 + 3) = 2$. The uniform monopoly price is $v^* = 2$. Under the uniform monopoly price, profit is $\pi^* = (2/3) \times 2 = 4/3$ and consumer surplus is $u^* = (1/3)(3 - 2) + (1/3)(2 - 2) = 1/3$.




Figure 5. The Bounds on Profits and Consumer Surplus in Third-Degree Price Discrimination

A segment x is a vector of probabilities of each valuation, thus $x = (x_1, x_2, x_3)$, and by $\sigma(x)$ we denote the total mass of a segment x . A *segmentation* of the market is therefore a collection of segments $x \in X$ and a probability distribution $\sigma(\cdot)$ over the segments. We give an example of a segmentation below. In the example, there are three segments and each segment is identified by its support on the valuations indicated by the set $\{\cdot\}$ in the superscript. The frequency of each segment x is given by $\sigma(x)$:

(16)

Segment	x_1	x_2	x_3	$\sigma(x)$
$x^{(1,2,3)}$	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{2}{3}$
$x^{(2,3)}$	0	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{1}{6}$
$x^{(1)}$	0	1	0	$\frac{1}{6}$
x^*	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	1

The particular segmentation has a number of interesting properties. First, in each segment, the seller is *indifferent* between

charging as price *any* valuation that is in the support of the segment. Second, the uniform monopoly price, $p^* = 2$, is in the support of every segment. Thus, this particular segmentation preserves the uniform monopoly profit. If the monopolist charges the uniform monopoly price on each segment, we get point A. If he charges the lowest value in the support of each segment (which is also an optimal price, by construction), we get point C; and if he charges the highest value in the support, we get point D.

Roesler and Szentes (2017) consider a related information design problem in which a single buyer can design her own information about her value before she is facing a monopolist seller. While the analysis of the third-degree price discrimination proceeds as a one-player application, the arguments extend to many-player settings. Bergemann, Brooks, and Morris (2017a) pursue the question of how private information may impact the pricing behavior in a many-buyer environment. There, we derive results about equilibrium behavior in the first-price auction that hold across all common-prior information structures. The results that we obtain can be used for a variety of applications, e.g., to partially identify the value distribution in settings where the information structure is unknown and to make informationally robust comparisons of mechanisms.

5.2 Information and Volatility

Bergemann, Heumann, and Morris (2015) revisit a classic issue in macroeconomics. Consider an economy of interacting agents—each of whom picks an action—where the agents are subject to idiosyncratic and aggregate shocks. A fundamental economic question in this environment is to ask how aggregate and idiosyncratic shocks map into “aggregate volatility”—the variance of the average action. Versions of this question arise in many different economic contexts. In particular, a central question in macro-

economics is how aggregate and individual productivity shocks translate into variation in GDP. Another classical question is when and how asymmetric information can influence this mapping, and in particular exacerbate aggregate volatility.

These questions are studied in a setting with a continuum of agents whose best responses are linear in the (expectation of the) average action of others and in the idiosyncratic as well as aggregate shocks. Shocks, actions, and signals are symmetrically normally distributed across agents, maintaining symmetry and normality of the information structure. The maximal aggregate volatility is attained in an information structure in which the agents confound idiosyncratic and aggregate shocks, and display excess response to the aggregate shocks, as in Lucas (1972) and more recently in Hellwig and Venkateswaran (2009), Venkateswaran (2013), and Angeletos and La’O (2013). Our contribution is to highlight that, in this setting with idiosyncratic and aggregate shocks, a class of noise-free confounding information structures are *extremal* and provide *global bounds* on how much volatility can arise. In particular, for any given variance of aggregate shocks, the upper bound on aggregate volatility is linearly increasing in the variance of the idiosyncratic shocks.

In this application, we do not think there is any economic agent who is able to or wants to maximize aggregate volatility. But because we are interested in bounds on aggregate volatility across equilibria of different information structures, the problem is naturally represented as an information design problem.

The basic setting is that the payoff shock θ_i of individual i is given by the sum of an *aggregate shock* θ and an *idiosyncratic shock* ε_i :

$$\theta_i \triangleq \theta + \varepsilon_i.$$

The aggregate shock θ is common to all agents and the idiosyncratic shock ε_i is independent

and identically distributed across agents, as well as independent of the aggregate shock. Each component of the payoff shock θ_i is normally distributed:

$$(17) \quad \begin{pmatrix} \theta \\ \varepsilon_i \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_\theta \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & 0 \\ 0 & \sigma_\varepsilon^2 \end{pmatrix}\right).$$

The variance of the individual payoff shock θ_i can be expressed in terms of the variance of the sum of the idiosyncratic and the aggregate shock: $\sigma_\theta^2 + \sigma_\varepsilon^2$. The correlation (coefficient) ρ_θ between the payoff shocks θ_i and θ_j of any two agents i and j is:

$$(18) \quad \rho_\theta \triangleq \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}.$$

The best response of agent i is given by a linear function

$$a_i = (1 - r) \mathbb{E}_i[\theta_i] + r \mathbb{E}_i[A],$$

where A is the average action. Now, the results described above regarding the structure of the extremal information structure hold independently of whether the weight on the average action, r , is negative (the strategic substitutes case), zero (the purely decision theoretic case), or positive (the strategic complements case). A striking property is that the set of feasible correlations between individual and average actions and individual and aggregate shocks is independent of r and determined only by statistical constraints. Here we will thus convey the flavor of the result in a setting where the decision of the agent is independent of any strategic considerations, thus the decision-theoretic case.

It suffices to consider the following *one-dimensional* class of signals:

$$(19) \quad s_i \triangleq \lambda \varepsilon_i + (1 - \lambda) \theta,$$

where the linear composition of the signal s_i is determined by the parameter $\lambda \in [0, 1]$.

The information structure λ is *noise free* in the sense that every signal s_i is a linear combination of the idiosyncratic and the aggregate shock, ε_i and θ , and no extraneous noise or error term enters the signal of each agent. Nonetheless, since the signal s_i combines the idiosyncratic and the aggregate shock with weights λ and $1 - \lambda$, each signal s_i leaves agent i with residual uncertainty about his true individual payoff shock θ_i , unless $\lambda = 1 - \lambda = 1/2$.

In the decision-theoretic case, $r = 0$, the best response of each agent simply reflects a statistical prediction problem, namely to predict the payoff shock θ_i given the signal s_i :

$$(20) \quad a_i = \mathbb{E}[\theta_i | s_i] = \frac{(1 - \lambda) \rho_\theta + \lambda(1 - \rho_\theta)}{(1 - \lambda)^2 \rho_\theta + \lambda^2(1 - \rho_\theta)} s_i.$$

The individual prediction problem is more responsive to the signal s_i , that is, assigns a larger weight to s_i if and only if the signal contains more information about the individual payoff shock θ_i . The noise-free information structure $\lambda = 1/2$ allows each agent to perfectly infer the individual payoff shock θ_i . It follows that the responsiveness, and hence the variance of the individual action σ_a^2 , is maximized at $\lambda = 1/2$:

$$\sigma_a^2 = \sigma_\theta^2 + \sigma_\varepsilon^2.$$

Now, to the extent that the individual payoff shocks, θ_i and θ_j , are correlated, we find that even though each agent i only solves an individual prediction problem, his actions are correlated by means of the underlying correlation of the individual payoff shocks, and the resulting aggregate volatility is:

$$\sigma_A^2 = \rho_a \sigma_a^2 = \sigma_\theta^2.$$

We can ask whether the aggregate volatility can reach higher levels under information structures different from $\lambda = 1/2$. As the information structure departs from $\lambda = 1/2$, we necessarily introduce a bias in the signal

s_i toward one of the two components of the payoff shock θ_i . Clearly, the signal s_i is losing its informational quality with respect to the individual payoff shock θ_i as λ moves away from $1/2$ in either direction. Thus, the individual prediction problem (20) is becoming noisier and in consequence, the response of the individual agent to the signal s_i is *attenuated*. But a larger weight, $1 - \lambda$, on the aggregate shock θ , may support correlation in the actions across agents, and thus support larger aggregate volatility. As the response of the agent is likely to be attenuated, a trade-off appears between bias and loss of information. We show that the maximal aggregate volatility:

$$(21) \quad \max_{\lambda} \{\text{var}(A)\} = \frac{(\sigma_{\theta} + \sqrt{\sigma_{\theta}^2 + \sigma_{\varepsilon}^2})^2}{4}$$

is achieved by the information structure λ^* :

$$(22) \quad \lambda^* \triangleq \arg \max_{\lambda} \{\text{var}(A)\} \\ = \frac{\sigma_{\theta}}{2\sigma_{\theta} + \sqrt{\sigma_{\theta}^2 + \sigma_{\varepsilon}^2}} < \frac{1}{2},$$

which biases the signal toward the aggregate shock. We can express the information structure that maximizes the aggregate volatility in terms of the correlation coefficient ρ_{θ} :

$$\arg \max_{\lambda} \{\text{var}(A)\} = \frac{\sqrt{\rho_{\theta}}}{1 + 2\sqrt{\rho_{\theta}}},$$

and the maximal volatility can be expressed as:

$$\max_{\lambda} \{\text{var}(A)\} = \frac{1}{4}(1 + \sqrt{\rho_{\theta}})^2(\sigma_{\theta}^2 + \sigma_{\varepsilon}^2).$$

Surprisingly, as we approach an environment with purely idiosyncratic shocks, the maximal *aggregate* volatility does not converge to zero; rather, it is bounded away from zero, and given by $\sigma_{\varepsilon}^2/4$. Thus, the economy can

maintain a large aggregate volatility even in the presence of vanishing aggregate payoff shocks by confounding the payoff-relevant information about the idiosyncratic shock with the (in the limit) payoff-irrelevant information about the aggregate shock.

6. Information Design with Private Information

We have thus far considered the scenario where the designer knows not only the true state θ , but also the players' prior information about the state. We now consider what happens when the information designer does not have access to players' prior information (but still can condition on the state). Here, we consider two alternative assumptions about the designer's ability to condition recommendations on players' prior information. If the designer can elicit the private information, then we have *information design with elicitation*. If the designer cannot elicit the private information, we have *information design without elicitation*. We present a self-contained discussion of these extensions and then discuss how—with these extensions—information design fits into the mechanism design and incomplete-information-correlated equilibrium literatures more broadly in sections 6.3 and 6.4.

6.1 Players' Prior Information is not Known to the Designer

When the information designer cannot observe the players' prior information, she may or may not be able to ask the players about it. In the case of information design with elicitation, she will be able to condition her recommendations on the reported types. In the case of information design without elicitation, she can only send a list of recommendations, namely one recommendation for each possible type of the player.

In the case of information design with elicitation, the revelation principle still implies

that we can restrict attention to the case where the information sent by the information designer consists of action recommendations. However, we will now require an incentive compatibility condition that entails *truth telling* as well as *obedience*, so that the information designer can only condition on a player's signal if the player can be given an incentive to report it truthfully. Following Myerson (1991, section 6.3), we can think of the information designer choosing a decision rule $\sigma : T \times \Theta \rightarrow \Delta(A)$, but each type of each player can choose a deviation $\delta_i : A_i \rightarrow A_i$ with the interpretation that $\delta_i(a_i)$ is the action chosen by player i if the information designer recommended action a_i . The decision rule σ is incentive compatible if each player does not have an incentive to deviate.

DEFINITION 2 (Incentive Compatible): A decision rule $\sigma : T \times \Theta \rightarrow \Delta(A)$ is incentive compatible for (G, S) if for each $i = 1, \dots, I$ and $t_i \in T_i$,

$$\begin{aligned} & \sum_{a_i, \theta} u_i((a_i, a_{-i}), \theta) \sigma((a_i, a_{-i}) | (t_i, t_{-i}), \theta) \pi((t_i, t_{-i}) | \theta) \psi(\theta) \\ & \geq \sum_{a_i, \theta} u_i((\delta_i(a_i), a_{-i}), \theta) \sigma((a_i, a_{-i}) | (t'_i, t_{-i}), \theta) \pi((t_i, t_{-i}) | \theta) \psi(\theta) \end{aligned}$$

for all $t'_i \in T_i$ and $\delta_i : A_i \rightarrow A_i$.

The displayed inequality will be referred to as player i 's type- t_i incentive constraint. It ensures that player i , after observing signal t_i , finds it optimal to report his signal truthfully and then, after observing and updating on the information contained in the resulting action recommendation a_i , finds it optimal to follow this recommendation. Thus it builds in both truth telling and obedience. In addition, the notion of incentive compatibility requires that the decision rule is immune to "double deviations" in which the player misreports his type to be t'_i (rather than t_i) and disobeys the recommendation of the

designer by choosing $\delta_i(a_i)$ rather than a_i . Thus, incentive compatibility implies, but is not implied by, separately requiring truth telling and obedience.

PROPOSITION 2: An information designer with elicitation can attain a decision rule if and only if it is incentive compatible.

In the case of information design *without* elicitation, the designer cannot condition the action recommendation on the reported type, but has to offer a *contingent* recommendation, a *vector* of action recommendations, where each individual entry is an action recommendation for a specific type of the player, hence contingent on the realized type. The set of feasible recommendations to player i is therefore given by $B_i = A_i^{T_i}$. The set of player i 's contingent recommendations therefore has a typical element

$$b_i : T_i \rightarrow A_i.$$

We define $B = \prod_{i=1}^I B_i$ and let a generic element be given by $b = (b_1, \dots, b_I) \in B$, and so $b(t) = (b_1(t_1), \dots, b_I(t_I))$,

We are now interested in contingent action recommendations $\phi : \Theta \rightarrow \Delta(B)$ rather than action recommendations.

DEFINITION 3 (Public Feasibility): A decision rule $\sigma : T \times \Theta \rightarrow \Delta(A)$ is publicly feasible if there exists a contingent recommendation $\phi : \Theta \rightarrow \Delta(B)$ such that for each $a \in A$, $t \in T$, and $\theta \in \Theta$ with $\pi(t | \theta) > 0$,

$$\sigma(a | t, \theta) = \sum_{\{b \in B : b(t) = a\}} \phi(b | \theta).$$

In this case, we say that σ is induced by ϕ .

Public feasibility is the restriction that a given player's contingent recommendation

cannot depend on the type of the player himself nor on the types of the other players. We refer to the above requirement as public feasibility, since the recommendation vector *cannot* be tailored to the private information of the players and it has to be feasible in the sense that it induces the decision rule σ . We emphasize that each contingent recommendation b_i is still communicated to each agent i separately and privately, and thus is *not* public in the sense of a public announcement to all players.

When $I = 1$, public feasibility is a vacuous restriction. Every decision rule σ is induced by the contingent recommendation ϕ given by

$$\phi(b|\theta) = \prod_{t \in T} \sigma(b(t)|t, \theta).$$

Under this choice of ϕ , the components of the contingent recommendation vector $b(t)$ are drawn independently. When $I > 1$, however, public feasibility is a substantive restriction. By recommending to a particular player a strategy rather than an action, the designer can condition that player's action on his type. By judiciously choosing a distribution over B , the designer can even correlate the players' strategies. But she cannot correlate one player's strategy on another player's *type*.

We are not interested in all contingent recommendations, but rather those that are obedient in the sense defined earlier, in definition 1. Below, we adapt the definition to account for the larger space of contingent recommendations, b , rather than action recommendations, a .

DEFINITION 4 (Publicly Feasible Obedience): *A decision rule $\sigma : T \times \Theta \rightarrow \Delta(A)$ is publicly feasible obedient if there exists a contingent recommendation $\phi : \Theta \rightarrow \Delta(B)$ such that (i) ϕ induces σ , and (ii) ϕ satisfies obedience in the sense that for each $i, t_i \in T_i$, and $b_i \in B_i$,*

$$\begin{aligned} & \sum_{b_i, t_i, \theta} u_i((b_i(t_i), b_{-i}(t_{-i})), \theta) \phi((b_i, b_{-i})|\theta) \pi((t_i, t_{-i})|\theta) \psi(\theta) \\ & \geq \sum_{b_i, t_i, \theta} u_i((a'_i, b_{-i}(t_{-i})), \theta) \phi((b_i, b_{-i})|\theta) \pi((t_i, t_{-i})|\theta) \psi(\theta) \end{aligned}$$

for all $a'_i \in A_i$.

The displayed inequality will be referred to as player i 's (t_i, b_i) -publicly feasible obedience constraint. It ensures that player i , after observing signal t_i and receiving and updating on the recommendation b_i , finds it optimal to take the action $b_i(t_i)$ prescribed by the vector b_i for his type t_i . Note that in so far as a contingent recommendation b_i reveals more information about the state θ and thus possibly and indirectly about the type profile t_{-i} than just an action recommendation, publicly feasible obedience will be a more demanding concept than mere obedience as it allows the agents to contemplate deviations based on more accurate information. After all, type t_i is able to observe the recommendation tailored toward him as well as the recommendation offered to all other types $t'_i \neq t_i$. To the extent that the recommendation across types is not perfectly correlated, it then follows that the type t_i will receive additional information through the contingent recommendation rather than the action recommendation.

PROPOSITION 3: *An information designer without elicitation can attain a decision rule if and only if it is publicly feasible obedient.*

6.2 The Investment Example Revisited

We now illustrate information design with private information using the investment example introduced earlier in section 3. The issues that arise in the one-player case with private information are already interesting and subtle. Here we do not present any examples of elicitation with many players. We thus do not consider any additional issues of public feasibility that only arise in the many-player case.

We therefore now allow the information of the firm to be private. The government does not know the realization of the signal that the firm observes but can or cannot elicit it.

6.2.1 Information Design with Elicitation

In the case of elicitation, we have a screening problem where the designer offers a recommendation that induces a probability of investing as a function of the reported signal and the true state. As noted above, we have three sets of constraints that need to be satisfied. First, each type has to truthfully report his signal; 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, must not be profitable. Kolotilin et al. (2017) refer to this informational environment as “private persuasion.”¹⁵ Kolotilin (2018) pursues a linear programming approach to identify the optimal information disclosure policy under a single-crossing assumption.

A decision rule now specifies the probability of investment $p_{\theta t}$ conditional on the true state $\theta \in \{B, G\}$ and the reported type $t \in \{b, g\}$. Thus, as before in section 3.2, a decision rule is now a vector $p = (p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg})$. The information designer offers a recommendation (stochastically) as a function of the true state and the reported type. The obedience conditions are as in section 3.2, where information was not private. A truthful reporting constraint is

described below for a good type $t = g$. The truth-telling constraint for the good type $t = g$ is:

$$(23) \quad qp_{Gg}x - (1 - q)p_{Bg} \geq qp_{Cb}x - (1 - q)p_{Bb},$$

and correspondingly for the bad type $t = b$:

$$(24) \quad (1 - q)p_{Cb}x - qp_{Bb} \geq (1 - q)p_{Gg}x - qp_{Bg}.$$

By misreporting and then following the resulting recommendation afterwards, each type can change the probability of investing. We can write the above two truth-telling constraints in terms of a bracketing inequality:

$$(25) \quad \frac{1 - q}{q}x(p_{Gg} - p_{Cb}) \leq p_{Bg} - p_{Bb} \\ \leq \frac{q}{1 - q}x(p_{Gg} - p_{Cb}).$$

These inequalities highlight how the differential between type $t = b$ and $t = g$ in the recommendation p_{Bt} in the bad state are bounded, below and above, by the differential in the recommendation p_{Gt} in the good state. Notice also that since

$$\frac{1 - q}{q} \frac{1}{x} < \frac{q}{1 - q}x,$$

the above bracketing inequality requires that

$$p_{Gg} - p_{Cb} \geq 0, \quad p_{Bg} - p_{Bb} \geq 0,$$

thus, the conditional probability of investing has to be larger for the good type than the bad type in either state. An implication specific to the binary action, binary state environment is the fact that double deviations do not impose any additional restriction on the behavior of the player. We state and prove this result formally in the appendix as proposition 6. In particular, this means that the obedience and truth-telling constraints discussed above imply the incentive compatibility condition of definition 2.

¹⁵Bergemann, Bonatti, and Smolin (2018) consider a model of private persuasion with quasilinear utility. Their main objective is to analyze the revenue maximizing solution to the information design problem subject to the elicitation constraints. Daskalakis, Papadimitriou, and Tzamos (2016) also consider an information design with quasilinear utility. The novel aspect of their analysis is that the object for sale has many attributes, and the seller chooses optimally how much to disclose about each individual attribute. Their analysis reveals a close relationship to the classic bundling problem of a multi-item monopolist.

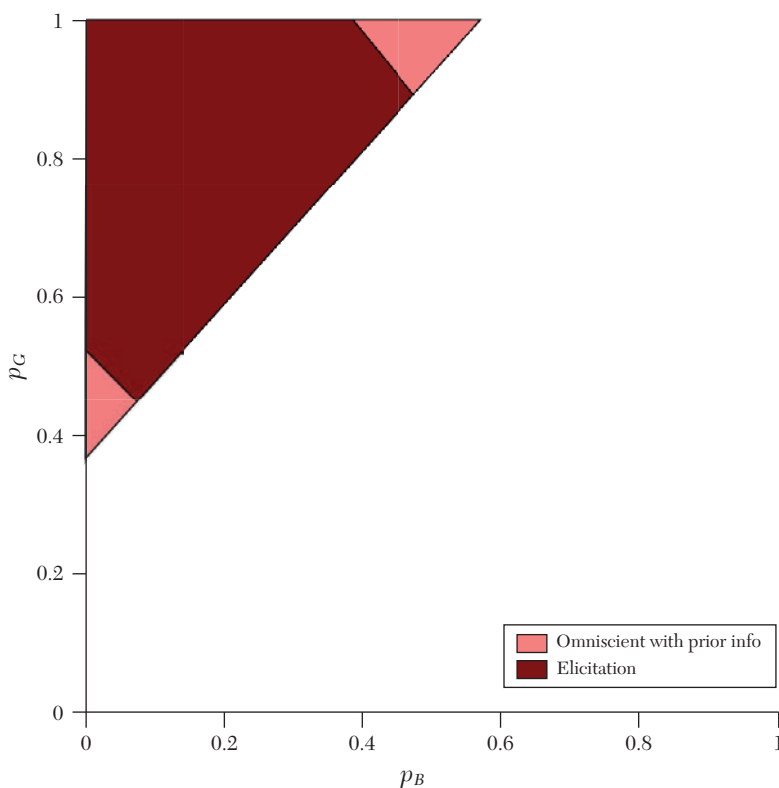


Figure 6. Investment Probability with Private Information

With these additional constraints, the set of outcomes that can arise in equilibrium under information design with elicitation is weakly, and typically strictly, smaller than under an omniscient designer, i.e., the case in the previous section where the designer can condition his recommendation directly on the players' prior information. The truth-telling constraints impose restrictions on how the differences in the conditional probabilities across types can vary across states. These impose additional restrictions on the ability of the government to attain either very low or very high investment probabilities in both states, as highlighted by equation (25).

Figure 6 illustrates the case where $x = 0.9$ and $q = 0.7$; the dark red region corresponds

to the outcomes that can arise under information design with elicitation; adding in the pink region, we get back to the triangle that corresponds to omniscient information design where the designer knows players' prior information.

6.2.2 Information Design without Elicitation

We could also consider a government, who *does not* know the signal of the firm and *cannot* even elicit it. Kolotilin et al. (2017) call this scenario "public persuasion." Such information design without elicitation has been the focus of the recent literature.

Clearly, the designer can replicate any decision rule without elicitation with a

decision rule with elicitation. This inclusion holds without any restrictions on the state space, the number of players, or the players' actions. In the specific investment example above, with a single player, two states, and two actions, the converse happens to be true as well. That is, any decision rule attainable with elicitation can also be attained without elicitation. In other words, in the binary setting and with a single player, there is no need for elicitation. The designer can induce any incentive-compatible decision rule by recommendations alone. We state and prove these two results in the appendix as propositions 5 and 6.¹⁶

The equivalence breaks down immediately if either of the binary assumptions regarding action and state are relaxed or we consider more than one player. We illustrate this failure of the equivalence result with a minor generalization of the investment example. In particular, in a single-player environment, we allow the player to either consider a small or a large investment. For completeness, we present, in the appendix, examples where one of the other two hypotheses fails. This modified example allows us to find a strict nesting of the set of outcomes without prior information, with prior information and an omniscient designer, with elicitation, and finally without elicitation. For the purpose of this example, it will be sufficient to focus on the case of a single player.

6.2.3 Beyond the Binary Setting

Consider the basic investment example with $I = 1$, $\Theta = \{B, G\}$, uniform prior, and symmetric types that are correct with probability $q > 1/2$. We now add an additional investment decision, to invest *small*, to the

set of feasible actions of the player. The decision to invest small comes with a higher rate of return but smaller total return than the (regular) investment decision. The payoff from a small investment is $-1/2$ in the bad state and $y \in (x/2, x)$ in the good state:

	bad state B	good state G
invest	-1	x
invest small	$-\frac{1}{2}$	y
not invest	0	0

For simplicity we restrict attention to decision rules that put zero probability on the small investment in equilibrium. We note that the small investment decision still plays a role in the characterization of incentive compatible decision rules as it is a feasible action to the player. It will hence generate additional obedience constraints that the designer has to respect as the player has now two possible deviations from the recommended action, one of which is to invest at a small scale. The decision rules—restricted to invest and not invest—can still be represented by a vector $p = (p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg})$ that records the probability of investing.

As a benchmark, first suppose the player has no prior information. Then a decision rule that *never* recommends the small investment can be represented as a pair $(p_B, p_G) \in [0, 1]^2$ that specifies the probability of the large investment in each state. When the firm has no prior information, there are two binding obedience constraints, one for the big investment against the small investment:

$$(26) \quad p_G x - p_B \geq p_G y - \frac{1}{2} p_B,$$

¹⁶Kolotilin et al. (2017) showed such an equivalence under a different set of assumptions. In their model, the designer and the player are privately informed about *distinct payoff states* rather than having *distinct information* about the same state as in the present setting.




Figure 7. Investment Probability under Different Information Design Scenarios

and one for no investment against the small investment:¹⁷

$$0 \geq (1 - p_C)y - \frac{1}{2}(1 - p_B).$$

The equilibrium regions are depicted in figure 7. If the firm has no prior information, the government faces only the above two constraints. The set of attainable decision rules is described by the light red area. In contrast to the setting with two investment

levels analyzed earlier, there is now a kink in the area of attainable decision rules that reflects a change in the binding obedience constraint, from zero investment to small investment.

If we consider the case in which the firm has prior information, then we have three different communication protocols for the government. An omniscient designer faces the obedience constraints that we analyzed earlier in section 3.2. By contrast, if the firm holds private information, then the information designer is no longer omniscient. Now the firm has two possible ways to disobey. If the government does not observe the signal but can elicit the information from the

¹⁷The other two possible incentive constraints, namely for the big investment against no investment and for no investment against the big investment, are supplanted by the above two.

firm, then we have truth-telling constraints as described by (23) and (24) in addition to the obedience constraints. Importantly, in this setting with more than two actions, the possibility of double deviations—misreporting and disobeying—generates additional constraints on the incentive compatible decision rules.

Finally, a designer without elicitation faces additional obedience constraints that rule out deviations conditional on a particular vector recommendation. The corresponding areas in figure 7 illustrate that the sequence of additional constraints from omniscient to elicitation to no elicitation impose increasingly more restrictions on the government and hence generate a sequence of strictly nested sets.

We already discussed how these three regimes offer an increasing number of constraints. It remains to discuss the specific impact of being able to elicit (or not) the private information. With elicitation, the player only learns the designer's recommendation for one type, namely the type that he reports. But a designer who cannot elicit must reveal her action recommendations for all types, hence the contingent recommendation. This enables a player to contemplate additional contingencies and hence deviations. With three possible actions, as in this example, there are two additional deviations that take advantage of this finer information. In particular, the high type can disobey the recommendation to invest by deviating to invest small only when the designer *also* recommends not to invest to the low type. Likewise, the low type can disobey the recommendation not to invest by deviating to invest small only when the designer *also* recommends investing to the high type. The additional options for the player induce further constraints on the information designer. Naturally, these additional deviations were not available in the binary action environment. And in fact, the absence of this large set of deviations accounts for the equivalence

between elicitation and no elicitation in the binary action and state environment.¹⁸

We conclude with a few observations about the comparative statics with respect to the information structure. As the precision of the information q decreases toward $1/2$, the inner three regions expand outward and converge to the no-prior-information equilibrium set. By contrast, as the precision q increases towards 1, the three inner regions contract and converge to the single coordinate vector $(0, 1)$.

6.3 *Information Design within Mechanism Design*

In the information design problem, the “information designer” can commit to providing information to the players to serve his ends, but has no ability to choose outcomes (or force the players to take particular actions). The set of available actions and a mapping from action profiles to outcomes and thus payoffs is fixed. How does this relate to “mechanism design”?

Myerson (1982) describes a class of *Bayes incentive problems*, which constitutes a leading definition of mechanism design (see also Myerson 1987, 1994). In this setting, players may have control over some actions affecting outcomes but the mechanism designer may be able to commit to pick other outcomes as a function of the players' reports. For example, in many classical mechanism design problems with individual rationality constraints, players do have control over some actions: participation versus nonparticipation. And even if the mechanism designer may not have any information that is unavailable to the players, he can—via the mechanism—implicitly control the

¹⁸We mentioned earlier that double deviations were not relevant in the binary environment in the sense that they do not add additional restrictions. This changes in the richer environment here, where the communicating designer indeed faces additional restrictions coming from the possibility of double deviations.

information that players have about each other. Myerson (1991) then labels the case where the mechanism designer has no direct control over outcomes “Bayesian games with communication” (section 6.3) and the setting where the designer has complete control over outcomes “Bayesian collective choice problems” (section 6.4). Thus, what we are calling information design corresponds to Myerson’s Bayesian games with communication with the proviso that the mediator brings his own information to the table, rather than merely redistributing others’ information.

There is also an important literature on an informed player (referred to as informed principal) who can commit to choose outcomes as a function of messages (see Myerson 1983). But in this setting, the information designer (principal) is typically assumed to be able to commit to a mechanism only *after* receiving his private information and the principal is not choosing action herself; see Mylovanov and Troeger (2012, 2014) and Perez-Richet (2014) for recent contributions. By contrast, in the information design setting there is a principal who cannot pick a contract/mechanism but can commit to a disclosure rule prior to observing her information.

6.4 *Correlated Equilibrium and Incomplete Information*

Aumann (1974, 1987) introduced correlated equilibrium as a solution concept for games with complete information (about the payoff matrix). He showed that the set of correlated equilibria equals the set of distributions over actions that could arise in a Bayes–Nash equilibrium if players observed some additional payoff-irrelevant signals (consistent with the common prior). Equivalently, the set of correlated equilibria corresponds to the set of outcomes that could be induced by an (uninformed) information designer. What we are calling “information design” thus corresponds to an incomplete information elaboration of this original

rationale for thinking about correlated equilibrium when the information designer has information of her own.¹⁹ In this section, we review the existing literature on incomplete information correlated equilibrium to relate it to the version of incomplete information correlated equilibrium—BCE—that is relevant for information design.

While Aumann (1987) provides an information design foundation for complete information correlated equilibrium, he offers a broader interpretation of the characterization, arguing that correlated equilibrium captures the implications of common knowledge of rationality in a complete information game, under the common prior assumption.²⁰ A large literature on the epistemic foundations of game theory has developed since then (Dekel and Siniscalchi 2015), elaborating on the formal language and questions suggested by Aumann’s work, although focused on the case of complete information without the common prior assumption. Formal treatment of the implications of common knowledge of rationality and the common prior assumption under incomplete information ties in with many of the issues discussed in this survey; we will discuss one issue that arises in this case in the next subsection.

To understand the literature on incomplete information correlated equilibrium, it is useful to identify two kinds of constraints in the literature on incomplete information correlated equilibrium: feasibility conditions

¹⁹Bergemann and Morris (2017) consider foundations for other solution concepts based on informational robustness and information design considerations, when the common prior assumption is not maintained.

²⁰Hillas, Kohlberg, and Pratt (2007) propose a related alternative foundation for correlated equilibrium. Consider an external observer who observes an infinite sequence of plays of a complete information game, has exchangeable beliefs about them, but does not believe he can offer beneficial advice to players on how to improve their payoffs. This observer must believe that play corresponds to a correlated equilibrium.

(constraints on what kind of information decision rules can condition on) and incentive compatibility conditions (what decision rules are consistent with optimal behavior). In this paper so far, we have introduced one feasibility condition—public feasibility (definition 3)—and three incentive constraints—obedience (definition 1), incentive compatibility (definition 2) and publicly feasible obedience (definition 4). Recall that BCE—our characterization of outcomes that can be induced by an omniscient information designer—imposed only obedience. We will discuss two further feasibility conditions to provide an overview of correlated equilibrium with incomplete information.

6.4.1 Belief Invariance

Consider the requirement that the information designer can correlate players' actions, but without changing players' beliefs and higher-order beliefs about the state of the world. This is formalized as follows.

DEFINITION 5 (Belief Invariant): Decision rule $\sigma : T \times \Theta \rightarrow \Delta(A)$ is belief invariant for (G, S) if $\sigma_i(a_i | (t_i, t_{-i}), \theta)$ is independent of t_{-i} for each i , a_i , t_i , and θ , where

$$\sigma_i(a_i | (t_i, t_{-i}), \theta) \triangleq \sum_{a_{-i} \in A_{-i}} \sigma_i((a_i, a_{-i}) | (t_i, t_{-i}), \theta).$$

We then say that a decision rule is a *belief-invariant BCE* if it satisfies belief invariance and obedience. It is not obvious how this feasibility condition arises under an information design interpretation: if the designer can condition his information on θ , why not allow him to change beliefs and higher-order beliefs?

There are two conceptual reasons why one might nonetheless be interested in belief-invariant BCE. First, Dekel, Fudenberg, and Morris (2007) introduced the solution concept of interim correlated rationalizability. They show that it characterizes the implica-

tions of common certainty of rationality and players' beliefs and higher-order beliefs. The solution concept by construction imposes belief invariance. Liu (2015) observes that the set of interim correlated rationalizable actions corresponds to the set of actions that can be played in a correlated equilibrium with incomplete information and subjective priors. If, then, the common prior assumption is imposed, this corresponds to the set of belief-invariant Bayes-correlated equilibria. Thus, the solution concept of belief-invariant BCE is the “right” one for understanding the implications of common knowledge assumptions under the common prior assumption.

Second, Mathevet, Perego, and Taneva (2017) consider a situation where the information designer can convey information only about beliefs and higher-order beliefs, but is not able to send additional information about correlation. Now the set of belief-invariant BCE, once some higher-order belief information has been sent, is equal to the set of BCE. Bergemann and Morris (2016a) describe how an arbitrary information structure can be decomposed into information about beliefs and higher-order beliefs, and additional belief-invariant signals.

6.4.2 Join Feasibility

Twenty five years ago, Forges (1993) (see also Forges 2006) gave an overview of incomplete information correlated equilibrium. A maintained assumption in that literature was that the information designer (or “mediator”) did not bring any information of her own to the table, but simply rearranged information, telling players privately about others' information. This can be formalized as follows.

DEFINITION 6 (Join Feasibility): Decision rule $\sigma : T \times \Theta \rightarrow A$ is join feasible for (G, S) if $\sigma(a | t, \theta)$ is independent of θ , i.e., $\sigma(a | t, \theta) = \sigma(a | t, \theta')$ for each $t \in T$, $a \in A$, and $\theta, \theta' \in \Theta$.

Thus, join feasibility allows the designer to use the *join* of the private information of all the players, the information contained in the entire type profile t . At the same time, it requires that the information designer can send information only about the type profile of the players and thus can only condition on the type profile, not on the state of the world θ or θ' . Join feasibility is imposed implicitly in some work on incomplete information correlated equilibrium—Forges (1993) integrates out uncertainty other than the players' types—but explicitly in others, e.g., Lehrer, Rosenberg, and Shmaya (2010).

As noted in the introduction, information design *adds* to the old incomplete information correlated literature the twist that the designer brings information of her own to the table. In turn, this allows the designer to choose the optimal design and provision of the information to the players.

Forges's 1993 paper was titled "Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information." The feasibility and incentive conditions described so far allow us to completely describe the five solution concepts she discusses:

- (i) A *Bayesian solution* is a decision rule satisfying join feasibility and obedience.
- (ii) A *belief invariant Bayesian solution* is a decision rule satisfying join feasibility, belief invariance, and obedience.
- (iii) An *agent normal form correlated equilibrium* is a decision rule satisfying join feasibility, public feasibility, (which implies belief invariance) and obedience.
- (iv) A *communication equilibrium* is a decision rule satisfying join feasibility and incentive compatibility (which implies obedience).

- (v) A *strategic form correlated equilibrium* is a decision rule satisfying join feasibility and publicly feasible obedience (which implies belief invariance, public feasibility, obedience, and incentive compatibility).

Thus, the Bayesian solution, communication equilibrium, and strategic form correlated equilibrium correspond to omniscient information design, information design with elicitation, and information design without elicitation, respectively. The belief invariant Bayesian solution and the agent normal form correlated equilibrium do not have natural information design interpretations.

Forges (1993) noted inclusions implied by these definitions. In particular, if we write (n) for the set of incomplete information correlated equilibria of type n above, we have

$$(5) \subseteq (3) \subseteq (2) \subseteq (1) \quad \text{and}$$

$$(5) \subseteq (4) \subseteq (1).$$

Forges (1993) reports examples showing that these inclusions are the only ones that can be shown, i.e., there exist decision rules that (i) are Bayesian solutions but not belief invariant BCE or communication equilibria; (ii) are belief invariant Bayes solutions but not communication equilibria or agent normal form correlated equilibria; (iii) are communication equilibria but not belief invariant Bayesian solutions; (iv) are belief invariant Bayesian solutions and communication equilibria but not agent normal form equilibria; (v) are agent normal form correlated equilibria but not communication equilibria; (vi) are agent normal form correlated equilibria and communication equilibria.

Forges (1993) discusses one more solution concept: the universal Bayesian solution. The universal Bayesian solution corresponds—in our language—to the set of Bayes-correlated

equilibria that would arise under joint feasibility if players had no information.

7. Information Design with Adversarial Equilibrium and Mechanism Selection

We have so far examined settings where the revelation principle holds: we can, without loss of generality, assume that the set of signals, or types, is equal to the set of actions. We now consider two natural extensions of information design where the revelation principle breaks down.

7.1 Adversarial Equilibrium Selection

In section 2, it was implicitly assumed that the information designer could, having designed the information structure, also select the equilibrium to be played. With one player, the equilibrium selection problem reduces to breaking ties and is not of substantive interest. However, Carroll (2016) and Mathevet, Perego, and Taneva (2017) highlighted that this issue is of first-order importance in the many-player case, and that the revelation principle argument breaks down and alternative arguments must be used; Mathevet, Perego, and Taneva (2017) formalize and analyze the case where players with no prior information will choose the worst equilibrium for the designer.

For our representation, we define a “communication rule” for the information designer. Players have the prior information encoded in the information structure $S = ((T_i)_{i=1}^I, \pi)$. The information designer sends each player i an extra message $m_i \in M_i$, according to rule $\phi: T \times \Theta \rightarrow \Delta(M)$, where $M = M_1 \times \dots \times M_I$. A communication rule is then $C = ((M_i)_{i=1}^I, \phi)$. Now the basic game G , the prior information structure S , and the communication rule C describe a Bayesian game (G, S, C) .²¹ A

strategy for player i in this game is a mapping $b_i: T_i \times M_i \rightarrow \Delta(A_i)$. A communication rule C and strategy profile b will now induce a decision rule

$$\sigma(a|t, \theta) = \sum_{m \in M} \phi(m|t, \theta) \left(\prod_i b_i(a_i|t_i, m_i) \right).$$

We will write $E(C)$ for the set of Bayes–Nash equilibria of the game with communication rule C . We can now give a more formal statement of proposition 1.

PROPOSITION 4: *Decision rule σ is a BCE of (G, S) if and only if there exists a communication rule C and a Bayes–Nash equilibrium $b \in E(C)$ that induce σ .*

This is a revelation principle argument that was formally stated as theorem 1 in Bergemann and Morris (2016a).

Recall that in section 2, we defined the information designer’s utility from BCE σ :

$$V(\sigma) = \sum_{a, t, \theta} v(a, \theta) \sigma(a|t, \theta) \pi(t|\theta) \psi(\theta).$$

We can also define the information designer’s utility from communication rule C and strategy profile b :

$$V^*(C, b) = \sum_{\theta, t, a} v(a, \theta) \left(\prod_i b_i(a_i|t_i, m_i) \sum_{m \in M} \phi(m|t, \theta) \right) \pi(t|\theta) \psi(\theta).$$

Let us consider the problem of an information designer who can pick both the communication rule and the equilibrium and is thus solving the problem

$$\max_C \max_{b \in E(C)} V^*(C, b).$$

Proposition 4 established that

$$\max_C \max_{b \in E(C)} V^*(C, b) = \max_{\sigma \in BCE} V(\sigma).$$

But one could also consider the problem of an information designer who can pick the communication rule, but wants to maximize

²¹ Bergemann and Morris (2016a) call the pair (S, C) an “expanded information structure.”

his utility in the worst equilibrium, and is thus solving the problem

$$\max_C \min_{b \in E(C)} V^*(C, b).^{22}$$

We now discuss three applications where maxmin information design problems have been motivated and studied; in each application, players have prior information. First, Carroll (2016) considers the problem of bilateral trade where he wants to know the worst possible gains from trade for a given distribution over the known private values of a buyer and a seller. If we picked the worst equilibrium, we could always support no trade with probability one, so instead he considers the best equilibrium. This is equivalent to having an information designer pick an information structure to minimize the efficiency of trade anticipating that the buyer and seller will play an equilibrium that maximizes efficiency (i.e., maximizes the gains from trade).

Second, Inostroza and Pavan (2017) consider global game models of regime change and the problem of an information designer trying to minimize the probability of regime change (they are motivated by the design of stress tests to minimize the probability of a run on a bank). What information should the information designer send—as a function of the state and players’ initial information—to minimize the probability of a run? They call this scenario “discriminatory” because the information designer can condition on players’ prior information. As in Carroll’s bilateral trade problem, the problem is not interesting if the designer is able to pick the equilibrium as well as the information structure: in this case, he can prevent the possibility

of inefficient outcomes by creating common knowledge of payoffs and picking the good equilibrium. To make the problem interesting, they then study the maxmin problem.

Finally, a literature on robustness to incomplete information (Kajii and Morris 1997) can be understood as an information design problem with adversarial equilibrium selection. We will give an example to illustrate this connection.²³ We will consider a slightly adapted version of the incomplete information investment game discussed earlier with payoffs:

$\theta = B$	invest	not invest
invest	x, x	$-1, 0$
not invest	$0, -1$	$0, 0$

$\theta = G$	invest	not invest
invest	x, x	$x, 0$
not invest	$0, -1$	$0, 0$

for some $0 < x < 1$, and the probability of state G is ε , and ε is small. Assume that the prior information is that player 1 knows the state and player 2 knows nothing. Thus, player 1 has a dominant strategy to invest in state G , while there are multiple equilibria in the complete information game corresponding to state B . In this setting, we can study the standard information design (with prior information) described above. Suppose that the information designer wants to maximize the probability that at least one player invests. Maintaining the assumption that the designer can pick the equilibrium, the answer is trivial: the information designer can simply give the players no additional

²²Of course, the minmax problem $\min_C \max_{b \in E(C)} V^*(C, b)$ is

a reinterpretation of a maxmin problem where the objective is replaced by $-V^*(C, b)$. The minmin problem is similarly a reinterpretation of the maxmax problem.

²³See also Hoshino (2017) on this connection.

information and there will be an equilibrium where players always invest.

But what if the information designer anticipated that the worst equilibrium would be played? This is information design with adversarial equilibrium selection. What information structure would the information designer choose and what would be the induced probability that at least one player invests? It is convenient and more transparent to describe information structures using the language of *partitions*.

Consider the information structure defined by a set $\Omega = \{1, 2, \dots, \infty\}$ where player 1 observes an element of the partition:

$$(\{1\}, \{2, 3\}, \{4, 5\}, \dots, \{\infty\}),$$

and player 2 observes an element of the partition:

$$(\{1, 2\}, \{3, 4\}, \dots, \{\infty\}).$$

Thus, an element of the partition now constitutes a signal realization. Let payoffs be given as defined by the payoff state $\theta = G$ if $1 \in \sigma$ is realized and by the payoff state $\theta = B$ everywhere else. For some $q \in (\frac{1}{2}, 1)$, let the probability of state $\omega \neq \infty$ be $\varepsilon \left(\frac{1-q}{q}\right)^\omega$ and so the probability of state ∞ is $1 - \frac{q}{2q-1}\varepsilon$ (if ε is sufficiently small). This information structure could arise from the prior information described above (only player 1 can distinguish between states B and G) and communicating additional information. Now suppose that $q > \frac{1}{1+x}$; this condition implies that a player assigning probability q to the other player investing will always have a strict incentive to invest. Following the induction argument of Rubinstein (1989), invest is the unique rationalizable action for both players at all states $\omega \neq \infty$. To see this, observe that at state 1, player 1 has a dominant strategy

to invest. Now player 2, with information set $\{1, 2\}$, must have a best response to invest, since he attaches probability q to player 1 investing. Now suppose that we have established that both players are investing at information sets of the form $\{\omega, \omega + 1\}$ if $\omega \leq k$. Now consider the player with information set $\{k + 1, k + 2\}$. He attaches probability q to the other player being at information set $\{k, k + 1\}$ and therefore investing. So the player with information set $\{k + 1, k + 2\}$ will invest. This argument establishes that it is possible to ensure that—if ε is sufficiently small—both players invest with probability $\frac{q}{2q-1}\varepsilon$. Since this is true for any $q > \frac{1}{1+x}$, it implies that it is possible to get both players to invest with probability arbitrarily close to

$$\frac{1/(1+x)}{2/(1+x)-1}\varepsilon = \frac{1}{1-x}\varepsilon.$$

The information structure we used to get arbitrarily close to this bound was (countably) infinite, but we can also get arbitrarily close using finite information structures as shown in Kajii and Morris (1997). Now, arguments from Kajii and Morris (1997) imply that this information structure is (arbitrarily close to) optimal for the information designer in this problem. To get a flavor of the argument, say that a player p -believes an event if he attaches probability at least p to the event occurring, and that there is common p -belief of that event if each player p -believes it, each player p -believes that both p -believe it, and so on. One can show that not invest is rationalizable only if there is common $\frac{x}{1+x}$ -belief that payoffs correspond to state B . But since $x < 1$, it follows that $\frac{x}{1+x} > \frac{1}{2}$ and one can show that if the event that payoffs are given by state B has probability at least $1 - \varepsilon$, then—for sufficiently small ε —the ex ante probability that there is common $\frac{x}{1+x}$ -belief that the state is B is

at least $1 - \frac{1}{1-x}\varepsilon$. This establishes that the bound is tight. If $x > 1$, similar arguments can be used to show that the information designer can ensure that both players invest with probability 1.

Arguments from Kajii and Morris (1997) and the follow-up literature (Ui 2001 and Morris and Ui 2005) can be used to analyze maxmin payoffs more generally when—as in the above example—the incomplete information game has each player either knowing that payoffs are given by a fixed complete information game or having a dominant strategy.

It is worth emphasizing that the above definition of the maxmin problem, and all three applications, correspond to the omniscient case where the information designer can condition on players' prior information as well as on the state. An alternative case that has been studied is when the information designer can only send public signals and only condition on the state (and not players' prior information). Goldstein and Huang (2016) and Inostroza and Pavan (2017) have studied this problem in global game models of regime change. (Inostroza and Pavan 2017 call this the nondiscriminatory case, to contrast with the discriminatory case described previously.) This case can be illustrated by our example above. An information designer interested in maximizing the probability of both investing would send a public signal to invest always if the state was good and with probability $\frac{\varepsilon}{1-\varepsilon}x$ if the state was bad. This would make player 2 indifferent between investing and not investing if he got the “invest” signal.

7.2 Adversarial Mechanism Selection

We considered an information designer who was choosing additional information for the players, holding fixed the basic game and players' prior information. But what if the information designer had to pick the information structure not knowing what the

basic game, or mechanism, was going to be? In particular, suppose that the choice of mechanism was adversarial. Again, we will lose the revelation principle. Once the information designer has picked the information structure (and thus the set of signals), the adversarial mechanism designer could pick a mechanism with a different set of messages.

Bergemann, Brooks, and Morris (2016) consider the problem of an information designer picking an information structure for a set of players with a common value of an object to *minimize* revenue, anticipating that an adversarial mechanism designer will then pick a mechanism to maximize revenue (a minmax problem). This gives an upper bound on the revenue of the seller of a single object who is picking a mechanism anticipating that the worst information structure will be chosen (a maxmin problem). Du (2018) constructs elegant bounds for the latter problem and shows that these bounds are sharp in the limit as the number of bidders increases. The former establishes the equivalence between minmax and maxmin exactly when there are two bidders and when the support of the value is binary, and the latter solves the auction design problem in the limit when the number of bidders goes to infinity. In a recent paper, Brooks and Du (2018) provide a general solution to the common value auction problem with a general common prior and a finite number of bidders. Both problems are studied *without* the common prior assumption by Chung and Ely (2007).

8. Conclusion

We have provided a unified perspective for a rapidly expanding literature on Bayesian persuasion and information design. In contrast with the recent literature on Bayesian persuasion that is concerned with a single player (receiver), we emphasized the implications of information design for

many-player strategic environments. We presented a two-step approach to information design: first identify the set of outcomes that are attainable under some information structure, then identify the optimal information structure. We have described the close connection between the information design problem and the earlier literature on correlated equilibrium with incomplete information; but whereas players are receiving real payoff-relevant information in the information design problem, in the older correlated equilibrium literature, the designer (mediator) was merely providing correlating devices.

We have focused on a pure version of the static information design problem where the designer has no ability to control outcomes. But—as argued in Myerson (1982) and Myerson (1987) and discussed in section 6.3—there are settings where a designer can control some outcomes (as a function of players' messages), but cannot control others and then can only use information to influence the outcomes outside her control. In other settings, the principal may be able to jointly choose the mechanism and the information structures. For example, Bergemann and Pesendorfer (2007) consider the optimal design of information structure and auction format in an independent private value environment. More recently, Daskalakis, Papadimitriou, and Tzamos (2016) solve for the optimal auction and information structure when the seller and the bidders each have some private information about the valuation of the object. Their analysis is motivated by online advertising auctions, where the two-way information asymmetry among seller and bidder is a central feature of the environment.

As one moves into dynamic settings, an overlap between the tools of information design and mechanism design more generally become more central. A specific setting where the tools of mechanism design and information design have recently been

studied in conjunction is the area of markets with resale. Here, the information that is disclosed in the first stage fundamentally affects the interaction in the resale market, see for example Calzolari and Pavan (2006); Dworzak (2017); Carroll and Segal (forthcoming); and Bergemann, Brooks, and Morris (2017b).

The information design problem—whether literal or metaphorical—identifies a mapping from the economic environment to possible outcomes, allowing for different choices of information structures. There is a second, reverse use of information design for robust identification, identifying a mapping from outcomes to possible parameters of the economic environment, allowing that different information structures might have generated the data.²⁴ For example, in an auction setting, one might consider a sample of bids from a sequence (or cross-section) of independent auctions. We can then ask what we can infer about the underlying distribution of valuations under weak assumptions on the information structure, that is without assuming a specific information structure. Syrgkanis, Tamer, and Ziani (2017) pursue such an approach for inference and identification in an auction setting. Magnolfi and Roncoroni (2017) adopt a similar perspective in the analysis of discrete games, in particular entry and exit games.

Many interesting avenues remain open in information design. There are many open methodological questions. The concavification approach has been very influential in the single-player (receiver) setting, so it is natural to ask if it can be as useful in the many-player setting. In both the linear programming and the concavification approach, the optimal information structure is identified by a global optimization problem. It might be insightful to find a more local approach that could

²⁴This reverse use is explicated in Bergemann and Morris (2013b).

identify the direction of valuable information provision. We briefly mentioned a number of applications of information design in the introduction. Digital information is becoming widely used in the allocation and distribution of services and commodities, as in traffic navigation apps such as Waze or Google Maps, recommender systems used by Netflix and

Amazon, or service platforms such as Uber or OpenTable. This suggests that information design will naturally be part of the solution of a large class of allocation problems. To the extent that the relevant information is arriving sequentially and improving over time, the resulting models will likely incorporate and address dynamic aspects.

APPENDIX

Additional Computation for Section 3.2

We observed in section 3.1 that, absent any additional information beyond the common prior, the firm does not invest. For any additional private information of the firm to change the “default” behavior, it has to be that the firm is investing after receiving the good signal, or that

$$(28) \quad qx + (1 - q)(-1) \geq 0 \Leftrightarrow x \geq \frac{1 - q}{q} \Leftrightarrow q \geq \frac{1}{1 + x}.$$

In other words, the information has to be sufficiently precise—thus q sufficiently large—to induce a change in the behavior.

Conditional on being type g , the firm will have an incentive to invest (when told to invest) under $p = (p_{Bb}, p_{Bg}, p_{Cb}, p_{Cg})$ if

$$(29) \quad -\frac{1}{2}(1 - q)p_{Bg} + \frac{1}{2}qp_{Cg}x \geq 0 \Leftrightarrow p_{Cg} \geq \frac{1 - q}{q} \frac{p_{Bg}}{x},$$

and an incentive to not invest (when told to not invest) if

$$(30) \quad 0 \geq -\frac{1}{2}(1 - q)(1 - p_{Bg}) + \frac{1}{2}q(1 - p_{Cg})x \Leftrightarrow p_{Cg} \geq \frac{1 - q}{q} \frac{p_{Bg}}{x} + 1 - \frac{1 - q}{q} \frac{1}{x}.$$

A similar pair of incentive constraints apply to the recommendations conditional on being type b .

As long as the private information of the firm is sufficiently noisy, or $q \leq 1/(1 + x)$, the binding constraint is (29) as in the uninformed case; otherwise it is the inequality (30) that determines the conditional probabilities. The obedience conditions for the firm observing a bad type b 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 $(p_{Bb}, p_{Bg}, p_{Cb}, p_{Cg})$ will induce behavior (p_B, p_G) integrating over types $t \in \{b, g\}$.

Additional Results for Section 6.2

For a given Bayesian game (G, S) , let $E(G, S)$, respectively, $NE(G, S)$, denote the set of decision rules that can be attained with and without elicitation, respectively.

PROPOSITION 5: *For each (G, S) , we have*

$$NE(G, S) \subseteq E(G, S).$$

PROOF:

Let $\sigma \in NE(G, S)$, and let ϕ be an obedient contingent recommendation that induces σ . To show that $\sigma \in E(G, S)$, we will verify that σ is incentive compatible.

Fix player i , types $t_i, t'_i \in T_i$, and a function $\delta_i: A_i \rightarrow A_i$. For each strategy $b_i \in B_i$, take $a'_i = \delta_i(b_i(t'_i))$ in player i 's (t_i, b_i) publicly feasible obedience constraint. Then sum the resulting inequalities over $b_i \in B_i$. After regrouping the summation, we have

$$\begin{aligned} & \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \left(\sum_{(b_i, b_{-i}) \in B} u_i((b_i(t_i), b_{-i}(t_{-i})), \theta) \phi((b_i, b_{-i}) | \theta) \right) \pi(t_i, t_{-i} | \theta) \psi(\theta) \\ & \geq \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \left(\sum_{(b_i, b_{-i}) \in B} u_i((\delta_i(b_i(t'_i)), b_{-i}(t_{-i})), \theta) \phi((b_i, b_{-i}) | \theta) \right) \pi(t_i, t_{-i} | \theta) \psi(\theta). \end{aligned}$$

We focus on the term in parentheses on each line. In the first line, group the summation according to the value of $(b_i(t_i), b_{-i}(t_{-i}))$ and use the fact that ϕ induces σ to obtain

$$\sum_{(a_i, a_{-i}) \in A} u_i((a_i, a_{-i}), \theta) \sigma((a_i, a_{-i}) | (t_i, t_{-i}), \theta).$$

In the second line, group the summation according to the value of $(b_i(t'_i), b_{-i}(t_{-i}))$ and use the fact that ϕ induces σ to obtain

$$\sum_{(a_i, a_{-i}) \in A} u_i((\delta_i(a_i), a_{-i}), \theta) \sigma((a_i, a_{-i}) | (t'_i, t_{-i}), \theta).$$

Substituting these expressions into the inequality gives player i 's type- t_i incentive constraint with deviation t'_i, δ_i . Since i, t_i, t'_i, δ_i are all arbitrary, the proof is complete. ■

PROPOSITION 6: *Let (G, S) be a Bayesian game with $I = 1$. If $|A| = 2$ and $|\Theta| = 2$, then*

$$NE(G, S) = E(G, S).$$

PROOF:

By proposition 5, it suffices to prove $NE(G, S) \supseteq E(G, S)$. First, we simplify the notation. Label the states and actions so that $\Theta = \{G, B\}$ and $A = \{0, 1\}$. If either action is weakly dominant, the desired result can be verified by directly computing $NE(G, S)$ and $E(G, S)$.

Therefore, we assume $u(1, G) - u(0, G)$ and $u(1, B) - u(0, B)$ are each nonzero and have opposite signs. Then without loss, we may assume the payoffs take the form $u(0, G) = u(0, B) = 0$, $u(1, B) = -1$, and $u(1, G) = x > 0$.²⁵ Action 1 can be interpreted as investment. We will represent a decision rule σ by a vector $p = (p_{\theta t})_{(\theta, t) \in \Theta \times T}$, where $p_{\theta t} = \sigma(1|\theta, t)$. For each signal $t \in T$, let

$$q(t) = \frac{\psi(G)\pi(t|G)}{\pi(t)},$$

where $\pi(t) = \psi(G)\pi(t|G) + \psi(B)\pi(t|B) > 0$ by assumption.

Let $p = (p_{\theta t}) \in E(G, S)$. To show that $p \in NE(G, S)$, we will explicitly construct an obedient contingent recommendation ϕ that induces p . Let $t, t' \in T$ and set $q = q(t)$ and $q' = q(t')$. Since p satisfies the truth-telling constraint,

$$\begin{aligned} qp_{Gt}x - (1 - q)p_{Bt} &\geq qp_{Gt'}x - (1 - q)p_{Bt'}, \\ q'p_{Gt'}x - (1 - q')p_{Bt'} &\geq q'p_{Gt}x - (1 - q')p_{Bt}. \end{aligned}$$

Taking $(1 - q', 1 - q)$ and (q', q) linear combinations of these two inequalities respectively yields

$$(q - q')(p_{Gt} - p_{Gt'})x \geq 0 \quad \text{and} \quad (q - q')(p_{Bt} - p_{Bt'}) \geq 0.$$

So $q(t) < q(t')$ implies $p_{\theta t} \leq p_{\theta t'}$ for $\theta \in \{G, B\}$. In the case $q = q'$, both inequalities must hold with equality so

$$q(p_{Gt} - p_{Gt'})x = (1 - q)(p_{Bt} - p_{Bt'}),$$

and hence $p_{Gt} \geq p_{Gt'}$ if and only if $p_{Bt} \geq p_{Bt'}$. Therefore, we can label the signals t_1, \dots, t_n so that

$$(31) \quad q(t_1) \leq \dots \leq q(t_n) \quad \text{and} \quad p_{\theta t_1} \leq \dots \leq p_{\theta t_n} \quad \text{for } \theta = B, G.$$

To simplify notation, define $q_l = q(t_l)$ for each $l = 1, \dots, n$; set $p_{\theta t_0} = 0$ and $p_{\theta t_{n+1}} = 1$ for all θ . For each $k = 1, \dots, n + 1$, define the cutoff strategy b^k by

$$b^k(t_l) = \begin{cases} 1 & \text{if } l \geq k \\ 0 & \text{otherwise.} \end{cases}$$

²⁵First, swap the labels G and B if needed to obtain $u(1, G) - u(0, G) > 0$. Then rescale the utility function so that $u(1, G) - u(0, G) = 1$. Finally, translate the functions $u(\cdot, G)$ and $u(\cdot, B)$ separately so that $u(0, G) = u(0, B) = 0$. The separate translations may change the agent's preferences over states but not over actions.

In particular, b^1 is unconditional investment and b^{n+1} is unconditional non-investment. Define the stochastic contingent recommendation $\phi: \Theta \rightarrow \Delta(B)$ by

$$\phi(b|\theta) = \begin{cases} p_{\theta t_k} - p_{\theta t_{k-1}} & \text{if } b = b^k \text{ for some } k = 1, \dots, n+1 \\ 0 & \text{otherwise.} \end{cases}$$

By (31), $p_{\theta t_k} - p_{\theta t_{k-1}} \geq 0$, so $\phi(\cdot|\theta)$ is a probability distribution for each $\theta \in \{G, B\}$. It is easy to check that ϕ induces the decision rule p .

To complete the proof, we verify that ϕ is obedient. For each $l = 1, \dots, n$ and $k = 1, \dots, n+1$, type t_l 's expected utility from investing if and only if being recommended b^k is:

$$U_{lk} = q_l(p_{Gt_k} - p_{Gt_{k-1}})x - (1 - q_l)(p_{Bt_k} - p_{Bt_{k-1}}).$$

Since both expressions in parentheses are nonnegative, U_{lk} is weakly increasing in l . Therefore, for types t_l with $l \geq k$,

$$U_{lk} \geq U_{k|k} = (q_k p_{Gt_k} x - (1 - q_k) p_{Bt_k}) - (q_k p_{Gt_{k-1}} x - (1 - q_k) p_{Bt_{k-1}}) \geq 0,$$

where the last inequality holds by truth telling for $k > 1$ and obedience for $k = 1$. Similarly, for types t_l with $l < k$,

$$U_{lk} \leq U_{k-1|k} = (q_{k-1} p_{Gt_k} x - (1 - q_{k-1}) p_{Bt_k}) - (q_{k-1} p_{Gt_{k-1}} x - (1 - q_{k-1}) p_{Bt_{k-1}}) \leq 0.$$

The last two inequalities establish the obedience of ϕ , so the proof is complete. ■

Now we return to the main problem of finding $(p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg})$. To compare these decision rules to the benchmark, we will ultimately integrate over the signals to compute the probability of investment in each state. Formally,

$$(p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg}) \rightarrow ((1 - q)p_{Bg} + qp_{Bb}, qp_{Gg} + (1 - q)p_{Gb}).$$

With an informed receiver, the omniscient designer faces four obedience constraints:

$$(32a) \quad \Delta_g^1 \triangleq qp_{Gg}x - (1 - q)p_{Bg} - \left(qp_{Gg}y - \frac{1 - q}{2}p_{Bg} \right) \geq 0,$$

$$(32b) \quad \Delta_b^1 \triangleq (1 - q)p_{Gb}x - qp_{Bb} - \left((1 - q)p_{Gb}y - \frac{q}{2}p_{Bb} \right) \geq 0,$$

$$(32c) \quad \Delta_g^0 \triangleq q(1 - p_{Gg})y - \frac{1 - q}{2}(1 - p_{Bg}) \leq 0,$$

$$(32d) \quad \Delta_b^0 \triangleq (1 - q)(1 - p_{Gb})y - \frac{q}{2}(1 - p_{Bb}) \leq 0.$$

More precisely, Δ_t^1 denotes the difference in type t 's utility between investing large and small when told to invest large and Δ_t^0 denotes the difference in type t 's utility between investing small and not investing when told not to invest. We thus have only ruled out profitable truthful deviations to the small investment, but it can be shown that this implies that there are no profitable truthful deviations to not invest or to invest large. An information designer with elicitation faces four additional constraints ruling out non-truthful deviations:

$$(33a) \quad qp_{Gg}x - (1-q)p_{Bg} \geq qp_{Gb}x - (1-q)p_{Bb},$$

$$(33b) \quad qp_{Gg}x - (1-q)p_{Bg} \geq q(p_{Gb}x + (1-p_{Gb})y) - (1-q)\frac{1+p_{Bb}}{2},$$

$$(33c) \quad (1-q)p_{Gb}x - qp_{Bb} \geq (1-q)p_{Gg}x - qp_{Bg},$$

$$(33d) \quad (1-q)p_{Gb}x - qp_{Bb} \geq (1-q)p_{Gg}y - \frac{q}{2}p_{Bg}.$$

Again, it is sufficient to consider a smaller class of deviations because the high type finds investment more attractive than the low type does. Formally, $E(G, S)$ is the set of $p \in [0, 1]^4$ satisfying (32a)–(33d).

Now we determine the additional constraints faced by an information designer without elicitation. Since there are only two signals, we may represent each strategy $b : T \rightarrow A$ as an ordered pair $(b(g), b(b)) \in A^2$. (In the second component, the letter b is used in two different ways, to denote a strategy and a signal.) The strategy $b = (\text{not invest}, \text{invest})$ can never be obedient for both types, so for any $p \in NE(G, S)$, there is only one candidate ϕ , namely

$$\phi((0, 0) | \theta) = 1 - p_{\theta g},$$

$$\phi((0, 1) | \theta) = 0,$$

$$\phi((1, 0) | \theta) = p_{\theta g} - p_{\theta b},$$

$$\phi((1, 1) | \theta) = p_{\theta b},$$

for each $\theta \in \{B, G\}$. A designer without elicitation faces two additional obedience constraints, which prevent deviations following the recommendation (invest, not invest):

$$(34a) \quad q(p_{Gg} - p_{Gb})x - (1-q)(p_{Bg} - p_{Bb}) \geq q(p_{Gg} - p_{Gb})y - \frac{1-q}{2}(p_{Bg} - p_{Bb}),$$

$$(34b) \quad (1-q)(p_{Gg} - p_{Gb})y - \frac{q}{2}(p_{Bg} - p_{Bb}) \leq 0.$$

Formally, $NE(G, S)$ is the set of decision rules in $E(G, S)$ satisfying (34a) and (34b).

After some algebra, we can see that (34a) is equivalent to

$$(35) \quad \Delta_g^1 \geq \Delta_b^1 + p_{Gb}(2q-1)(x-y) + p_{Bb}(2q-1)/2.$$

When p puts positive probability on investing after a bad signal, (35) eliminates decision rules for which (32a) has little slack. Similarly, (34b) is equivalent to

$$(36) \quad \Delta_b^0 \leq \Delta_g^0 - (1 - p_{Gg})(2q - 1)y - (1 - p_{Bg})(2q - 1)/2.$$

When p puts positive probability on not investing after a good signal, (36) eliminates decision rules for which (32d) does not have too much slack.

EXAMPLE 1 (Two Agents): Suppose $I = 2$, $\Theta = \{B, G\}$, $A_1 = A_2 = \{\text{invest, not invest}\}$, and $u_i(a_i, a_{-i}, \theta) = u(a_i, \theta)$ with u as in the opening example given by (5). Each player i receives a conditionally independent signal $t_i \in \{g, b\}$ that is correct with probability $q_i > 1/2$. Suppose $q_1 > q_2$, so that player 1 receives a more accurate signal. Consider the following decision rule: both players invest if player 1's signal is good and neither agent invests if player 1's signal is bad. For x sufficiently near one, this decision rule is incentive compatible. However, it is not even publicly feasible because following any contingent recommendation, player 2's choice of action will depend on her own signal, not on player 1's.

EXAMPLE 2 (Three States): Consider the single player, single investment setting of the opening example given by (5), but now split the bad state into two bad states B_1 and B_2 , each with prior probability $1/4$ and the same payoffs as in state B of the original example. Suppose the agents receive a completely uninformative binary signal t taking values t_1 and t_2 with equal probability. Consider the following decision rule: type t_i invests precisely in states G and B_i . For $x \in (1/2, 1)$, this decision rule is incentive compatible. It is uniquely induced by recommending $(b(t_1), b(t_2)) = (\text{invest, invest})$ in state G ; $(b(t_1), b(t_2)) = (\text{invest, not invest})$ in state B_1 ; and $(b(t_1), b(t_2)) = (\text{not invest, invest})$ in state B_2 . However, this contingent recommendation perfectly reveals the state of the world, so the agent can profitably deviate to his first-best strategy of investing if and only if the state is G . Therefore, the decision rule is not publicly feasible obedient.

REFERENCES

- Alonso, Ricardo, and Odilon Câmara. 2016. "Persuading Voters." *American Economic Review* 106 (11): 3590–605.
- Angeletos, George-Marios, and Jennifer La'O. 2013. "Sentiments." *Econometrica* 81 (2): 739–79.
- Arieli, Itai, and Yakov Babichenko. 2018. "Private Bayesian Persuasion." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2721307.
- Asquith, Paul, Thomas R. Covert, and Parag A. Pathak. 2013. "The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market." National Bureau of Economic Research Working Paper 19417.
- Aumann, Robert J. 1974. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics* 1 (1): 67–96.
- Aumann, Robert J. 1987. "Correlated Equilibrium as an Expression of Bayesian Rationality." *Econometrica* 55 (1): 1–18.
- Aumann, Robert J., and Michael B. Maschler. 1995. *Repeated Games with Incomplete Information*. Cambridge and London: MIT Press.
- Ball, Ian. 2018. "Dynamic Information Provision: Rewarding the Past and Guiding the Future." Unpublished.
- Bassan, Bruno, Olivier Gossner, Marco Scarsini, and Shmuel Zamir. 2003. "Positive Value of Information in Games." *International Journal of Game Theory* 32 (1): 17–31.
- Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin. 2018. "The Design and Price of Information." *American Economic Review* 108 (1): 1–48.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. "The Limits of Price Discrimination." *American Economic Review* 105 (3): 921–57.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2016. "Informationally Robust Optimal Auction Design." Cowles Foundation Discussion Paper

- 2065.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2017a. "First-Price Auctions with General Information Structures: Implications for Bidding and Revenue." *Econometrica* 85 (1): 107–43.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2017b. "Selling to Intermediaries: Optimal Auction Design in a Common Value Model." Cowles Foundation Discussion Paper 2064R.
- Bergemann, Dirk, Tibor Heumann, and Stephen Morris. 2015. "Information and Volatility." *Journal of Economic Theory* 158 (Part B): 427–65.
- Bergemann, Dirk, and Stephen Morris. 2013a. "Bayes Correlated Equilibrium and the Comparison of Information Structures." Cowles Foundation Discussion Paper 1822R.
- Bergemann, Dirk, and Stephen Morris. 2013b. "Robust Predictions in Games with Incomplete Information." *Econometrica* 81 (4): 1251–308.
- Bergemann, Dirk, and Stephen Morris. 2016a. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics* 11 (2): 487–522.
- Bergemann, Dirk, and Stephen Morris. 2016b. "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium." *American Economic Review* 106 (5): 586–91.
- Bergemann, Dirk, and Stephen Morris. 2017. "Belief-Free Rationalizability and Informational Robustness." *Games and Economic Behavior* 104: 744–59.
- Bergemann, Dirk, and Martin Pesendorfer. 2007. "Information Structures in Optimal Auctions." *Journal of Economic Theory* 137 (1): 580–609.
- Blackwell, David. 1951. "Comparison of Experiments." *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, edited by Jerzy Neyman, 93–102. Berkeley: University of California Press.
- Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics* 24 (2): 265–72.
- Brocas, Isabelle, and Juan D. Carrillo. 2007. "Influence through Ignorance." *RAND Journal of Economics* 38 (4): 931–47.
- Brooks, Benjamin, and Songzi Du. 2018. "Optimal Auction Design with Common Values: An Informationally-Robust Approach." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3137227.
- Calzolari, Giacomo, and Alessandro Pavan. 2006. "Monopoly with Resale." *RAND Journal of Economics* 37 (2): 362–75.
- Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review* 105 (7): 2183–203.
- Caplin, Andrew, and Daniel Martin. 2015. "A Testable Theory of Imperfect Perception." *Economic Journal* 125 (582): 184–202.
- Carroll, Gabriel. 2016. "Informationally Robust Trade and Limits to Contagion." *Journal of Economic Theory* 166: 334–61.
- Carroll, Gabriel, and Ilya Segal. Forthcoming. "Robustly Optimal Auctions with Unknown Resale Opportunities." *Review of Economic Studies*.
- Chung, Kim-Sau, and J. C. Ely. 2007. "Foundations of Dominant-Strategy Mechanisms." *Review of Economic Studies* 74 (2): 447–76.
- Das, Sammy, Emir Kamenica, and Renee Mirka. 2017. "Reducing Congestion through Information Design." In *Proceedings of the 2017 55th Annual Allerton Conference on Communication, Control, and Computing*, 1279–84. Piscataway, NJ: IEEE.
- Daskalakis, Constantinos, Christos Papadimitriou, and Christos Tzamos. 2016. "Does Information Revelation Improve Revenue?" In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 233–50. New York: Association for Computing Machinery.
- Dekel, Eddie, Drew Fudenberg, and Stephen Morris. 2007. "Interim Correlated Rationalizability." *Theoretical Economics* 2: 15–40.
- Dekel, Eddie, and Marciano Siniscalchi. 2015. "Epistemic Game Theory." In *Handbook of Game Theory with Economic Applications: Volume 4*, edited by H. Peyton Young and Shmuel Zamir, 619–702. Amsterdam and Lausanne: Elsevier, North-Holland.
- Doval, Laura, and Jeffrey C. Ely. 2016. "Sequential Information Design." Unpublished.
- Du, Songzi. 2018. "Robust Mechanisms under Common Valuation." *Econometrica* 86 (5): 1569–88.
- Duffie, Darrell, Piotr Dworczak, and Haoxiang Zhu. 2017. "Benchmarks in Search Markets." *Journal of Finance* 72 (5): 1983–2044.
- Dughmi, Shaddin. 2017. "Algorithmic Information Structure Design: A Survey." *ACM SIGecom Exchanges* 15 (2): 2–24.
- Dworczak, Piotr. 2017. "Mechanism Design with Aftermarkets: Cutoff Mechanisms." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2859206.
- Ely, Jeffrey C. 2017. "Beeps." *American Economic Review* 107 (1): 31–53.
- Ely, Jeffrey C., Alexander Frankel, and Emir Kamenica. 2015. "Suspense and Surprise." *Journal of Political Economy* 123 (1): 215–60.
- Ely, Jeffrey C., and Martin Szydlowski. 2017. "Moving the Goalposts." *Semantic Scholar*. <https://www.semanticscholar.org/paper/Moving-the-Goalposts-%E2%88%97-Ely-Szydlowski/c081b130ef941ba2ea338591a73e8480ccad3c03>.
- Forges, Françoise. 1993. "Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information." *Theory and Decision* 35 (3): 277–310.
- Forges, Françoise. 2006. "Correlated Equilibrium in Games with Incomplete Information Revisited." *Theory and Decision* 61 (4): 329–44.
- Forges, Françoise, and Frédéric Koessler. 2005. "Communication Equilibria with Partially Verifiable Types." *Journal of Mathematical Economics* 41 (7): 793–811.
- Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge and London: MIT Press.
- Gentzkow, Matthew, and Emir Kamenica. 2014.

- "Costly Persuasion." *American Economic Review* 104 (5): 457–62.
- Gentzkow, Matthew, and Emir Kamenica. 2017. "Competition in Persuasion." *Review of Economic Studies* 84: 300–322.
- Goldstein, Itay, and Chong Huang. 2016. "Bayesian Persuasion in Coordination Games." *American Economic Review* 106 (5): 592–96.
- Gossner, Olivier. 2000. "Comparison of Information Structures." *Games and Economic Behavior* 30 (1): 44–63.
- Hellwig, Christian, and Venky Venkateswaran. 2009. "Setting the Right Prices for the Wrong Reasons." *Journal of Monetary Economics* 56 (Supplement): S57–77.
- Hillas, John, Elon Kohlberg, and John Pratt. 2007. "Correlated Equilibrium and Nash Observer's Assessment of the Game." Harvard Business School Working Paper 08-005.
- Hirshleifer, Jack. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *American Economic Review* 61 (4): 561–74.
- Horner, Johannes, and Andrzej Skrzypacz. 2016. "Learning, Experimentation and Information Design." Unpublished.
- Hoshino, T. 2017. "Multi-agent Persuasion: Leveraging Strategic Uncertainty." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2978998.
- Inostroza, Nicolas, and Alessandro Pavan. 2017. "Persuasion in Global Games with Application to Stress Testing." Unpublished.
- Kajii, Atsushi, and Stephen Morris. 1997. "The Robustness of Equilibria to Incomplete Information." *Econometrica* 65 (6): 1283–309.
- Kamenica, Emir. Forthcoming. "Bayesian Persuasion and Information Design." *Annual Review of Economics*.
- Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (6): 2590–615.
- Kolotilin, Anton. 2018. "Optimal Information Disclosure: A Linear Programming Approach." *Theoretical Economics* 13 (2): 607–36.
- Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapecheinyuk, and Ming Li. 2017. "Persuasion of a Privately Informed Receiver." *Econometrica* 85 (6): 1949–64.
- Laffont, Jean-Jacques. 1989. *The Economics of Uncertainty and Information*. Cambridge and London: MIT Press.
- Lehrer, Ehud, Dinah Rosenberg, and Eran Shmaya. 2010. "Signaling and Mediation in Games with Common Interests." *Games and Economic Behavior* 68 (2): 670–82.
- Liu, Qingmin. 2015. "Correlation and Common Priors in Games with Incomplete Information." *Journal of Economic Theory* 157: 49–75.
- Lucas, Robert E., Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (2): 103–24.
- Magnolfi, Lorenzo, and Camilla Roncoroni. 2017. "Estimation of Discrete Games with Weak Assumptions on Information." Unpublished.
- Makris, Miltiadis, and Ludovic Renou. 2018. "Information Design in Multi-stage Games." Queen Mary, University of London School of Economics and Finance Working Paper 861.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. Oxford and New York: Oxford University Press.
- Maskin, Eric. 1999. "Nash Equilibrium and Welfare Optimality." *Review of Economic Studies* 66 (1): 23–38.
- Mathevet, Laurent, Jacopo Peregó, and Ina Taneva. 2017. "On Information Design in Games." *Semantic Scholar*. <https://www.semanticscholar.org/paper/On-Information-Design-in-Games-%E2%88%97-Mathevet-Peregó/de58caf63cb718d319855eea035468ad55e7b5b8>.
- Mertens, Jean-François, and Shmuel Zamir. 1985. "Formulation of Bayesian Analysis for Games with Incomplete Information." *International Journal of Game Theory* 14 (1): 1–29.
- Meyer, Margaret. 2017. "Information Design: Insights from Orderings of Dependence and Heterogeneity." Unpublished.
- Milgrom, Paul R., and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089–122.
- Morris, Stephen, and Takashi Ui. 2005. "Generalized Potentials and Robust Sets of Equilibria." *Journal of Economic Theory* 124 (1): 45–78.
- Myerson, Roger B. 1982. "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems." *Journal of Mathematical Economics* 10 (1): 67–81.
- Myerson, Roger B. 1983. "Mechanism Design By an Informed Principal." *Econometrica* 51 (6): 1767–97.
- Myerson, Roger B. 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323–58.
- Myerson, Roger B. 1987. "Bayesian Equilibrium and Incentive-Compatibility: An Introduction." In *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, edited by Leonid Hurwicz, David Schmeidler, and Hugo Sonnenschein, 229–59. Cambridge and New York: Cambridge University Press.
- Myerson, Roger B. 1991. *Game Theory: Analysis of Conflict*. Cambridge and London: Harvard University Press.
- Myerson, Roger B. 1994. "Communication, Correlated Equilibria and Incentive Compatibility." In *Handbook of Game Theory with Economic Applications: Volume 2*, edited by Robert Aumann and Sergiu Hart, 827–48. Amsterdam and Boston: Elsevier, North-Holland.
- Myerson, Roger B., and Mark A. Satterthwaite. 1983. "Efficient Mechanisms for Bilateral Trading." *Journal of Economic Theory* 29 (2): 265–81.
- Mylovanov, Tymofiy, and Thomas Tröger. 2012. "Informed-Principal Problems in Environments with Generalized Private Values." *Theoretical Economics* 7 (3): 465–88.

- Mylovanyov, Tymofiy, and Thomas Tröger. 2014. "Mechanism Design By an Informed Principal: Private Values with Transferable Utility." *Review of Economic Studies* 81 (4): 1668–707.
- Newman, Abraham. 1991. "The Positive Value of Information." *Games and Economic Behavior* 3 (3): 350–55.
- Novshek, William, and Hugo Sonnenschein. 1982. "Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release." *Bell Journal of Economics* 13 (1): 214–18.
- Ostrovsky, Michael, and Michael Schwarz. 2010. "Information Disclosure and Unraveling in Matching Markets." *American Economic Journal: Microeconomics* 2 (2): 34–63.
- Passadore, Juan, and Juan Pablo Xandri. 2014. "Robust Conditional Predictions in Dynamic Games: An Application to Sovereign Debt." *Semantic Scholar*. <https://www.semanticscholar.org/paper/Robust-Conditional-Predictions-in-Dynamic-Games-%3A-Passadore-Xandri/272ea4246dd7d4475838dd3aaa803962876f198b>.
- Perez-Richet, Eduardo. 2014. "Interim Bayesian Persuasion: First Steps." *American Economic Review* 104 (5): 469–74.
- Rayo, Luis, and Ilya Segal. 2010. "Optimal Information Disclosure." *Journal of Political Economy* 118 (5): 949–87.
- Roesler, Anne-Katrin, and Balázs Szentes. 2017. "Buyer-Optimal Learning and Monopoly Pricing." *American Economic Review* 107 (7): 2072–80.
- Rubinstein, Ariel. 1989. "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge'." *American Economic Review* 79 (3): 385–91.
- Syrkkanis, Vasilis, Elie Tamer, and Juba Ziani. 2017. "Inference on Auctions with Weak Assumptions on Information." <https://arxiv.org/abs/1710.03830>.
- Taneva, Ina. 2015. "Information Design." University of Edinburgh Discussion Paper 256.
- Ui, Takashi. 2001. "Robust Equilibria of Potential Games." *Econometrica* 69 (5): 1373–80.
- Venkateswaran, Venky. 2013. "Heterogeneous Information and Labor Market Fluctuations." Unpublished.