Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Cowles Foundation Discussion Papers | Cowles Foundation

2-1-2017

Information Design: A Unified Perspective

Dirk Bergemann

Stephen Morris

Follow this and additional works at: https://elischolar.library.yale.edu/cowles-discussion-paper-series

 Part of the Economics Commons

Recommended Citation

Bergemann, Dirk and Morris, Stephen, "Information Design: A Unified Perspective" (2017). *Cowles Foundation Discussion Papers*. 2544.
https://elischolar.library.yale.edu/cowles-discussion-paper-series/2544

This Discussion Paper is brought to you for free and open access by the Cowles Foundation at EliScholar – A Digital Platform for Scholarly Publishing at Yale. It has been accepted for inclusion in Cowles Foundation Discussion Papers by an authorized administrator of EliScholar – A Digital Platform for Scholarly Publishing at Yale. For more information, please contact elischolar@yale.edu.

INFORMATION DESIGN: A UNIFIED PERSPECTIVE

By

Dirk Bergemann and Stephen Morris

February 2017
Revised November 2017

COWLES FOUNDATION DISCUSSION PAPER NO. 2075R2



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

http://cowles.yale.edu/

Information Design: A Unified Perspective*

Dirk Bergemann† Stephen Morris‡

November 27, 2017

Abstract

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 into 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.

KEYWORDS: Information design, Bayesian persuasion, correlated equilibrium, incomplete information, robust predictions, information structure.

JEL CLASSIFICATION: C72, D82, D83.

*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 are thankful for the extensive comments from the Editor, Steven Durlauf, and five anonymous referees. We have benefited from discussions and correspondence about the material in this paper with Ben Brooks, Laura Doval, Jeff Ely, Francoise 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 Musolff, Denis Shishkin, Áron Tóbiás and Xinyang Wang.

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

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

Contents

1 Introduction **4**

2 The Information Design Problem **9**

3 An Investment Example **13**

3.1 Single Player without Prior Information 14

3.2 Single Player with Prior Information 15

3.3 Many Players without Prior Information 17

3.4 Many Players with Prior Information 19

4 Issues in Information Design Illustrated by the Examples **20**

4.1 Public versus Private Signals; and Instrumental versus Intrinsic Motivation for Preferences Over Correlation 21

4.2 Tightening Obedience Constraints and Bayes Correlated Equilibrium Outcomes 23

4.3 Bayes Correlated Equilibrium Outcomes and Information Design without Concavification 26

5 Metaphorical Information Design and Applications **27**

5.1 The Limits of Price Discrimination 28

5.2 Information and Volatility 31

6 Information Design with Private Information **34**

6.1 Players' Prior Information is not Known to the Designer 34

6.2 The Investment Example Re-Visited 37

6.2.1 Information Design with Elicitation 37

6.2.2 Information Design without Elicitation 39

6.2.3 Beyond the Binary Setting 40

6.3 Information Design and the Relation to Mechanism Design 42

6.4 Correlated Equilibrium and Incomplete Information 43

6.4.1 Belief Invariance 44

6.4.2 Join Feasibility 45

7 Information Design with Adversarial Equilibrium and Mechanism Selection **47**

7.1 Adversarial Equilibrium Selection 47
7.2 Adversarial Mechanism Selection 51

8 Conclusion **52**

9 Appendix **54**

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 the game itself, or force the players to choose particular actions.[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 auction design (Daskalakis, Papadimitriou, and Tzamos (2016)), grade disclosure and matching markets (Ostrovsky and Schwarz (2010)), voter mobilization (Alonso and Camara (2016)), traffic routing (Das, Kamenica, and Mirka (2017)), rating systems (Duffie, Dworczak, and Zhu (2017)) and transparency regulation (Asquith, Covert, and Pathak (2014)) 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 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") 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

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

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.

A third purpose is to contrast –and at the same time highlight the parallels between–the "information design" problem with the basic "mechanism design" problem. A "mechanism designer" can commit to a mechanism for the players to serve her ends, but has no ability to provide the players with any additional information (or force the players to make particular action choices).[2] An "information designer" can commit to provide information about the states to the players to serve her ends, but has no ability to change the game of the players (or force the players to make particular action choices). In either information design or mechanism design, the problem is sometimes studied with a restricted choice set for the designer. In the information design problem, we could restrict the designer to *only* choose whether the players are given no information or complete information about the environment. In the mechanism design problem, we could restrict the designer to *only* choose between a first price or a second price auction. However, in each design problem, there is a revelation principle argument that allows for the analysis of all information structures or all mechanisms respectively. For the mechanism design problem, we can restrict attention to direct mechanisms where the players' action sets are equal to their type sets. Conversely, for the information design problem, we can restrict attention to information structures where the players' type sets are equal to their action sets. In this paper, using this observation, we consider static information design problems when all information structures are available to the designer.

To survey the literature, and to provide some graphical illustrations, we use a family of two player, two action, and two state examples. We start with the leading example of Bayesian persuasion (with

[2] This statement of the mechanism design problem is a narrow one corresponding to what Myerson (1991) (Section 6.4) calls "Bayesian collective choice problems." We discuss how what we are calling information design fits into the broader mechanism design literature in Section 6.3.

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 private information of the agents (types) into possibly stochastic signals to the players.[3] As the mapping essentially recommends actions to the players, we refer to it as

[3]Whether 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' private signals. For example, a prosecutor might never know whether the defendant is guilty or innocent, but can nevertheless set up an

a *decision rule*. We will initially focus on what we will sometimes call the *omniscient* case, where 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

investigation process which 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.

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

(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 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 (2016), Doval and Ely (2016), Ely (2017), Ely and Szydlowski (2017), Ball (2017) and Makris and Renou (2017). 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) surveys work on information design in dynamic settings. Kamenica and Gentzkow (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.

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 the broad overview of mechanism design proposed by Myerson (1982) and with 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, a mediator, a principal and a 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, $1, 2, ..., I$, and we write i for a typical player. We write Θ for the payoff states of the world and θ for a typical element of 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

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

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

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

An "information structure" S consists of (1) for each player i, a finite set of types $t_i \in T_i$; and (2) a type distribution $\pi : \Theta \to \Delta(T)$, where we write $T = T_1 \times \cdots \times T_I$. Thus $S = \left((T_i)_{i=1}^I, \pi\right)$.

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 a 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.

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

1. the information designer picks and commits to a rule for providing the players with extra messages;
2. the true state θ is realized;
3. each agent's type t_i is privately realized;

4. the players receive extra messages according to the information designer's rule;

5. the players pick their actions based on their prior information and the messages provided by the information designer;

6. 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 signalling environments where the informed party commits to 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*

$$\sigma : T \times \Theta \to \Delta(A). \tag{1}$$

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 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 to the desired precision and accuracy.

The decision rule encodes the information that players may receive about the realized state of the world and the types of the other players. The dependence of the action profile a on these is created by the recommended action that the information designer conveys 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 recommendation action a_i rather than choose any other available action a_i'.

Definition 1 (Obedience)

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

$$\begin{aligned} & \sum_{a_{-i} \in A_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} u_i\left(\left(a_i, a_{-i}\right), \theta\right) \sigma\left(\left(a_i, a_{-i}\right) \mid \left(t_i, t_{-i}\right), \theta\right) \pi\left(\left(t_i, t_{-i}\right) \mid \theta\right) \psi(\theta) \\ \geq & \sum_{a_{-i} \in A_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} u_i\left(\left(a_i', a_{-i}\right), \theta\right) \sigma\left(\left(a_i, a_{-i}\right) \mid \left(t_i, t_{-i}\right), \theta\right) \pi\left(\left(t_i, t_{-i}\right) \mid \theta\right) \psi(\theta), \end{aligned} \tag{2}$$

for all $a_i' \in A_i$.

The obedience inequality requires that each player i, after receiving his recommendation a_i, assesses that no other action a_i' could yield him a strictly higher utility. While the above inequality is written from an ex ante perspective, we emphasize that each player when computing his expected utility is indeed using the information contained in the action recommendation a_i. This becomes apparent when we 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' rule explicitly:

$$\begin{aligned} & \sum_{a_{-i} \in A_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} u_i\left(\left(a_i, a_{-i}\right), \theta\right) \frac{\sigma\left(\left(a_i, a_{-i}\right) \mid \left(t_i, t_{-i}\right), \theta\right) \pi\left(\left(t_i, t_{-i}\right) \mid \theta\right) \psi(\theta)}{\sum_{a_{-i}'' \in A_{-i}, t_{-i}'' \in T_{-i}, \theta'' \in \Theta} \sigma\left(\left(a_i, a_{-i}''\right) \mid \left(t_i, t_{-i}''\right), \theta''\right) \pi\left(\left(t_i, t_{-i}''\right) \mid \theta''\right) \psi\left(\theta''\right)} \\ \geq & \\ & \sum_{a_{-i} \in A_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} u_i\left(\left(a_i', a_{-i}\right), \theta\right) \frac{\sigma\left(\left(a_i, a_{-i}\right) \mid \left(t_i, t_{-i}\right), \theta\right) \pi\left(\left(t_i, t_{-i}\right) \mid \theta\right) \psi(\theta)}{\sum_{a_{-i}'' \in A_{-i}, t_{-i}'' \in T_{-i}, \theta'' \in \Theta} \sigma\left(\left(a_i, a_{-i}''\right) \mid \left(t_i, t_{-i}''\right), \theta''\right) \pi\left(\left(t_i, t_{-i}''\right) \mid \theta''\right) \psi\left(\theta''\right)}. \end{aligned}$$

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 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 provided that the denominator is strictly positive we can multiply through to obtain the earlier inequality (2).

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 *An omniscient information designer can attain decision rule* σ *if and only if it is a Bayes correlated equilibrium, 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.[5] 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 private information t:

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

In this (and later) propositions, we omit formal statements and proofs that correspond to revelation principle arguments. Bergemann and Morris (2016a) give a formal statement[6] 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, that is if the set Θ is a singleton, then the notion of Bayes correlated equilibrium exactly coincides with the complete information correlated equilibrium as introduced in the seminal contributions of Aumann (1974) and (1987). In the absence of payoff uncertainty, we can simply suppress the dependence of the payoff function on θ as θ does not vary. Moreover 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 $\sigma \in \Delta\left(A\right)$. Now, a distribution $\sigma \in \Delta\left(A\right)$ is defined to be a correlated equilibrium if for each i and $a_i \in A_i$, we have:

$$\sum_{a_{-i} \in A_{-i}} u_i\left(a_i, a_{-i}\right) \sigma\left(\left(a_i, a_{-i}\right)\right) \geq \sum_{a_{-i} \in A_{-i}} u_i\left(a_i', a_{-i}\right) \sigma\left(a_i, a_{-i}\right). \tag{3}$$

The above obedience condition (2) thus collapses to the best response property (3) in the absence of 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); this is an interpretation of the statement that correlated equilibrium captures the set of outcomes attainable by an information designer in the absence of payoff. 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)

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

[6] A formal statement also appears in Section 7.

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 design 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, her ex ante utility of the information designer from decision rule σ for a given game of incomplete information (G, S) is:

$$V(\sigma) = \sum_{a \in A, t \in T, \theta \in \Theta} v(a, \theta) \sigma(a|t, \theta) \pi(t|\theta) \psi(\theta). \tag{4}$$

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. 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:

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

. (5)

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. 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:

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

where we can omit the type space T. 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:

$\sigma(a\,\vert\theta)\,\psi(\theta)$	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' 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:

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

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 ex-ante probabilities, which we shall always do from now on:

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

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 is 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 $v(a, \theta)$.

INSERT FIGURE 1: INVESTMENT PROBABILITY WITH UNINFORMED PLAYER: $x = 55/100$

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:

$\sigma(a\|\theta)\psi(\theta)$	bad state B	good state G
invest	$\frac{1}{2}x$	$\frac{1}{2}$
not invest	$\frac{1}{2}(1-x)$	0

. (8)

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 from 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 of the government.[7] In particular,

[7] Some detailed calculations for this example appear in the Appendix.

the firm has a type (or receives a signal) which 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:

$\pi(t\,\vert\theta)$	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 vector, a quadruple:

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

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:

$$(1-q)\,p_{Bg}\,(-1) + qp_{Gg}x \geq 0. \tag{10}$$

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

$$p_G = qp_{Gg} + (1-q)\,p_{Gb}.$$

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

$$p_G \geq q - \frac{1-q}{x}, \tag{11}$$

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 private information by the agent does not constrain the government in its recommendation policy. In consequence, private 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 as 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' law.

INSERT FIGURE 2: INVESTMENT PROBABILITY WITH INFORMED PLAYER: $x = 55/100$

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.[8] 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	
	$\theta = B$	invest	not invest
firm 1	invest	$-1+\varepsilon$	-1
	not invest	0	0

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

(12)

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

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

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

$\theta = B$	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	$p_G - r_G$	$1 + r_G - 2p_G$

. (13)

To ensure that all probabilities are non-negative, 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 invest if

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

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

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

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

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

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 $x > \frac{1}{2}$ and restrict attention to $|\varepsilon| \leq x - \frac{1}{2}$. 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.

Insert Figure 3: Investment Probability with Negative or Positive Strategic Term ε, $x = 55/100$.

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_G) unlike in the single player case depicted earlier in Figure 2.

Insert Figure 4: Investment Probability with Two Players with Prior Information, with Strategic Term $\varepsilon = 3/10, x = 55/100$.

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.

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 generated the insight that the designer would like to generate public signals when there are strategic complementarities and to generate 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 non-strategic 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 (non-cooperative 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.[9] 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.[10]

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 for intrinsic but not instrumental reasons. We can illustrate this case with the example

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

[10] In 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.

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 $x > \frac{1}{2}$ - it is possible to ensure that one firm always enters. 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 $\frac{2}{3}\left(x - \frac{1}{2}\right) > 0$. 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.[11]

4.2 Tightening Obedience Constraints and Bayes Correlated Equilibrium 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.

[11] In 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.

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 adding additional signals reduces the set of outcomes that can occur; it is 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, 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.[12]

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

[12] Blackwell (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.

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 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 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 distribution 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, Gossner, Scarsini, and Zamir (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 focussing on common interest games. Here, more information in the sense of individual sufficiency translates into more attainable outcomes. But looking at Bayes correlated equilibria 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 Bayes Correlated Equilibrium 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 Bayes correlated equilibria (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 sub-sections 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 state, 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 set of posteriors over the states of the world, subject to the constraint that the prior over states is a convex combination of those posteriors. 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 matter, 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 on (one player) price discrimination discussed in the next section, Bergemann, Brooks, and Morris (2015), 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 (2017), 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 simplest 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.[13] 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 Bayes correlated equilibria is precisely the set of outcomes that can arise with extra information, for a given basic game and prior information structure. If there are properties that hold for all Bayes correlated equilibria, we have identified predictions that are robust to the exact information structure. Identifying the best or worst outcome under some criterion that can arise under some information structure 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 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 segmen-

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

tation. First, consumer surplus must be non-negative 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 always to offer the product with probability one at a given price to all buyers. We therefore refer to it as *uniform* monopoly price, and correspondingly uniform monopoly profit. Third, the sum of consumer and producer surplus cannot exceed the total social value that consumers receive from the good, value minus unit cost of production. The shaded right angled triangle in Figure 5 illustrates these three bounds.

Insert Figure 5: The Bounds on Profits and Consumer Surplus in Third Degree Price Discrimination

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 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 for some information designer preferences over producer and consumer surplus.

The information design has a very clear literal interpretation in one case. If the monopolist knew the consumer's valuation, he can 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 $u^* = (1/3)(3 - 2) + (1/3)(2 - 2) = 1/3$. 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 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 $\sigma(x)$:

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

(16)

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 (2017) pursue the question of how private

information may impacts 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 macroeconomics 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 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 in 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 identically and inde-

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

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

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 of any two agents i and j, θ_i and θ_j is:

$$\rho_\theta \triangleq \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}. \tag{18}$$

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

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

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 complementarities 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:

$$s_i \triangleq \lambda\varepsilon_i + (1 - \lambda)\theta, \tag{19}$$

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:

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

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, their actions are correlated by means of the underlying correlation of the individual payoff shocks, and the 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 towards 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:

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

is achieved by the information structure λ^* :

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

which biases the signal towards the aggregate shock. We can express the information structure that maximizes the aggregate volatility in terms of the correlation coefficient ρ_θ:

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

and the maximal volatility can be expressed as:

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

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 0, 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 classical 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 $\sigma\colon T \times \Theta \to \Delta(A)$ but each type of each player can choose a deviation $\delta_i : A_i \to A_i$ 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 $\sigma\colon T \times \Theta \to \Delta(A)$ *is* incentive compatible *for* (G, S) *if for each* $i = 1, \ldots, I$ *and* $t_i \in T_i$,

$$\sum_{(a_i, a_{-i}) \in A,\, t_{-i} \in T_{-i},\, \theta \in \Theta} u_i((a_i, a_{-i}), \theta)\sigma((a_i, a_{-i})|(t_i, t_{-i}), \theta)\pi((t_i, t_{-i})|\theta)\psi(\theta)$$
$$\geq \sum_{(a_i, a_{-i}) \in A,\, t_{-i} \in T_{-i},\, \theta \in \Theta} u_i((\delta_i(a_i), a_{-i}), \theta)\sigma((a_i, a_{-i})|(t_i', t_{-i}), \theta)\pi((t_i, t_{-i})|\theta)\psi(\theta)$$

for all $t_i' \in T_i$ *and* $\delta_i\colon A_i \to A_i$.

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 truthtelling *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, that is 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

$$b_i : T_i \to A_i.$$

We define $B = \times_{i=1}^{I} B_i$ and let a generic element be given by $b = (b_1, \ldots, b_I) \in B$.

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

Definition 3 (Public Feasibility)

A decision rule $\sigma\colon T \times \Theta \to \Delta(A)$ *is* publicly feasible *if there exists a contingent recommendation*

$\phi\colon \Theta \to \Delta(B)$ *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 players'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 player, hence only the publicly available information about the player, namely the set of possible types and their common prior distribution; and it has to be feasible in the sense that it induces the decision rule σ.[14] 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 $b(t)$ of the strategy for different types 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 $\sigma\colon T \times \Theta \to \Delta(A)$ *is* publicly feasible obedient *if there exists a contingent recommendation* $\phi\colon \Theta \to \Delta(B)$ *such that* (*i*) ϕ *induces* σ, *and* (*ii*) ϕ *satisfies* obedience *in the sense that for*

[14] Kolotilin, Li, Mylovanov, and Zapechelnyuk (2017) offer a related interpretation when they distinguish between "public" versus "private" persuasion. If there is continuum population of ex ante identical players (without strategic interaction) and the types then correspond to individuals in the population, "public" persuasion refers to the setting when all types receive the *same*, hence public recommendation, one that is again typically a vector.

each $i = 1, \ldots, I$, $t_i \in T_i$, *and* $b_i \in B_i$,

$$\sum_{b_{-i} \in B_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} u_i((b_i(t_i), b_{-i}(t_{-i})), \theta) \phi((b_i, b_{-i})|\theta) \pi((t_i, t_{-i})|\theta) \psi(\theta)$$
$$\geq \sum_{b_{-i} \in B_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} u_i((a_i', b_{-i}(t_{-i})), \theta) \phi((b_i, b_{-i})|\theta) \pi((t_i, t_{-i})|\theta) \psi(\theta)$$

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 towards 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 Re-Visited

We reconsider the investment example introduced earlier in Section 3 but now allow the firm's information 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 which 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, Li, Mylovanov, and Zapechelnyuk (2017) refer to this informational environment

as "private persuasion".[15] Kolotilin (2017) 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 truthtelling constraint for the good type $t = g$ is:

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

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

$$(1-q)\,p_{Gb}x - qp_{Bb} \geq (1-q)\,p_{Gg}x - qp_{Bg}. \tag{24}$$

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

$$\frac{1-q}{q}x\,(p_{Gg} - p_{Gb}) \leq (p_{Bg} - p_{Bb}) \leq \frac{q}{1-q}x\,(p_{Gg} - p_{Gb})\,. \tag{25}$$

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

$$(p_{Gg} - p_{Gb}) \geq 0, (p_{Bg} - p_{Bb}) \geq 0,$$

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

[15] Bergemann, Bonatti, and Smolin (2017) 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 mult-item monopolist.

double deviations do not impose any additional restriction on the behavior of the player. We state and prove this results formally in the appendix as Proposition 6.

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 knows the players' prior information. The truthtelling 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.

Insert Figure 6: Investment Probability with Private 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, Li, Mylovanov, and Zapechelnyuk (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, the information designer can attain any decision rule with elicitation with one that does not use 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 Proposition 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

[16] Kolotilin, Li, Mylovanov, and Zapechelnyuk (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.

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 examples where one of the other two hypotheses fails in the appendix. 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 $p_{\theta t} = (p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg})$ 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:

$$p_G x - p_B \geq p_G y - \frac{1}{2} p_B, \tag{26}$$

and one for the no investment against the small investment:[17]

$$0 \geq (1 - p_G)y - \frac{1}{2}(1 - p_B)\,. \tag{27}$$

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 rule 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 firm then we have truthtelling 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 imposing 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

[17] The other two possible incentive constraints, namely for the big investment recommendation not to invest all, and for no investment recommendation to invest big are supplanted by the above two.

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 towards $1/2$, the inner three regions expand outwards 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 singleton $(0, 1)$.

Insert Figure 7: Investment Probability under Different Information Design Scenarios.

6.3 Information Design and the Relation to 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 change the "mechanism" (or force the players to take particular action choices). A mechanism here describes the set of players, their available actions and a mapping from action profiles to outcomes. Contrast this "information design" problem with the basic "mechanism design" problem, where a "mechanism designer" can commit to a mechanism for the players to serve his ends, but has no ability to provide the players with any additional information (or force the players to make particular action choices).

This narrow definition of mechanism design assumes that players have no control over their actions and that the mechanism designer has no information of her own to provide to the players. Myerson (1982) describes a class of *Bayes incentive problems*, which constitutes a very broad definition of mechanism design (see also Myerson (1987) and Myerson (1994)). In this broad vision, 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 non-participation. 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 information that players have about each other. Myerson (1991) then labels the case where the mechanism designer has

[18] We 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.

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 re-distributing others' information.

There is also an important literature on informed players who can commit to choosing outcomes as a function of messages, where they are called informed principals (Myerson (1983)). But in this setting, the information designer (principal) is typically assumed to be able to commit to a mechanism but can do so only after receiving his information; players do not interact strategically; and the principal is choosing the contract/mechanism; see Mylovanov and Troeger (2012), (2014) and Perez-Richet (2014) for recent contributions. By contrast, we have a principal who cannot pick a contract/mechanism but can commit to a disclosure rule prior to observing her information, in an environment with multiple players interacting strategically.

6.4 Correlated Equilibrium and Incomplete Information

Aumann (1987) introduced correlated equilibrium as a solution concept for game of complete information. He showed that the set of correlated equilibria equalled 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 - Bayes correlated equilibrium - 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

[19] Bergemann 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.

[20] Hillas, 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

theory has developed since then (Dekel and Siniscalchi (2014)), elaborating on the formal language and questions suggested by Aumann's work, although focussed 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 (constraints on what kind of information decision rules can condition on) and incentive compatibility conditions (what decision rules are consistent with optimal behavior). In the 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 Bayes correlated equilibrium - 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:

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

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

for each $t_{-i} \in T_{-i}$*.*

We then say that a decision rule is a *belief invariant Bayes correlated equilibrium* 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 couple of 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 implications of common certainty of rationality

must believe that play corresponds to a correlated equilibrium.

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 re-arranged information, telling players privately about others' information. This can be formalized as:

Definition 6 (Join Feasibility)
Decision rule $\sigma : T \times \Theta \to A$ *is join feasible for* (G, S) *if* $\sigma(a|t, \theta)$ *is independent of* θ*, i.e.,* $\sigma(a|t, \theta) = \sigma\left(a|t, \theta'\right)$ *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, and 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' 1993 paper was titled "Five Legitimate Definitions of Correlated Equilibrium in Incomplete Information Games". The feasibility and incentive conditions described so far allow us to completely describe the five solution concepts she discusses:

1. A *Bayesian solution* is a decision rule satisfying join feasibility and obedience.

2. A *belief invariant Bayesian solution* is a decision rule satisfying join feasibility, belief invariance and obedience.

3. An *agent normal form correlated equilibrium* is a decision rule satisfying join feasibility, public feasibility (which implies belief invariance) and obedience.

4. A *communication equilibrium* is a decision rule satisfying join feasibility and incentive compatibility (which implies obedience)

5. 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

$$\begin{aligned} (5) &\subseteq (3) \subseteq (2) \subseteq (1) \text{ and} \\ (5) &\subseteq (4) \subseteq (1) \end{aligned}$$

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 solutions corresponds - in our language - to the set of Bayes correlated equilibria that would arise under join 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 $S = \left(\left(T_i\right)_{i=1}^{I}, \pi\right)$. The information designer sends each player i an extra message $m_i \in M_i$, according to rule $\phi : T \times \Theta \to \Delta\left(M\right)$, where $M = M_1 \times \cdots \times M_I$. A communication rule is then $C = \left(\left(M_i\right)_{i=1}^{I}, \phi\right)$. 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 $b_i : T_i \times M_i \to \Delta\left(A_i\right)$. A communication rule C and strategy profile b will now induce a decision rule

$$\sigma\left(a|t,\theta\right) = \sum_{m \in M} \phi\left(m|t,\theta\right) \left(\prod_{i=1,\ldots,I} b_i\left(a_i|t_i,m_i\right) \right).$$

We will write $E\left(C\right)$ 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:

[21] Bergemann and Morris (2016a) call the pair (S, C) an "expanded information structure".

Proposition 4

Decision rule σ is a Bayes correlated equilibrium of (G,S) if and only if there exists a communication rule C and a Bayes Nash equilibrium $b \in E(C)$ which 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:

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

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

$$V^*(C,b) = \sum_{\theta,t,a} \psi(\theta)\pi(t|\theta)\left(\sum_{m\in M}\phi(m|t,\theta)\prod_{i=1,...,I} b_i(a_i|t_i,m_i)\right)v(a,\theta).$$

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 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 maximize the inefficiency of trade anticipating that the buyer and seller will play an equilibrium that minimizes inefficiency (i.e., maximizes the gains from trade).

[22] Of course, the minmax problem $\min_C \max_{b\in E(C)} V^*(C,b)$. is a reinterpretation of maxmin problem where objective is replaced by $-V^*(C,b)$. The minmin problem is similarly a reinterpretation of the maxmax problem.

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 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 an 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 to describe information structures using the language of partitions. Consider the information structure with state space $\Omega = \{1, 2, ..., .., \infty\}$ where player 1 observes the partition $(\{1\}, \{2, 3\}, \{4, 5\}, ..., \{\infty\})$ and player 2 observes the partition $(\{1, 2\}, \{3, 4\}, ..., \{\infty\})$. Let payoffs be given by $\theta = G$ at state 1 and by $\theta = B$ everywhere else. For some $q > \frac{1}{2}$, let the probability of

[23] See also Hoshino (2017) on this connection.

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. 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 $x < 1$, $\frac{x}{1+x} > \frac{1}{2}$ 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 max min problem, and all the three applica-

tions, 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 non-discriminatory case to contrast with the discriminatory case described above). 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 (2016a) 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 (2016) 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 (2017) 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. We have highlighted a duality with mechanism design: in mechanism design, a mechanism is picked for a given information structure, while in information design, an information structure is picked for a given mechanism. 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 attainable outcomes by means of some information structure; then identify the optimal information structure. We have described the close connection between information design 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 drawn a sharp contrast between information design and mechanism design. 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. As one moves into dynamic settings, the overlap becomes more central, as in Calzolari and Pavan (2006b) and Bergemann and Pesendorfer (2007). A specific setting where the interaction between mechanism design and information design has recently been studied is the area of markets with resale in which the information disclosed in the first stage fundamentally affects the interaction in the resale market, see for example Calzolari and Pavan (2006a), Dworczak (2016), Carroll and Segal (2016) and Bergemann, Brooks, and Morris (2016b).

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 can we infer about the underlying distribution of valuations under weak assumption 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.

[24] This reverse use is exposited in Bergemann and Morris (2012).

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, it is natural to ask if it can be as useful in the many player setting. In either the linear programming or 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 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 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.

9 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

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

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 $p = (p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg})$ if

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

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

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

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 $(p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg})$ will induce behavior (p_B, p_G) integrating over types $t \in \{b, g\}$.

The behavior of the equilibrium set is illustrated in Figure 2. Note that the sets becomes smaller as the firm's private information improves. Intuitively, the firm's private information limits the government's ability to influence the firm's decision as the private information tightens the obedience constraints. We observe that the boundary that describes the sets of obedient decision rules maintains a constant slope independent of q, namely $1/x$, and it is only the intercept that moves upward. Moreover, the slope is identical with the one described in the problem of the uninformed firm. The lowest probability of investing in the good state is achieved if there is zero probability of investing in the bad state as derived earlier in (11).

Additional Results for Section 6.2 For a given Bayesian game (G, S), let $\mathrm{E}(G, S)$, respectively, $\mathrm{NE}(G, S)$, denote the set of decision rules that can be attained with and without elicitation, respectively.

Proposition 5

For each (G, S), we have

$$\mathrm{NE}(G, S) \subseteq \mathrm{E}(G, S).$$

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

Fix player i, types $t_i, t_i' \in T_i$, and a function $\delta_i \colon A_i \to A_i$. 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

$$\sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \left(\sum_{(b_i, b_{-i}) \in B} u_i \left((b_i(t_i), b_{-i}(t_{-i})), \theta \right) \phi((b_i, b_{-i}) | \theta) \right) \pi(t_i, t_{-i} | \theta) \psi(\theta)$$

$$\geq \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \left(\sum_{(b_i, b_{-i}) \in B} u_i \left((\delta_i(b_i(t_i')), b_{-i}(t_{-i})), \theta \right) \phi((b_i, b_{-i}) | \theta) \right) \pi(t_i, t_{-i} | \theta) \psi(\theta).$$

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

$$\mathrm{NE}(G, S) = \mathrm{E}(G, S).$$

Proof. By Proposition 5, it suffices to prove $NE(G, S) \supseteq \mathrm{E}(G, S)$. 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 $\mathrm{E}(G, S)$. 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 $p = (p_{\theta t}) \in \mathrm{E}(G, S)$. 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 truthtelling 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

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

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

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

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

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

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

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

[25] First, 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.

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 $l = 1, \ldots, n$ and $k = 1, \ldots, n+1$, type t_l's expected utility from investing if and only if being recommended b^k is:

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

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

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

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

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

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,

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

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

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

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

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

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

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:

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

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

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

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

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, $\mathrm{E}(G,S)$ 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 $b\colon T \to A$ 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 $b =$ (not invest, invest) can never be obedient for both types, so for any $p \in NE(G,S)$, there is only one candidate ϕ, namely

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

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

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

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

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

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

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

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

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

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* $I = 2$, $\Theta = \{B, G\}$, $A_1 = A_2 =$*{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* $(b(t_1), b(t_2)) =$ (*invest*, *invest*) *in state* G; $(b(t_1), b(t_2)) =$ (*invest*,*not invest*) *in state* B_1; *and* $(b(t_1), b(t_2)) =$ (*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 iff the state is* G. *Therefore, the decision rule is not publicly feasible obedient.*

References

Alonso, R., and O. Camara (2016): "Persuading Voters," *American Economic Review*, 106, 3590–3605.

Angeletos, G.-M., and J. La'O (2013): "Sentiments," *Econometrica*, 81, 739–779.

Arieli, I., and Y. Babichenko (2016): "Private Bayesian Persuasion," Discussion paper, Technion.

Asquith, P., T. Covert, and P. Pathak (2014): "The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market," Discussion paper, NBER.

Aumann, R. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.

Aumann, R. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18.

Aumann, R., and M. Maschler (1995): *Repeated Games with Incomplete Information*. MIT Press.

Ball, I. (2017): "Optimal Continuous Disclosure," Yale University.

Bassan, B., O. Gossner, M. Scarsini, and S. Zamir (2003): "Positive Value of Information in Games," *International Journal of Game Theory*, 32, 17–31.

Bergemann, D., A. Bonatti, and A. Smolin (2017): "The Design and Price of Information," *American Economic Review*, forthcoming.

Bergemann, D., B. Brooks, and S. Morris (2015): "The Limits of Price Discrimination," *American Economic Review*, 105, 921–957.

——— (2016a): "Informationally Robust Auction Design," Discussion paper, Cowles Foundation Discussion Paper 2065.

——— (2016b): "Optimal Auction Design in a Common Value Model," Discussion paper, Yale University, University of Chicago and Princeton University.

——— (2017): "First Price Auctions with General Information Structures: Implications for Bidding and Revenue," *Econometrica*, 85, 107–143.

Bergemann, D., T. Heumann, and S. Morris (2015): "Information and Volatility," *Journal of Economic Theory*, 158, 427–465.

Bergemann, D., and S. Morris (2012): "Robust Predictions in Games with Incomplete Information," Discussion Paper 1821RR, Cowles Foundation for Research in Economics, Yale University.

——— (2013a): "Bayes Correlated Equilibrium and Comparison of Information Structures," Discussion Paper 1822R, Cowles Foundation for Research in Economics, Yale University.

——— (2013b): "Robust Predictions in Games with Incomplete Information," *Econometrica*, 81, 1251–1308.

——— (2016a): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11, 487–522.

——— (2016b): "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium," *American Economic Review Papers and Proceedings*, 106(5), 586–591.

——— (2017): "Belief-Free Rationalizability and Informational Robustness," *Games and Economic Behavior*, 104, 744–759.

Bergemann, D., and M. Pesendorfer (2007): "Information Structures in Optimal Auctions," *Journal of Economic Theory*, 137, 580–609.

Blackwell, D. (1951): "Comparison of Experiments," in *Proc. Second Berkeley Symp. Math. Statist. Probab.*, pp. 93–102. University of California Press, Berkeley.

——— (1953): "Equivalent Comparison of Experiments," *Annals of Mathematics and Statistics*, 24, 265–272.

Brocas, I., and J. Carrillo (2007): "Influence through Ignorance," *RAND Journal of Economics*, 38, 931–947.

Brooks, B., and S. Du (2017): "Optimal Auction Design with Common Values: An Informationally Robust Approach," Discussion paper, University of Chicago.

Calzolari, G., and A. Pavan (2006a): "Monopoly with Resale," *RAND Journal of Economics*, 37, 362–375.

——— (2006b): "On the Optimality of Privacy in Sequential Contracting," *Journal of Economic Theory*, 130, 168–204.

Caplin, A., and M. Dean (2015): "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105, 2183–2203.

Caplin, A., and D. Martin (2015): "A Testable Theory of Imperfect Perception," *Economic Journal*, 125, 184–202.

Carroll, G. (2016): "Informationally Robust Trade and Limits to Contagion," *Journal of Economic Theory*, 166, 334–361.

Carroll, G., and I. Segal (2016): "Robustly Optimal Auctions with Unknown Resale Opportunities," Discussion paper, Stanford University.

Chung, K.-S., and J. Ely (2007): "Foundations of Dominant Strategy Mechanisms," *Review of Economic Studies*, 74, 447–476.

Das, S., E. Kamenica, and R. Mirka (2017): "Reducing Congestion Through Information Design," in *Proceedings of the 53rd Allerton Conference on Communication, Control, and Computing*.

Daskalakis, C., C. Papadimitriou, and C. Tzamos (2016): "Does Information Revelation Improve Revenue," 17th ACM Conference on Economics and Computation.

Dekel, E., D. Fudenberg, and S. Morris (2007): "Interim Correlated Rationalizability," *Theoretical Economics*, 2, 15–40.

Dekel, E., and M. Siniscalchi (2014): "Epistemic Game Theory," Discussion paper, Northwestern University.

Doval, L., and J. Ely (2016): "Sequential Information Design," Discussion paper, California Institute of Technology and Northwestern University.

Du, S. (2016): "Robust Mechanisms under Common Valuation," Discussion paper, Simon Fraser University.

Duffie, D., P. Dworczak, and H. Zhu (2017): "Benchmarks in Search Markets," *Journal of Finance*, 72, 1983–2044.

Dughmi, S. (2017): "Algorithmic Information Structure Design: A Survey," *ACM SIGecom Exchanges*, 15, 2–24.

Dworczak, P. (2016): "Mechanism Design with Aftermarkets: Cutoff Mechanisms," Discussion paper, Stanford University.

Ely, J. (2017): "Beeps," *American Economic Review*, 107, 31–53.

Ely, J., A. Frankel, and E. Kamenica (2015): "Suspense and Surprise," *Journal of Political Economy*, 123, 215–260.

Ely, J., and M. Szydlowski (2017): "Moving the Goalposts," Discussion paper, Northwestern University and University of Minnesota.

Forges, F. (1993): "Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information," *Theory and Decision*, 35, 277–310.

——— (2006): "Correlated Equilibrium in Games with Incomplete Information Revisited," *Theory and Decision*, 61, 329–344.

Forges, F., and F. Koessler (2005): "Communication Equilibria with Partially Verifiable Types," *Journal of Mathematical Economics*, 41, 793–811.

Fudenberg, D., and J. Tirole (1991): *Game Theory*. MIT Press, Cambridge.

Gentzkow, M., and E. Kamenica (2017): "Competition in Persuasion," *Review of Economic Studies*, 84, 300–322.

Goldstein, I., and C. Huang (2016): "Bayesian Persuasion in Coordination Games," *American Economic Review*, 106, 592–596.

Gossner, O. (2000): "Comparison of Information Structures," *Games and Economic Behavior*, 30, 44–63.

Hellwig, C., and V. Venkateswaran (2009): "Setting the Right prices for the Wrong Reasons," *Journal of Monetary Economics*, 56, 57–77.

Hillas, J., E. Kohlberg, and J. Pratt (2007): "Correlated Equilibrium and Nash Equilibrium as an Observer's Assessment of the Game," Discussion Paper 08-005, Harvard University.

Hirshleifer, J. (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61, 561–574.

Horner, J., and A. Skrzypacz (2016): "Learning, Experimentation and Information Design," Discussion paper, Yale University and Stanford University.

Hoshino, T. (2017): "Multi-agent Persuasion: Leveraging Strategic Uncertainty," Discussion paper, Pennsylvania State University.

Inostroza, N., and A. Pavan (2017): "Persuasion in Global Games with Application to Stress Testing," Discussion paper, Northwestern University.

Kajii, A., and S. Morris (1997): "The Robustness of Equilibria to Incomplete Information," *Econometrica*, 65, 1283–1309.

Kamenica, E., and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Kamenica, E., and M. Gentzkow (2014): "Costly Persuasion," *American Economic Review*, 104, 457–462.

Kolotilin, A. (2017): "Optimal Information Disclosure: A Linear Programming Approach," *Theoretical Economics*, forthcoming.

Kolotilin, A., M. Li, T. Mylovanov, and A. Zapechelnyuk (2017): "Persuasion of Privately Informed Receiver," *Econometrica*, forthcoming.

Laffont, J. (1989): *The Economics of Uncertainty and Information*. MIT Press, Cambridge.

Lehrer, E., D. Rosenberg, and E. Shmaya (2010): "Signaling and Mediation in Games with Common Interest," *Games and Economic Behavior*, 68, 670–682.

Liu, Q. (2015): "Correlation and Common Priors in Games with Incomplete Information," *Journal of Economic Theory*, 157, 49–75.

Lucas, R. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4, 103–124.

Magnolfi, L., and C. Roncoroni (2017): "Estimation of Discrete Games with Weak Assumptions on Information," Discussion paper.

Makris, M., and L. Renou (2017): "Information Design in Multi-Stage Games," Discussion paper, University of Southampton and Queen Maru University.

Mas-Collel, A., M. Whinston, and J. Green (1995): *Microeconomic Theory*. Oxford University Press, Oxford.

Maskin, E. (1999): "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66, 23–38.

Mathevet, L., L. Perego, and I. Taneva (2017): "On Information Design in Games," Discussion paper, New York University.

Mertens, J., and S. Zamir (1985): "Formalization of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1–29.

Meyer, M. (2017): "Information Design: Insights from Orderings of Dependence and Heterogeneity," Discussion paper, Oxford University.

Milgrom, P., and R. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

Morris, S., and T. Ui (2005): "Generalized Potentials and Robust Sets of Equilibria," *Journal of Economic Theory*, 124, 45–78.

Myerson, R. (1982): "Optimal Coordination Mechanism in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 10, 67–81.

——— (1983): "Mechanism Design by an Informed Principal," *Econometrica*, 51, 1767–1797.

——— (1986): "Multistage Games with Communication," *Econometrica*, 54, 323–358.

——— (1987): "Bayesian Equilibrium and Incentive-Compatibility: An Introduction," in *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, ed. by L. Hurwicz, D. Schmeidler, and H. Sonnenschein, pp. 229–259. Cambridge University Press, Cambridge.

——— (1991): *Game Theory*. Harvard University Press, Cambridge.

Myerson, R., and M. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 28, 265–281.

Myerson, R. B. (1994): "Communication, Correlated Equilibria and Incentive Compatability," in *Handbook of Game Theory Volume 2*, ed. by R. Aumann, and S. Hart, vol. 2, chap. 24. Elsevier.

Mylovanov, T., and T. Troeger (2012): "Informed Principal Problems in Environments with Generalized Private Values," *Theoretical Economics*, 7, 465–488.

——— (2014): "Mechanism Design by An Informed Principal: Private Values with Transferable Utility," *Review of Economic Studies*, 81, 1668–1707.

Neyman, A. (1991): "The Positive Value of Information," *Games and Economic Behavior*, 3, 350–355.

Novshek, W., and H. Sonnenschein (1982): "Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release," *Bell Journal of Economics*, 13, 214–218.

Ostrovsky, M., and M. Schwarz (2010): "Information Disclosure and Unraveling in Matching Markets," *American Economic Journal: Microeconomics*, 2, 34–63.

Passadore, J., and J. P. Xandri (2016): "Robust Conditional Predictions in a Model of Sovereign Debt," Discussion paper, Princeton University.

Perez-Richet, E. (2014): "Interim Bayesian Persuasion: First Steps," *American Economic Review Papers and Proceedings*, 104, 469–474.

Rayo, L., and I. Segal (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987.

Roesler, A., and B. Szentes (2017): "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107, 2072–2080.

Rubinstein, A. (1989): "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge'," *American Economic Review*, 79, 385–391.

Syrgkanis, V., E. Tamer, and J. Ziani (2017): "Inference on Auctions under Weak Assumptions on Information," Discussion paper, Harvard University.

Taneva, I. (2015): "Information Design," Discussion paper, University of Edinburgh.

Ui, T. (2001): "Robust Equilibria of Potential Games," *Econometrica*, 69, 1373–1380.

Venkateswaran, V. (2013): "Heterogeous Information and Labor Market Fluctuations," Discussion paper, New York University.

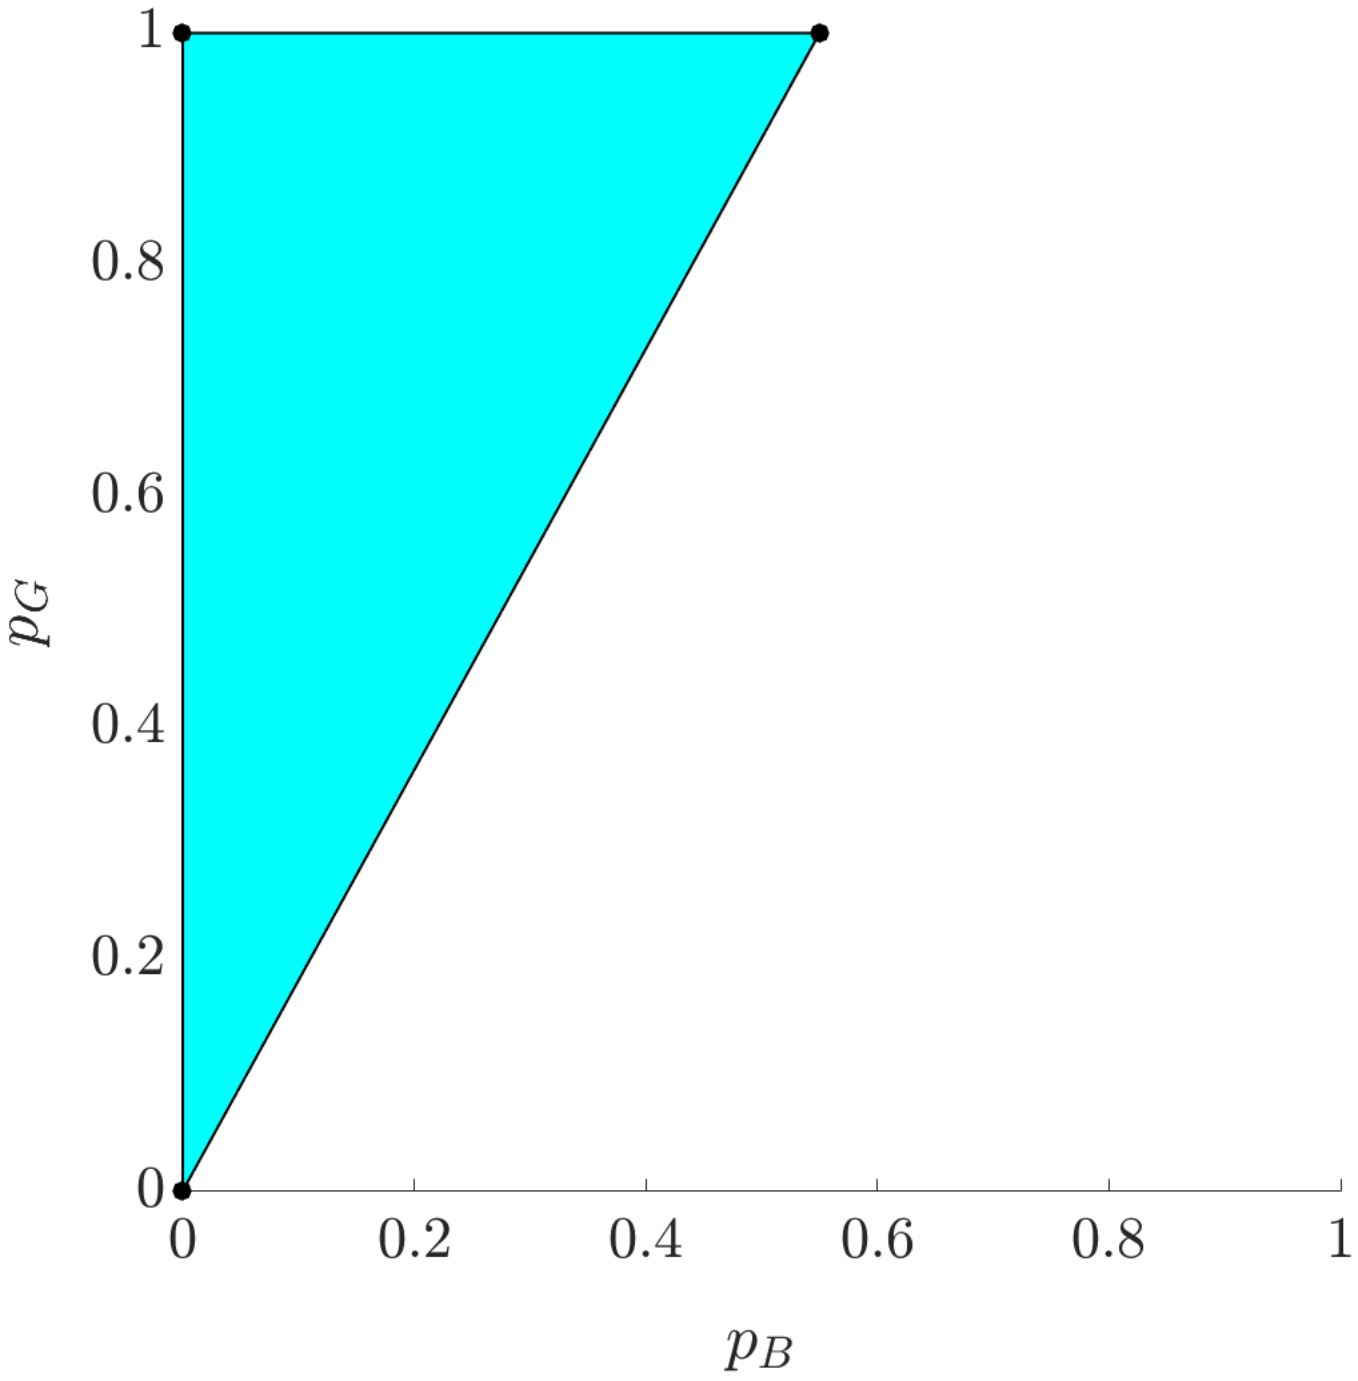


Figure 1: Investment Probabilities with Uninformed Player: x=55/100

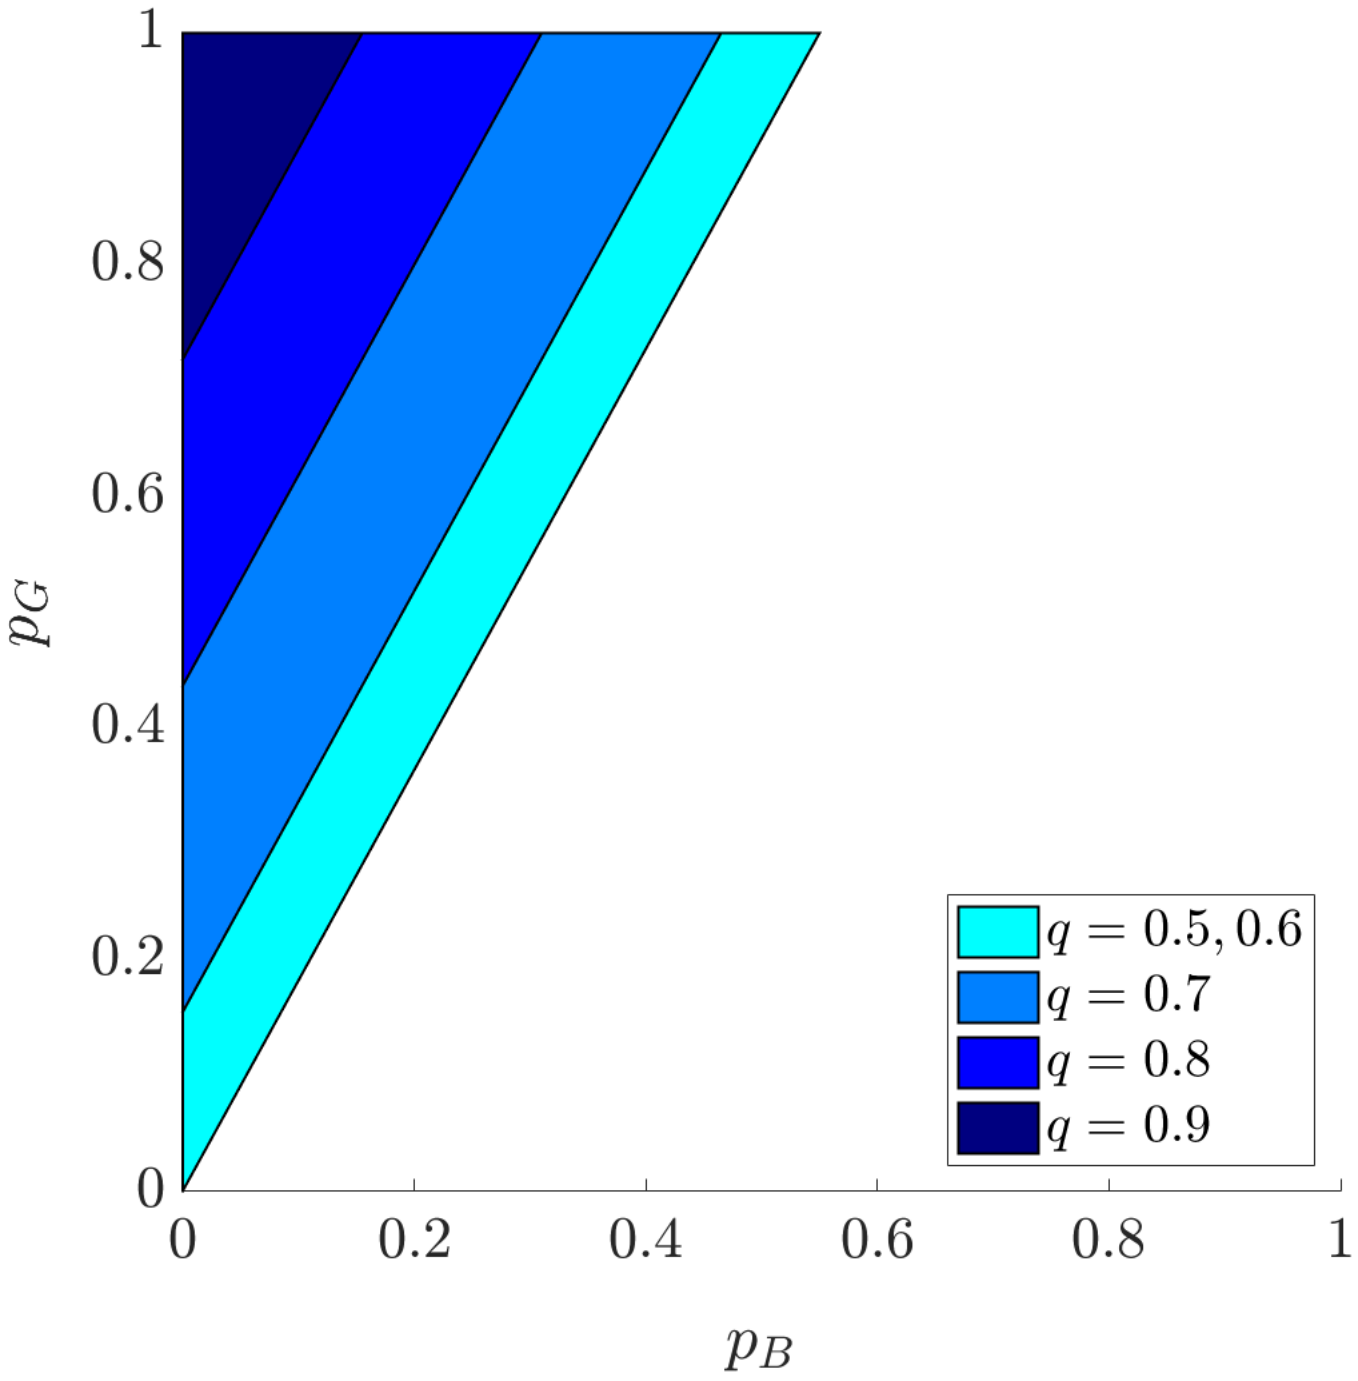


Figure 2: Investment Probability with Informed Player: x=55/100.

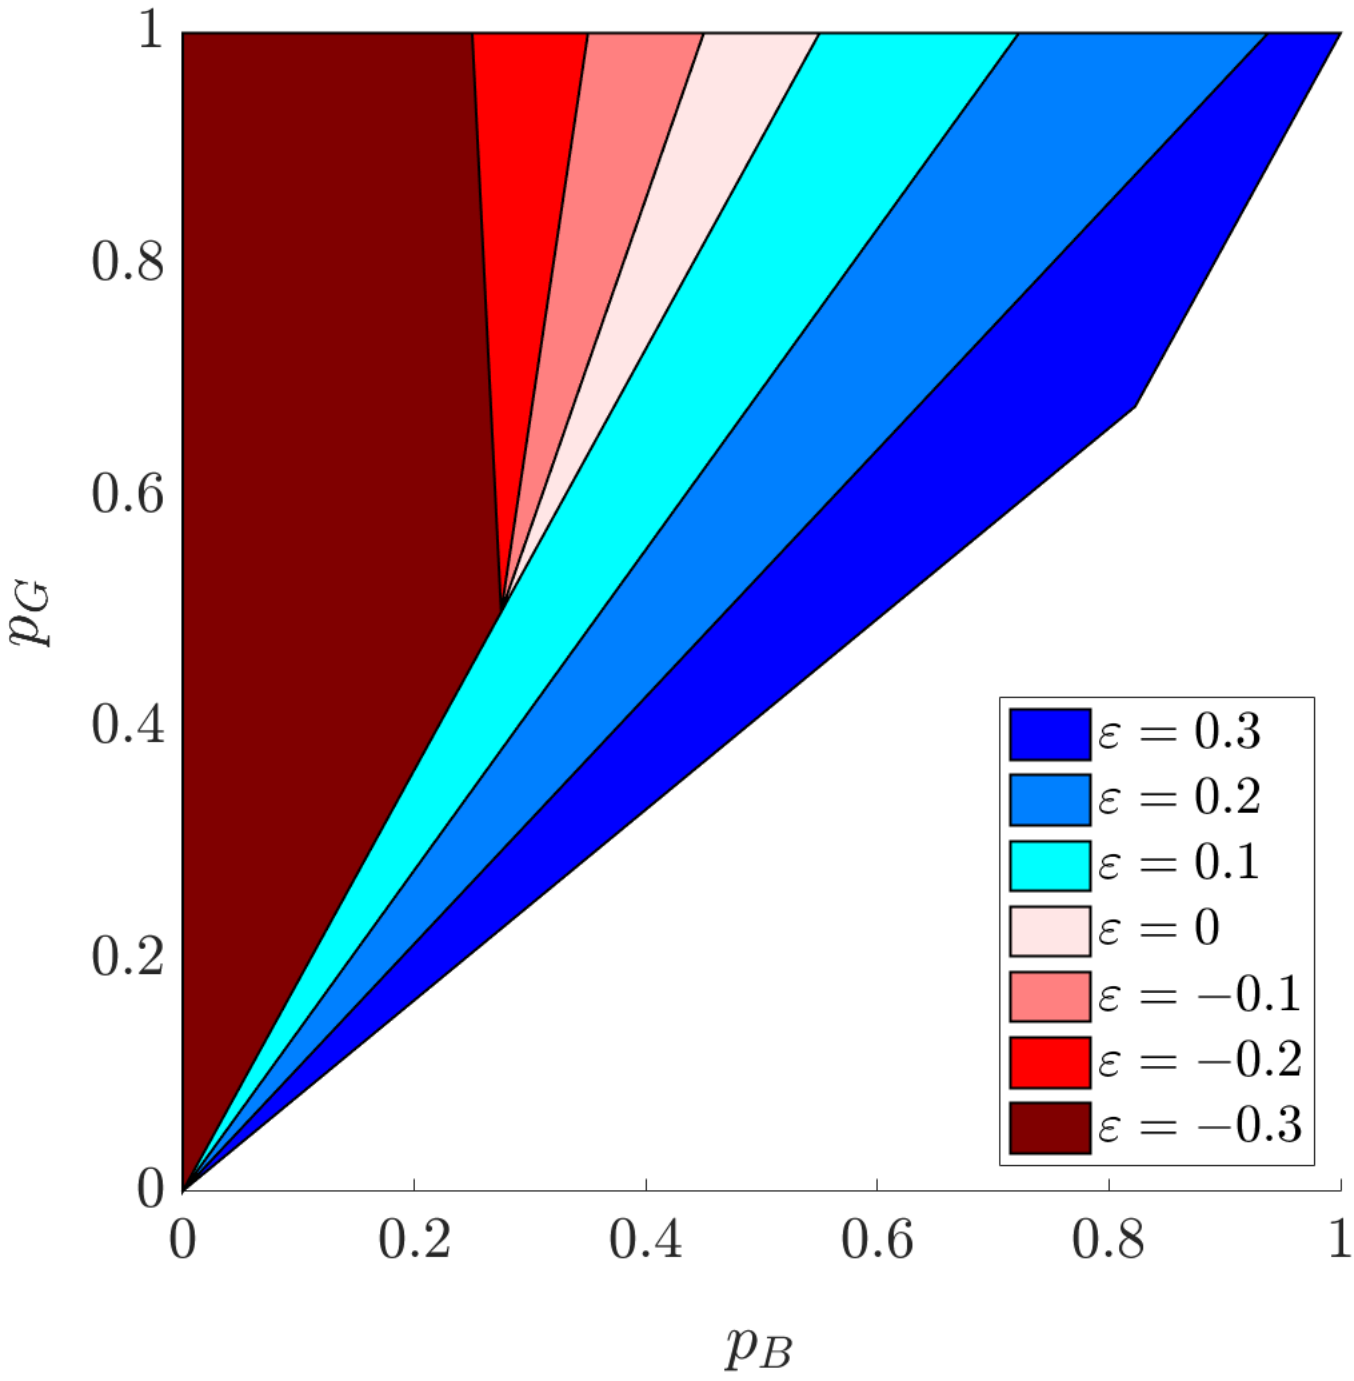


Figure 3: Investment Probability with Negative or Positive Strategic Term ε.

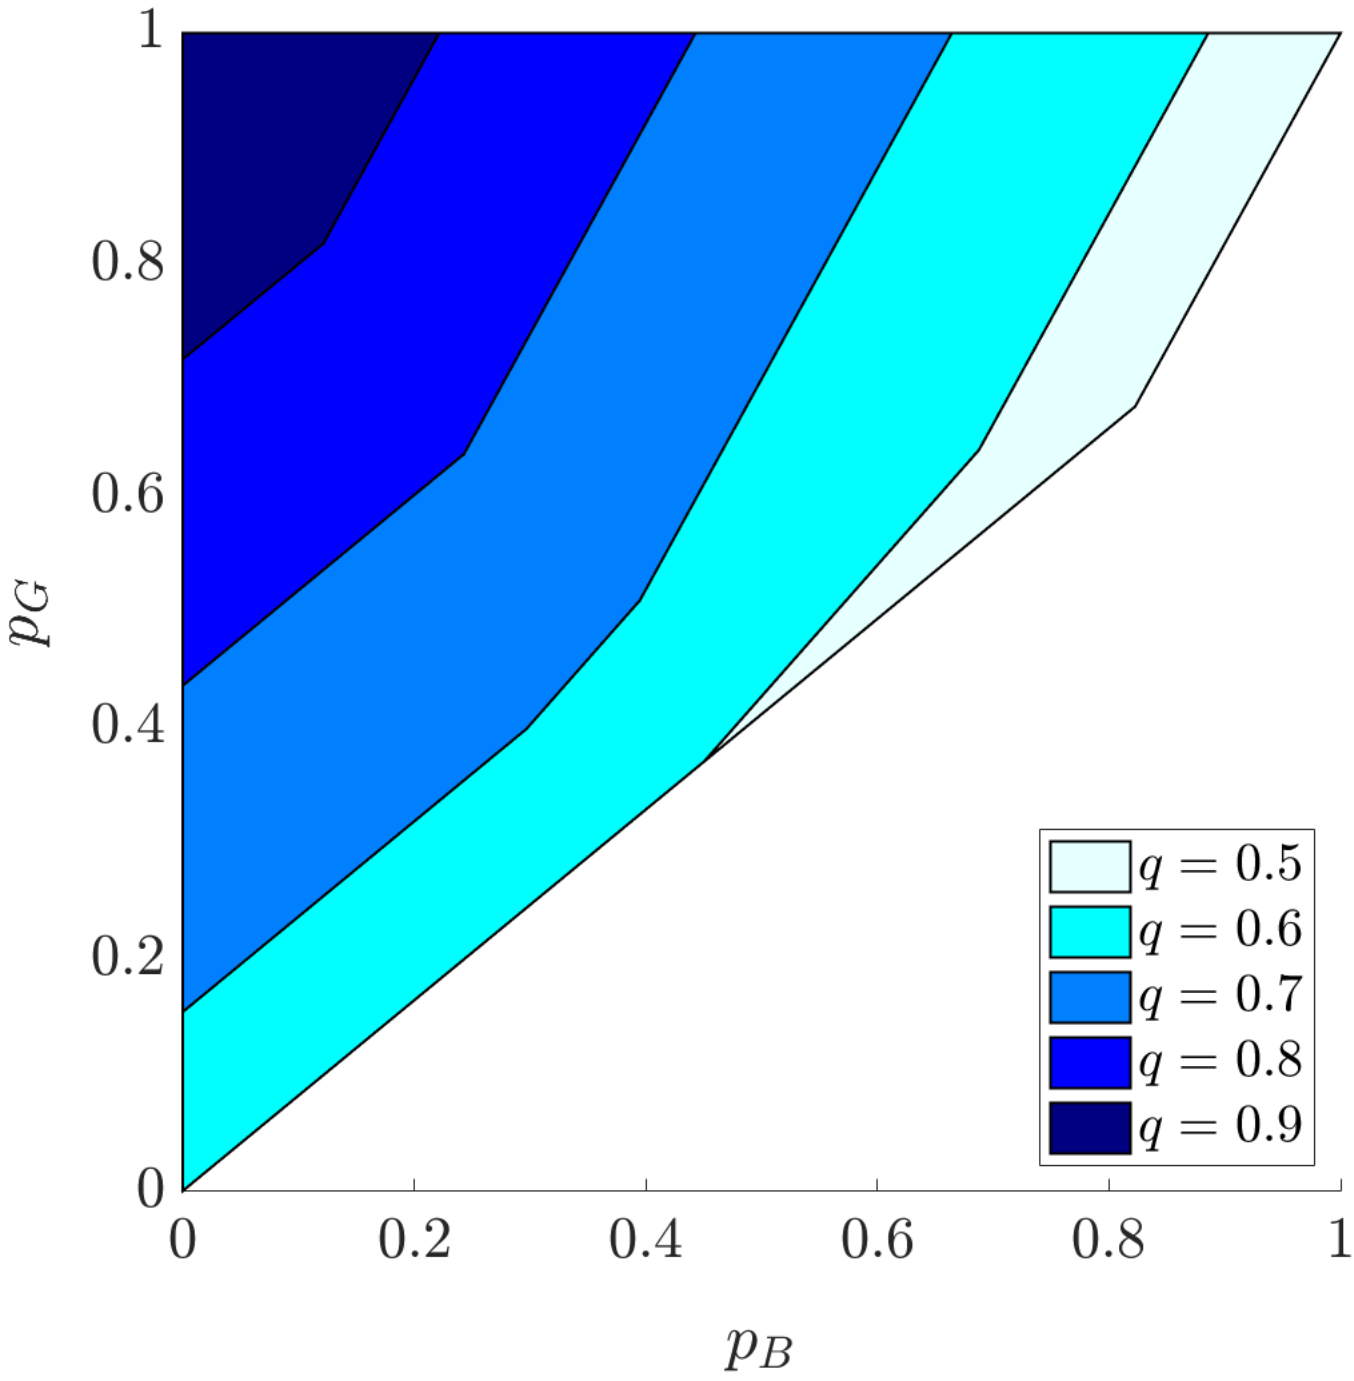


Figure 4: Investment Probability with Two Players with Prior Information, with Strategic Term ε=3/10.

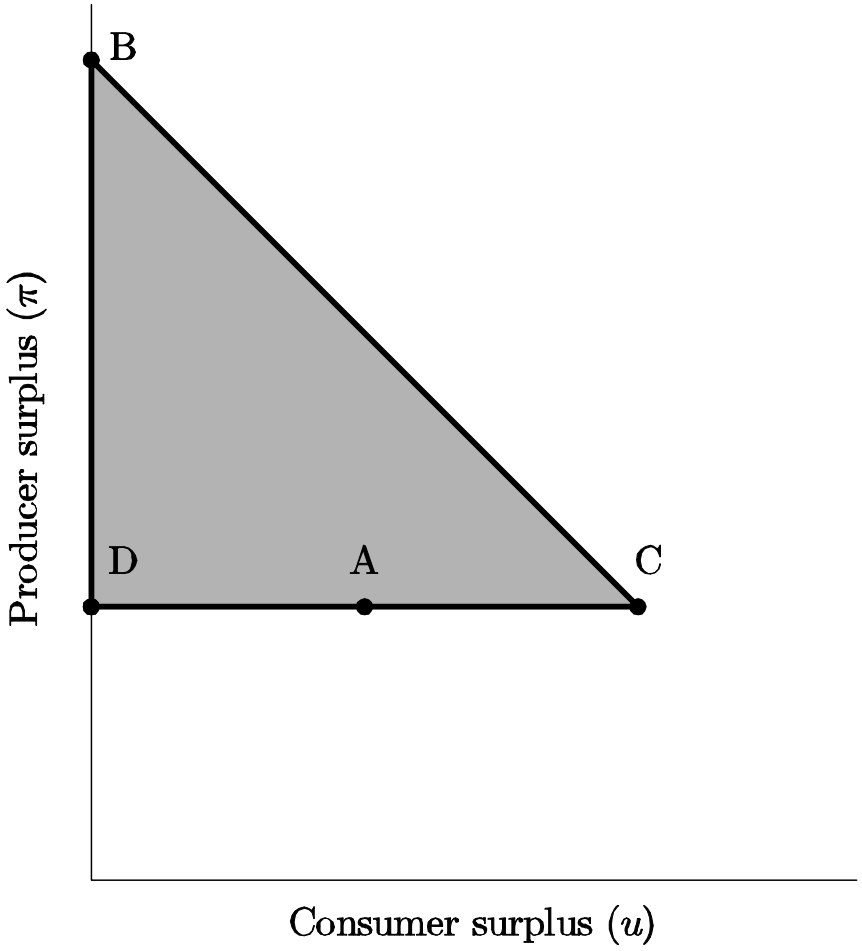


Figure 5: The Bounds on Profits and Consumer Surplus
in Third Degree Price Discrimination

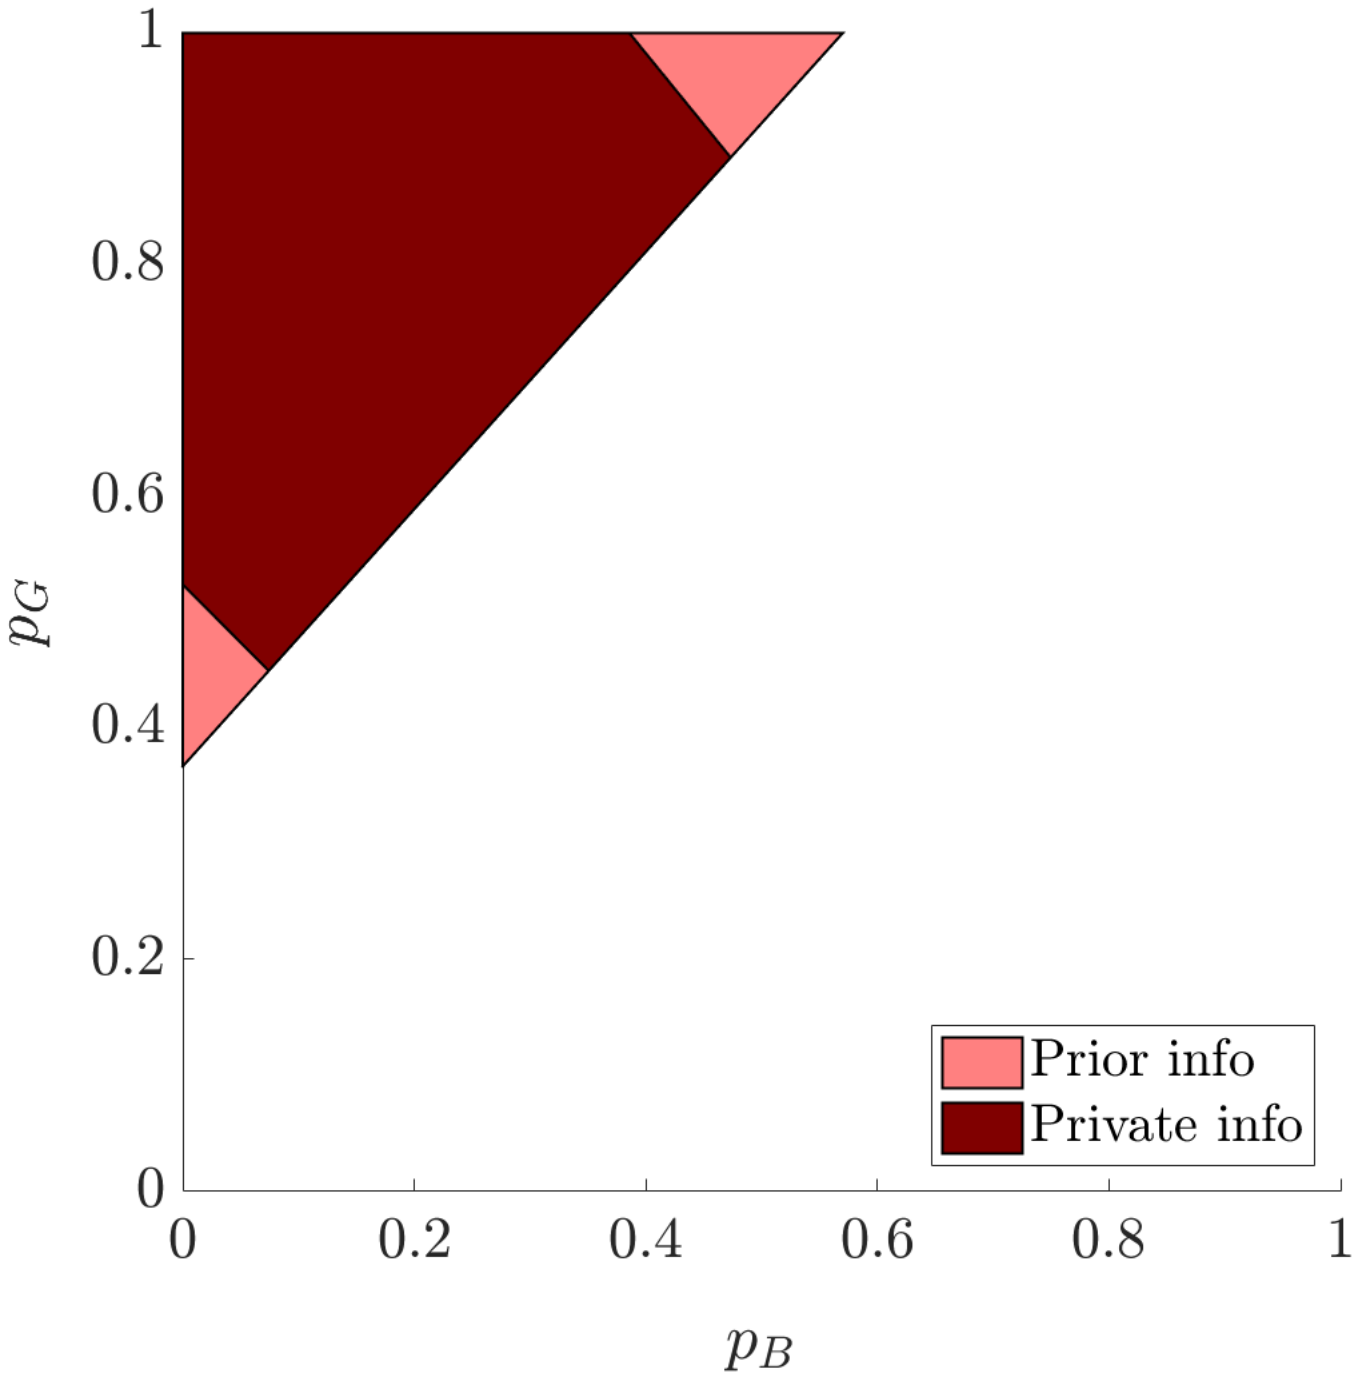


Figure 6: Investment Probability with Private Information

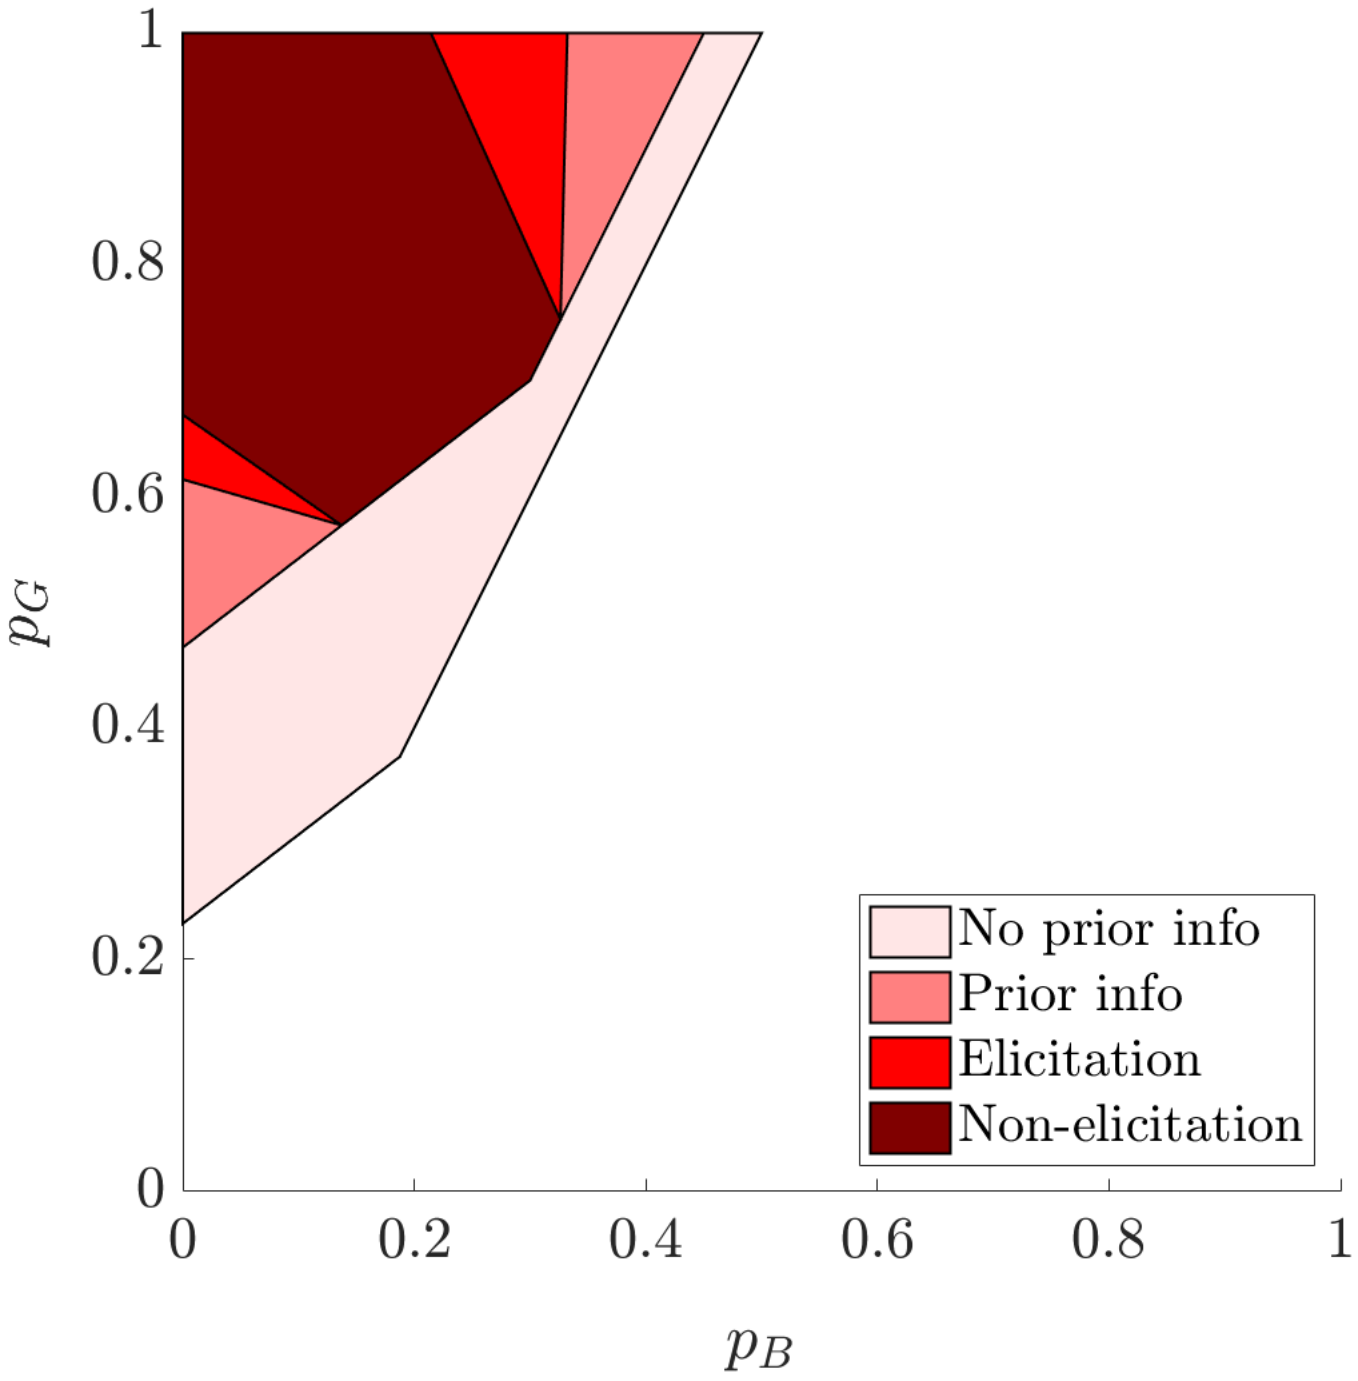


Figure 7: Investment Probability under Different Information Design Scenarios.