

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*. 2543.

<https://elischolar.library.yale.edu/cowles-discussion-paper-series/2543>

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 March 2017

COWLES FOUNDATION DISCUSSION PAPER NO. 2075R



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[‡]

March 22, 2017

Abstract

Fixing a game with uncertain payoffs, *information design* identifies the information structure and equilibrium that maximizes the payoff of an information designer. We show how this perspective unifies existing work, including that on communication in games (Myerson (1991)), Bayesian persuasion (Kamenica and Gentzkow (2011)) and some of our own recent work. 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.

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. This material has been presented in lectures at the 2015 Istanbul meetings of the Society of Economic Design, the 2015 Delhi School of Economics Winter School, the 2016 AEA meetings, the North American Summer Meeting of the Econometric Society, the Gerzensee Summer Symposium in Economic Theory, the Becker Friedman Institute conference on Frontiers and Economic Theory and Computer Science, the Harris lecture at Harvard University and the 2016 fall Finance Theory group conference at Princeton. Some of the material in this paper was previewed in Bergemann and Morris (2016b). We have benefited from discussions and correspondence about the material in this paper with Ben Brooks, Laura Doval, Jeff Ely, Françoise Forges, Drew Fudenberg, Olivier Gossner, Tibor Heumann, Atsushi Kajii, Emir Kamenica, Rohit Lamba, Laurent Mathevet, Roger Myerson, Tymofiy Mylanov, Alessandro Pavan, Eran Shmaya, Jonathan Weinstein and Juan Pablo Xandri; and from valuable research assistance from Ian Ball, 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	3
2	Information Design	7
2.1	An Investment Example	9
2.1.1	Single Player without Prior Information	9
2.1.2	Single Player with Prior Information	11
2.1.3	Many Players without Prior Information	12
2.1.4	Many Players with Prior Information	14
2.2	Issues in Information Design illustrated by the Examples	15
2.2.1	Public versus Private Signals; and Instrumental versus Intrinsic Motivation for Preferences Over Correlation	15
2.2.2	Bayes Correlated Equilibrium Outcomes and Information Design without Concavification	19
2.2.3	Tightening Obedience Constraints and Bayes Correlated Equilibrium Outcomes .	20
2.2.4	Metaphorical Information Design: Robust Predictions and Maxmin Objectives .	23
3	Information Design with Private Information and the Relation to Mechanism Design	24
3.1	Information Design when Players' Prior Information is not known to the Designer . . .	24
3.2	The Investment Example Re-Visited	27
3.3	General Mechanism Design	31
3.4	Incomplete Information Correlated Equilibrium	32
3.4.1	Belief Invariance	32
3.4.2	Join Feasibility	33
4	Information Design with Adversarial Equilibrium and Mechanism Selection	35
4.1	Adversarial Equilibrium Selection	35
4.2	Adversarial Mechanism Design	39
5	Conclusion	40
6	Appendix	41

1 Introduction

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

One purpose of the paper is to provide an overview of information design that unifies 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 (Kamenica and Gentzkow (2011)). Some of our recent work corresponds to the information design problem when there are *both* many players *and* the information designer has an informational advantage over the players (Bergemann and Morris (2013b), (2016a)). 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

¹We follow Taneva (2015) in our use of the term "information design" in this context. 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 3.3.

conditioned on states that the players do not know.

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

To survey the literature, and to provide some graphical illustrations, we use a family of two player, two action, two state and two signal examples to survey the literature. We start with the leading example of Bayesian persuasion (with a single player/receiver with no prior information) from the work of Kamenica and Gentzkow (2011). We can use extensions of this example - with many players and prior information - to illustrate many of the key ideas in the survey. Three key substantive general insights are illustrated in these examples. First, it is often optimal for the information designer to selectively obfuscate information. This insight is familiar from the one player without prior information case. 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 will 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 outcome 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.

Our analysis of information design focusses what we will sometimes call the *omniscient* case: the information designer knows not only the payoff-relevant state of the world, but also knows the players' prior information about the state and others' information. We also consider information design with private information, where players' prior information is not known by the information designer, even though she knows everything about the payoff-relevant states (which are in turn not known by the players). 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 know them (information design with elicitation) or she may be unable to do so (information design without elicitation). If the information designer has no information of her own, 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.² We follow Mathevet, Perego, and Taneva (2016) in formally describing a notion of maxmin information design,

²This point has been highlighted by Carroll (2016) and Mathevet, Perego, and Taneva (2016).

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)).

This paper focusses on static settings when all information structures are allowed, and free, and there is a single information designer. This allows an appeal to the revelation principle. Of course, there are many (static) settings where the impact of different information structures has been studied, without allowing all information structures. Two classic examples would be information sharing in oligopoly (Novshek and Sonnenschein (1982)) and the linkage principle in auction theory (Milgrom and Weber (1982)). Optimal information design in dynamic settings has been studied recently in applications of Ely, Frankel, and Kamenica (2015), Ely (2017) and Passadore and Xandri (2016). Doval and Ely (2016) study a general class of optimal dynamic information design problems. Horner and Skrzypacz (2016) surveys work on information design more generally 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 survey applications.

Our main results are in Section 2, where we describe the information design problem and review it using our examples. We also discuss private versus public signals, intrinsic versus instrumental preferences over correlation, the two step procedure for solving information design problems, ordering information, the widely used concavification in information design (instead of pure linear programming methods) and metaphorical information design interpretations. In Section 3, 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 4, 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 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 eclectic in our labelling, compromising between the use of familiar terminology and having a unified language.

2 Information Design

Throughout the paper, we will fix a finite set of players and a finite set of payoff states of the world. There are I players, $1, 2, \dots, I$, and we write i for a typical player. We write Θ for the 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 a utility function $u_i : A \times \Theta \rightarrow \mathbb{R}$, where we write $A = A_1 \times \dots \times A_I$; and (2) a full support prior $\psi \in \Delta(\Theta)$. Thus $G = \left((A_i, u_i)_{i=1}^I, \Theta, \psi \right)$. An "information structure" S consists of (1) for each player i , a finite set of types or "signals" T_i ; and (2) a signal distribution $\pi : \Theta \rightarrow \Delta(T)$, where we write $T = T_1 \times \dots \times T_I$. 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 set of decision rules $\sigma : T \times \Theta \rightarrow \Delta(A)$ that can be induced by an information designer. In this section, we will consider the leading case where the designer can condition on the state *and* on the players' prior information, if they have any. In this context, we will sometimes refer to the *omniscient* information designer. In the next section, we will consider the case where prior information of the players is private, and hence the information designer cannot condition on their prior information unless she is able to induce them to reveal it.

Assuming that the information designer knows the true state $\theta \in \Theta$ and the signal vector $t \in T$, obedience will be the key restriction on decision rules. Obedience is the requirement that if the information designer privately communicated information as stochastic action recommendations according to σ , the players would want to follow the recommendation.

Definition 1 (Obedience)

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

$$\begin{aligned} & \sum_{a_{-i} \in A_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} u_i((a_i, a_{-i}), \theta) \sigma((a_i, a_{-i}) | (t_i, t_{-i}), \theta) \pi((t_i, t_{-i}) | \theta) \psi(\theta) \\ & \geq \sum_{a_{-i} \in A_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} u_i((a'_i, a_{-i}), \theta) \sigma((a_i, a_{-i}) | (t_i, t_{-i}), \theta) \pi((t_i, t_{-i}) | \theta) \psi(\theta); \end{aligned} \quad (1)$$

for all $a'_i \in A_i$.

Bergemann and Morris (2016a) define a *Bayes correlated equilibrium* (BCE) to be a decision rule satisfying obedience.

Proposition 1 *An omniscient information designer can attain decision rule σ if and only if it is a Bayes correlated equilibrium.*

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.³ In this (and later) propositions, we omit formal statements and proofs that correspond to revelation principle arguments. Bergemann and Morris (2016a) give a formal statement⁴ and proof of this proposition as Theorem 1, adapting standard characterizations of complete and incomplete information correlated equilibrium. But these correlated equilibrium characterizations are themselves revelation principle arguments, 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 there is "incentive compatibility". In the case of complete information, Myerson (1991) (section 6.2) describes this as the "revelation principle for strategic form games". Note that the expression "revelation principle" is often limited to the case where agents are sending messages rather than receiving them (e.e.g, Fudenberg and Tirole (1991) and Mas-Collel, Whinston, and Green (1995)), but we use the broader meaning throughout the paper.

Now if we let $v : A \times \Theta \rightarrow \mathbb{R}$ be the ex post objective of the information designer, then the utility of the information designer from decision rule σ is

$$V_{G,S}(\sigma) = \sum_{a \in A, t \in T, \theta \in \Theta} v(a, \theta) \sigma(a|t, \theta) \pi(t|\theta) \psi(\theta). \quad (2)$$

The (omniscient) information design problem is then to pick a BCE σ to maximize $V_{G,S}(\sigma)$. 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).⁵ In this case, the single player is called the "receiver" and the information designer is called the "sender". In the next subsection, we describe in detail a series of examples to illustrate some key ideas about information design.

³We do not discuss information design under solution concepts other than Bayes Nash equilibrium in this paper. Mathevet, Perego, and Taneva (2016) study information design under bounded level rationalizability and Inostroza and Pavan (2017) under full rationalizability.

⁴A formal statement also appears in Section 4.

⁵This problem was also studied by Rayo and Segal (2010) and Ostrovsky and Schwarz (2010).

2.1 An Investment Example

We will 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 are summarized in the following matrix:

	bad state B	good state G	
invest	-1	x	.
not invest	0	0	

(3)

2.1.1 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 firm has no information about the state. A decision rule $\sigma : \Theta \rightarrow \Delta(A)$ now specifies the probability of investment p_θ conditional on the true state $\theta \in \{B, G\}$. Thus a decision rule is a pair (p_B, p_G) . We can think of a decision rule as a (stochastic) action recommendation from the government. If the recommendations are obeyed, the *outcome* - the ex ante distribution over states and actions - is given by

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

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

$$\frac{\frac{1}{2}p_B}{\frac{1}{2}p_B + \frac{1}{2}p_G}(-1) + \frac{\frac{1}{2}p_G}{\frac{1}{2}p_B + \frac{1}{2}p_G}x \geq 0. \quad (4)$$

This can be rearranged to give:

$$p_G \geq \frac{p_B}{x}. \quad (5)$$

There is an analogous obedience constraint corresponding to the recommendation not to invest.

The obedience conditions reflect the fact that the firm may be given information (that we do not need to describe explicitly by the revelation principle) that leads it to act differently across the two states, hence p_B and p_G may differ. Because the firm would not invest in either state with no information - by our maintained assumption that $x < 1$ - the binding obedience constraint will be the one corresponding to investment, i.e., inequality (5). 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) that is 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.

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

Now any 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:

	bad state B	good state G	
invest	$\frac{1}{2}x$	$\frac{1}{2}$. (6)
not invest	$\frac{1}{2}(1 - x)$	0	

This will give rise to the outcome that has the firm investing whenever the signal is good and not investing when the signal is bad. Thus a government trying to encourage investment will *obfuscate* the states of the world in order to maximize investment. By pooling bad and good states in this way, the firm is made indifferent between investing or not after the good signal realization, and the indifference is broken in favor of investment. The bad state is completely isolated in the bad signal. 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.

2.1.2 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.⁶ In particular, 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:

S	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 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:

$$p_{\theta t} = (p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg}). \quad (7)$$

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:

$$\frac{(1 - q)p_{Bg}}{(1 - q)p_{Bg} + qp_{Gg}}(-1) + \frac{qp_{Gg}}{(1 - q)p_{Bg} + qp_{Gg}}x \geq 0. \quad (8)$$

However, we are interested in what we can say about the joint distribution of states and actions ex ante, integrating out the types. One can show that there is a lower bound on investment in the good state given by:

$$p_G = q - \frac{1 - q}{x}, \quad (9)$$

which approaches 1 as q approaches 1. 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. We return to the issue of comparing information structures in Section 2.2.3.

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

⁶Some detailed calculations for this example appear in the Appendix.

2.1.3 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.⁷ 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. Thus we can write firm 1's payoff for the incomplete information game as:

$$\begin{array}{c}
 \text{firm 1} \\
 \begin{array}{|c|c|c|}
 \hline
 & \text{firm 2} & \\
 \hline
 \theta = B & \text{invest} & \text{not invest} \\
 \hline
 \text{invest} & -1 + \varepsilon & -1 \\
 \hline
 \text{not invest} & 0 & 0 \\
 \hline
 \end{array}
 \end{array}
 \quad
 \begin{array}{c}
 \text{firm 1} \\
 \begin{array}{|c|c|c|}
 \hline
 & \text{firm 2} & \\
 \hline
 \theta = G & \text{invest} & \text{not invest} \\
 \hline
 \text{invest} & x + \varepsilon & x \\
 \hline
 \text{not invest} & 0 & 0 \\
 \hline
 \end{array}
 \end{array}
 \quad (10)$$

We can focus on symmetric decision rules, given the symmetry of the basic game, for any symmetric objective of the information designer. To see why, note that if we found an asymmetric maximizing decision rule, the decision rule changing the names of the firms would also be optimal and so would the (symmetric) average of the two decision rules. Therefore, we will continue to write p_θ for the probability that each firm will invest in state $\theta \in \{G, B\}$; but we will now write r_θ for the probability that both invest. Thus a decision rule is a vector (p_B, r_B, p_G, r_G) . A decision rule can now be represented in a table as

$$\begin{array}{|c|c|c|}
 \hline
 \theta = B & \text{invest} & \text{not invest} \\
 \hline
 \text{invest} & r_B & p_B - r_B \\
 \hline
 \text{not invest} & p_B - r_B & 1 + r_B - 2p_B \\
 \hline
 \end{array}
 \quad
 \begin{array}{|c|c|c|}
 \hline
 \theta = G & \text{invest} & \text{not invest} \\
 \hline
 \text{invest} & r_G & p_G - r_G \\
 \hline
 \text{not invest} & p_G - r_G & 1 + r_G - 2p_G \\
 \hline
 \end{array}
 \quad (11)$$

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

$$-\frac{1}{2}p_B + \frac{1}{2}p_Gx + \frac{1}{2}(r_B + r_G)\varepsilon \geq 0, \quad (12)$$

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

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

⁷Other two player, two action and two state examples appear in Bergemann and Morris (2013a), (2016a) and Taneva (2015).

Since $x < 1$, (12) is always the binding constraint and - for $|\varepsilon|$ sufficiently close to 0 - we can rewrite it by the same reasoning as in Section 2.1.1 as

$$p_G \geq \frac{p_B}{x} - (r_B + r_G) \frac{\varepsilon}{x}. \quad (13)$$

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 (13). 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 $\varepsilon > 0$, (13) 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 (13), 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	$\theta = G$	invest	not invest
invest	$\frac{x+\varepsilon}{1-\varepsilon}$	0	invest	1	0
not invest	0	$\frac{1-x-2\varepsilon}{1-\varepsilon}$	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 restriction attention to $|\varepsilon| \leq x - \frac{1}{2}$. In this case, it will be optimal to set $r_B = 0$. Substituting these expressions into (13), 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	$\theta = G$	invest	not invest
invest	0	$x + \varepsilon$	invest	1	0
not invest	$x + \varepsilon$	$1 - 2x - 2\varepsilon$	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 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 .

2.1.4 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 one player case. In Figure 4 we illustrate the set of attainable investment probabilities under increasing prior information with strategic complementarities. The strategic complementarities gives 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 $\varepsilon=3/10$.

2.2 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 one player with no prior information case.

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 our one player with prior information case, more prior information constrains the ability of the designer to control outcomes; we discuss the many player generalization of this observation. Third, we note that the information design problem can be used to address many important questions where there is not a literal information designer. In particular, understanding the set of outcomes that an information designer can induce corresponds to identifying that set of all outcomes that could arise for some information structure. Finally, we discuss the elegant approach of using "concavification" to characterize and provide insight into the information designer's problem; we did not use this above, rather we gave a purely linear programming representation of the problem. We 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.

2.2.1 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 many players without prior information case 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 two player without prior information case, 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 under strategic substitutes.

We now describe three environments where there will be only instrumental concerns about correlation. First, Mathevet, Perego, and Taneva (2016) 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 (2016), 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 2.1.3, 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	$\theta = G$	invest	not invest
invest	$-1 + \varepsilon + z$	-1	invest	$x + \varepsilon + z$	x
not invest	z	0	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.⁸ 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 too 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.⁹

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 of Section 2.1.3 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

⁸This corresponds to a large literature on information sharing in oligopoly following Novshek and Sonnenschein (1982).

⁹In 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.

ensure that one firm always enters. Consider the following decision rule:

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

If ε were equal to 0, this decision rule would be obedient, with all constraints holding strictly: a firm told to not invest would have a strict incentive to obey, since it would know that the state was bad; a firm told to invest would have a strict incentive to obey, since its expected payoff will be $\frac{2}{3} (x - \frac{1}{2}) > 0$. Because the obedience constraints hold strictly, this decision rule will continue to be obedient, for positive or negative ε , as long as $|\varepsilon|$ is sufficiently small. Note that the government’s objective, of maximizing the probability that at least one firm invests, necessitates private signals. Bergemann, Heumann, and Morris (2015) show that in a world where each player wants to set his action equal to the sum of a common shock and an idiosyncratic shock, aggregate volatility is maximized when players observe signals that are a weighted sum of the shocks, but with more weight on the common shock. This is true in a setting where there is no strategic interaction between the players (a special case in that paper), and information cannot influence each player’s average action, but only their correlation.¹⁰ In a dynamic setting, Ely (2017) shows how an information designer with intrinsic preferences for negative correlated will optimally use private signals to induce it (he also shows that this is consistent with players strategic objectives). Arieli and Babichenko (2016) provides an elegant characterization 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.

In many cases, the designer’s preferences over correlation will not be exclusively instrumental or intrinsic. In the paper described above on bounding aggregate variance (Bergemann, Heumann, and Morris (2015)), we also consider the strategic case where firms do care about other firms’ level of output. In this case, strategic substitutes will act as a constraint on the ability of the designer to maximize aggregate volatility, and he will have an incentive to make information more private than he otherwise would.

¹⁰In this work and others, we have considered the information design problem in basic games with linear best responses and normally distributed payoff states, and where we restrict attention to normal and symmetric information structures. These works thus illustrate information design ideas within a rich but not complete general class of information structures.

2.2.2 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 an important "one step" approach to information design: concavification. 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. We call this a "one step" approach, because it avoids the first step of characterizing the BCE, and works from the beginning with the information designer's objective. 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 (2016) 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, 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. Although 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 results are used in stating or proving our results (although linear programming played an important role in supplying conjectures for the results).

2.2.3 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 that 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 where the extra information imposes more obedience constraints. Recall that in our language, communication in games corresponds to information design when the information designer has no information of her own.

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

This was illustrated in our one player with prior information example. 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 but not by simply giving the player more signals. We will now informally describe how this observation can be generalized in many directions. First, this result will continue to hold in the one player case for more general games (i.e., decision problems) and orderings on information. In the one player case, an information structure reduces to an experiment in the sense of Blackwell (1951), (1953). If an experiment is more informative than another in Blackwell's sense, then - in any decision problem - the set of BCE outcomes for a given experiment is smaller under the more informative experiment. There

is also a converse. If an experiment is not more informative than another, then one can find a decision problem and an outcome that is a BCE for the first experiment but not for the second.

Taken together, there is now an elegant set of connections between Blackwell's theorem and the information design problem. Blackwell describe a natural *statistical* ordering on experiments: one experiment is more informative than another if the former is sufficient for the latter, meaning that 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). Generalizing the above example, one can show in general that this incentive ordering is equivalent to the statistical ordering. (This result is the one player special case of the main result (for many player information structures) from Bergemann and Morris (2016a).)

The incentive ordering is conceptually different from the *feasibility* ordering studied in Blackwell (1951), (1953). Say that one experiment is *more valuable than* another if the set of outcomes (joint distributions over actions and states) that can be induced by decision rules mapping signals to action is larger - in any decision problem - under the first experiment. One can show that there is a three way equivalence between (i) the sufficiency ordering; (ii) the more valuable ordering; and (iii) the incentive constrained ordering. Blackwell's theorem shows an equivalence between (i) and (ii). The result described in the previous paragraph showed the equivalence between (i) and (iii). Bergemann and Morris (2013a) discuss the trichotomy of statistical, feasibility and more incentive constrained orderings, as well as the one player special case, in more detail.

The definition of the incentive ordering generalizes to the many player case. Bergemann and Morris (2016a) characterize the many player statistical ordering (*individual sufficiency*) 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 Blackwell's sense) for their joint information in the latter case. But the resulting

ordering 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.

How does this relate to a many player feasibility ordering? Suppose that we consider belief invariant decision rules, where a player's recommended action under the decision rule does not reveal more information to a player about the state and other players' signals than he had before the recommendation (belief invariance is discussed more formally in Section 3.4). Say that an information structure is more valuable than another if the set of outcomes (joint distribution over action profiles and states) that can be induced by belief invariant decision rules is larger - in all basic games - under the former information structure. Now Bergemann and Morris (2016a) show that there is an equivalence in the many player case between individual sufficiency and more valuable orderings. In Section 3.4, we will say that a decision rule is a *belief invariant Bayes correlated equilibrium* if it satisfies obedience and belief invariance. Now if we look at games of common interests (where players have identical payoffs), an information structure is more valuable than another if it gives a higher (common) payoff in the best belief invariant Bayes correlated equilibrium. One can show that our individual sufficiency ordering is equivalent to the more valuable than ordering. This result closely follows Lehrer, Rosenberg, and Shmaya (2010), who give a statistical characterizations of the more valuable than ordering under different versions of incomplete information correlated equilibrium.

A common observation is that in strategic situations, there is no many player analogue of Blackwell's ordering: see, for example, 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.

2.2.4 Metaphorical Information Design: Robust Predictions and Maxmin Objectives

Mechanism design sometimes has a literal interpretation. Thus - 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.¹¹

In both cases, the role of the "designer" may be metaphorical rather than literal. In our own applications of information design, we have mostly been interested in such problems. Consider an information designer who minimizes revenue in a first price auction, for a given symmetric distribution of values. By finding the solution to this problem, we are identifying a lower bound on revenue for the seller that holds across all information structures (Bergemann, Brooks, and Morris (2017)). Consider an information designer who seeks to maximize the variance of aggregate output in an economy where players face common and idiosyncratic shocks. By finding a solution to this problem, we identify an upper bound on the variance of aggregate output (Bergemann, Heumann, and Morris (2015)).

More generally, the set of Bayes correlated equilibria characterizes 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. In our work on third degree price discrimination (Bergemann, Brooks, and Morris (2015)), we characterize the set of welfare outcomes that can arise across all information structures (or market segmentations). Thus we show that there are essentially no robust predictions about welfare. On the other hand, it is a robust prediction in our work on auctions that revenue can at least attain a certain strictly positive lower bound. In Bergemann, Brooks, and Morris (2017), we highlight this

¹¹Forges and Koessler (2005) observe that conditioning of players' exogenous information makes sense if players' types are ex post verifiable.

interpretation of the information design problem.

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 is subject to imperfect perception 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 attributed to optimal costly information acquisition by the player in the presence of imperfect information.

3 Information Design with Private Information and the Relation to Mechanism Design

In the previous section, we 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 knows the state. Here we consider two alternative assumptions about the designer's ability to condition recommendations on players' prior information. If the designer does not know the players' prior information but can elicit their information, then we have *information design with elicitation*. If the designer does not know the players' information and cannot elicit it, we have *information design without elicitation*.

3.1 Information Design when Players' Prior Information is not known to the Designer

When the information designer cannot observe players' prior information, she may or may not be able to ask the players about it. In the former case of information design with elicitation, she will be able to condition the information that she provides on reports from players about their information. In the latter case of information design without elicitation, she can only send information that does not depend on players' prior information.

In the case of information design with elicitation, the revelation principle still implies that we can restrict attention to the case where the information sent by the information designer consists of action recommendations. However, we will now require an incentive compatibility condition that entails *truth-telling* as well as *obedience*, so that the information designer can only condition on a player's signal if

the player can be given an incentive to report it truthfully. Following Myerson (1991) (Section 6.3), we can think of the information designer choosing a decision rule $\sigma: T \times \Theta \rightarrow \Delta(A)$ but each player can choose a deviation $\delta_i: A_i \rightarrow A_i$ with the interpretation that $\delta_i(a_i)$ is the action chosen by player i if the information designer recommended action a_i . The decision rule σ is incentive compatible if each player does not have an incentive to deviate:

Definition 2 (Incentive Compatible)

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

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

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

The displayed inequality will be referred to as player i 's type- t_i incentive constraint. It ensures that player i , after observing signal t_i , finds it optimal to report his signal truthfully and then takes whichever action the designer recommends. 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 *and* disobeys the recommendation by the designer. Thus, incentive compatibility implies but is not implied by separately requiring truth-telling *and* obedience.

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

In the case of information design without elicitation, the designer cannot condition on players' types, but can offer a vector of recommendations to each player, where the individual entry is the recommendation to a specific type of the player. We denote player i 's set of pure strategies by $B_i = A_i^{T_i}$ with generic element $b_i: T_i \rightarrow A_i$. We denote by $B = \times_{i=1}^I B_i$ and a generic element is given by $b = (b_1, \dots, b_I) \in B$. We will abuse notation somewhat by writing $b(t)$ or $(b_i(t_i), b_{-i}(t_{-i}))$ to mean $(b_1(t_1), \dots, b_I(t_I)) \in A$.

We are interested in strategy recommendations $\phi: \Theta \rightarrow \Delta(B)$.

Definition 3 (Public Feasibility)

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

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

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

When $I = 1$, public feasibility is a vacuous restriction. Every decision rule σ is induced by the strategy 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 strategy 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 strategies, b , rather than actions, a .

Definition 4 (Publicly Feasible Obedience)

A decision rule $\sigma: T \times \Theta \rightarrow \Delta(A)$ is publicly feasible obedient if it is publicly feasible and the associated strategy recommendation $\phi: \Theta \rightarrow \Delta(B)$ satisfies obedience in the sense that for each $i = 1, \dots, I$, $t_i \in T_i$, and $b_i \in B_i$,

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

for all $a'_i \in A_i$.

The displayed inequality will be referred to as player i 's (t_i, b_i) - publicly feasible obedience constraint. It ensures that player i , after observing signal t_i and receiving the recommendation b_i , finds it optimal to take the action $b_i(t_i)$ prescribed by the strategy b_i for his type t_i .

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

3.2 The Investment Example Re-Visited

We reconsider the investment example introduced earlier in Section 2.1 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 elicit it. So we have a screening problem where the government offers a recommendation which induces a probability of investing as a function of the reported signal and the true state. Kolotilin, Li, Mylovanov, and Zapechelnuyk (2015) refer to this informational environment as "private persuasion".¹² As noted above, we have three sets of constraints 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.

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 2.1.2, a decision rule is now a vector $p_{\theta t} = (p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg})$. 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 2.1.2 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:

$$qp_{Gg}x - (1 - q)p_{Bg} \geq qp_{Gb}x - (1 - q)p_{Bb} \quad (15)$$

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

$$(1 - q)p_{Gb}x - qp_{Bb} \geq (1 - q)p_{Gg}x - qp_{Bg}. \quad (16)$$

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:

$$\frac{1 - q}{q}x(p_{Gg} - p_{Gb}) \leq (p_{Bg} - p_{Bb}) \leq \frac{q}{1 - q}x(p_{Gg} - p_{Gb}). \quad (17)$$

These inequalities are useful to highlight how in the bad state the differential in the recommendation for bad and good type are bounded, below and above, by the differential in the recommendation in the good state. Notice also that since

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

¹²The name "private persuasion" is motivated by an alternative interpretation: there is a continuum population corresponding to the player, and types correspond to individuals in the population. Bergemann, Bonatti, and Smolin (2015) analyze this environment with monetary transfers.

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.¹³

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 truth-telling constraints impose restrictions on how the differences in the conditional probabilities across types can vary across states. These impose additional restrictions on the ability of the government to attain either very low or very high investment probabilities in both states as highlighted by equation (17).

Figure 5 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 5: INVESTMENT PROBABILITY WITH PRIVATE INFORMATION.

We could also consider a government, who *does not* know the signal of the firm and *cannot* even elicit it. Kolotilin, Li, Mylovanov, and Zapechelnuyk (2015) 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. Kolotilin, Li, Mylovanov, and Zapechelnuyk (2015) showed such an equivalence under a different set of assumptions.

¹³As the difference in the probability of investing has to be nonnegative in both states, it can be shown that the possibility of double deviations does not impose any additional restriction on the behavior of the player. This is a special feature of the present binary action, binary state environment.

The equivalence breaks down if the assumptions are relaxed. We illustrate this failure of the equivalence result with a minor generalization of the investment example. In particular, in a single player environment, we allow the player to either consider a small or a large investment. For completeness, we present examples where one of the other two hypotheses fails in the appendix.

We end this section with an extension of the earlier investment example that allows us to observe the proper nesting between the set of outcomes without prior information, with prior information but 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.

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, in our subsequent decision 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 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:

$$p_G x - p_B \geq p_G y - \frac{1}{2} p_B, \tag{18}$$

and one for the no investment against the small investment:¹⁴

$$0 \geq (1 - p_G)y - \frac{1}{2}(1 - p_B). \quad (19)$$

The equilibrium regions are depicted in Figure 6. 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 consider the three different communication protocols for the government. An omniscient designer faces the obedience constraints that we analyzed earlier in Section 2.1.2, except that 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 (15) and (16). Finally, a designer without elicitation faces additional obedience constraints that rule out deviations conditional on a particular strategy recommendation. The corresponding areas in Figure 6 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 the first three regimes offer an increasing number of constraints. It remains to discuss the impact of elicitation. 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. 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 good type can disobey the recommendation to invest by deviating to invest small only when the designer *also* recommends not to invest to the bad type. Likewise, the bad type can disobey the recommendation not to invest by deviating to invest small only when the designer *also* recommends to invest to the high type. The additional options for the player induce further constraints on the information designer. Naturally, these additional deviations were not available in the binary action environment. And in fact the absence of these large set of deviations accounts for the equivalence between elicitation and no elicitation in the binary action and state environment.¹⁵

¹⁴The 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.

¹⁵We 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.

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 6: INVESTMENT PROBABILITY UNDER DIFFERENT INFORMATION DESIGN SCENARIOS.

3.3 General Mechanism Design

In the introduction, we contrasted information design with mechanism design as follows:

Consider the problem of an "information designer" who 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 "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)). In this broad vision, players may have control over some actions effecting 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 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.

3.4 Incomplete Information Correlated Equilibrium

We can also relate information design to old and new work on incomplete information correlated equilibrium. This literature describes how different communication and recommendation protocols under private information affect the set of the attainable equilibrium outcomes. For our purpose, 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. To conclude a discussion of incomplete information correlated equilibrium, we will discuss two further feasibility conditions.

3.4.1 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 $\sigma : T \times \Theta \rightarrow \Delta(A)$ 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) have shown that the solution concept of interim correlated rationalizability characterizes the implications of common certainty of rationality and players' beliefs and higher order beliefs. The solution concept builds in 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 subjective incomplete information correlated equilibrium. If the common prior assumption is then imposed, then this corresponds to 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 (2016) 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.

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

Thus join feasibility 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. 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" and Forges (2006) introduced a sixth one. The feasibility and incentive conditions described so far allow us to completely describe the six solution concepts she discusses:

1. A *Bayesian solution* is a decision rule satisfying joint feasibility and obedience.
2. A *belief invariant Bayesian solution* is a decision rule satisfying joint feasibility, belief invariance and obedience.
3. A *agent normal form correlated equilibrium* is a decision rule satisfying joint feasibility, public feasibility (which implies belief invariance) and obedience.
4. A *communication equilibrium* is a decision rule satisfying joint feasibility and incentive compatibility (which implies obedience)
5. A *strategic form correlated equilibrium* is a decision rule satisfying joint feasibility and publicly feasible obedience (which implies belief invariance, public feasibility, obedience and incentive compatibility).

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

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

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

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

Forges (1993) reports examples showing that these inclusions are the only ones that can be shown, i.e., there exist decision rules that (i) are Bayesian solutions but not belief invariant BCE or communication

equilibria; (ii) are belief invariant Bayes solutions but not a communication equilibria or an 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 equilibria corresponds - in our language - to the set of Bayes correlated equilibria that would arise under joint feasibility if players had no information.

4 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.

4.1 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 (2016) 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 (2016) formalize and analyze the case where uninformed players 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 $S = \left((T_i)_{i=1}^I, \pi \right)$. The information designer sends each player i an extra message $m_i \in M_i$, according to rule $\phi : T \times \Theta \rightarrow \Delta(M)$, where $M = M_1 \times \dots \times M_I$. A communication rule is then $C = \left((M_i)_{i=1}^I, \phi \right)$. Now the basic game G , the prior information structure S and the communication rule C describe a Bayesian game (G, S, C) .¹⁶ A strategy for player i in this game is a mapping $b_i : T_i \times M_i \rightarrow \Delta(A_i)$. A communication rule C and

¹⁶Bergemann and Morris (2016b) call the pair (S, C) an "expanded information structure".

strategy profile b will now induce a decision rule

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

We will write $E_{G,S}(C)$ for the set of Bayes Nash equilibria of the game with communication rule C . We can now give a more formal statement of Proposition 1:

Proposition 4

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

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

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

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

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

$$\max_C \max_{b \in E_{G,S}(C)} V_{G,S}^*(C, b).$$

Proposition 4 established that

$$\max_C \max_{b \in E_{G,S}(C)} V_{G,S}^*(C, b) = \max_{\sigma \in BCE(G,S)} V_{G,S}(\sigma).$$

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

$$\max_C \min_{b \in E_{G,S}(C)} V_{G,S}^*(C, b).^{17}$$

¹⁷Of course, the minmax problem

$$\min_C \max_{b \in E_{G,S}(C)} V_{G,S}^*(C, b).$$

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

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).

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.¹⁸ We will consider a slightly adapted version of the incomplete information investment game discussed earlier with payoffs:

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

for some $0 < x < 1$, where player 1 knows the true state 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

¹⁸See also Hoshino (2017) on this connection.

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 $\Omega = \{1, 2, \dots, \dots, \infty\}$ where player 1 observes the partition $(\{1\}, \{2, 3\}, \{4, 5\}, \dots, \{\infty\})$ and player 2 observes the partition $(\{1, 2\}, \{3, 4\}, \dots, \{\infty\})$. 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 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 applications, correspond to the omniscient case where the information designer can condition on players' prior information as well as on the state. An alternative case that has been studied is when the information designer can only send public signals and only condition on the state (and not players' prior information). Goldstein and Huang (2016) and Inostroza and Pavan (2017) have studied this problem in global game models of regime change (Inostroza and Pavan (2017) call this the 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.

4.2 Adversarial Mechanism Design

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). Bergemann, Brooks, and Morris (2016a) identify circumstances where these are equal. Du (2016) constructs elegant bounds for the latter problem. Both problems are studied *without* the common

prior assumption by Chung and Ely (2007).

5 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 3.3 - 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), Dworzak (2016), Carroll and Segal (2016) and Bergemann, Brooks, and Morris (2016b).

6 Appendix

Additional Computation for Section 2.1.2 We observed in Section 2.1.1 that absent any information the firm does not invest. For the private information of the firm to change the "default" behavior under no private information, it has to be that the firm is investing after receiving the good signal, or that

$$qx + (1 - q)(-1) \geq 0 \Leftrightarrow x \geq \frac{1 - q}{q} \Leftrightarrow q \geq \frac{1}{1 + x}. \quad (20)$$

In other words, the information q has to be sufficiently precise to induce a change in the behavior if absent any information the expected payoff from investing exceeded the expected payoff from not investing.

Now, as the information designer is adopting a recommendation policy $(p_{Bb}, p_{Bg}, p_{Gb}, p_{Gg})$, the firm will have an incentive to invest (when told to invest) if

$$-\frac{1}{2}(1 - q)p_{Bg} + \frac{1}{2}qp_{Gg}x \geq 0 \Leftrightarrow p_{Gg} \geq \frac{1 - q}{q} \frac{p_{Bg}}{x}, \quad (21)$$

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

$$0 \geq -\frac{1}{2}(1 - q)(1 - p_{Bg}) + \frac{1}{2}q(1 - p_{Gg})x \Leftrightarrow p_{Gg} \geq \frac{1 - q}{q} \frac{p_{Bg}}{x} + 1 - \frac{1 - q}{q} \frac{1}{x}. \quad (22)$$

The above two incentive constraints pertain to the recommendations conditional on having a good type g . A similar pair of incentive constraints apply to the recommendations conditional on having a bad type b .

As long as the private information of the firm is sufficiently noisy, or $q \leq 1/(1 + x)$, the binding constraint is (21) as in the uninformed case; otherwise it is the inequality (22) 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, $(1 - q)/qx$, 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 for the good state is achieved if there is zero probability of investing in the bad state as derived earlier in (9).

Additional Results for Section 3.2 For a given Bayesian game (G, S) , let $E(G, S)$, respectively, $NE(G, S)$, denote the set of decision rules that can be attained with and without elicitation, respectively.

Proposition 5

For each (G, S) , we have

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

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

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

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

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

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

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

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

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

Proposition 6

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

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

Proof. By Proposition 5, it suffices to prove $NE(G, S) \supseteq E(G, S)$. First, we simplify the notation. Label the states and actions so that $\Theta = \{G, B\}$ and $A = \{0, 1\}$. If either action is weakly dominant, the desired result can be verified by directly computing $NE(G, S)$ and $E(G, S)$. Therefore, we assume $u(1, G) - u(0, G)$ and $u(1, B) - u(0, B)$ are each nonzero and have opposite signs. Then without loss, we may assume the payoffs take the form $u(0, G) = u(0, B) = 0$, $u(1, B) = -1$ and $u(1, G) = x > 0$.¹⁹ Action 1 can be interpreted as investment. We will represent a decision rule σ by a vector $p = (p_{\theta t})_{(\theta, t) \in \Theta \times T}$, where $p_{\theta t} = \sigma(1|\theta, t)$. For each signal $t \in T$, let

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

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

Let $p = (p_{\theta t}) \in E(G, S)$. To show that $p \in NE(G, S)$, we will explicitly construct an obedient strategy recommendation ϕ that induces p . Let $t, t' \in T$ and set $q = q(t)$ and $q' = q(t')$. The incentive compatibility of p implies

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

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, \dots, t_n so that

$$q(t_1) \leq \dots \leq q(t_n) \quad \text{and} \quad p_{\theta t_1} \leq \dots \leq p_{\theta t_n} \quad \text{for } \theta = B, G. \quad (23)$$

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

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

¹⁹First, swap the labels G and B if needed to obtain $u(1, G) - u(0, G) > 0$. Then rescale the utility function so that $u(1, G) - u(0, G) = 1$. Finally, translate the functions $u(\cdot, G)$ and $u(\cdot, B)$ separately so that $u(0, G) = u(0, B) = 0$. The separate translations may change the agent's preferences over states but not over actions.

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

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

By (23), $p_{\theta t_k} - p_{\theta t_{k-1}} \geq 0$, 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 $\ell = 1, \dots, n$ and $k = 1, \dots, n+1$, type t_ℓ 's expected utility from investing, conditional on being recommended b^k , is

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

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

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

where the last inequality holds by incentive compatibility. Similarly, for types t_ℓ with $\ell < k$,

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

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

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

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

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

$$\Delta_g^1 \triangleq qp_{Gg}x - (1 - q)p_{Bg} - \left(qp_{Gg}y - \frac{1 - q}{2}p_{Bg} \right) \geq 0, \quad (24a)$$

$$\Delta_b^1 \triangleq (1 - q)p_{Gb}x - qp_{Bb} - \left((1 - q)p_{Gb}y - \frac{q}{2}p_{Bb} \right) \geq 0. \quad (24b)$$

$$\Delta_g^0 \triangleq q(1 - p_{Gg})y - \frac{1 - q}{2}(1 - p_{Bg}) \leq 0, \quad (24c)$$

$$\Delta_b^0 \triangleq (1 - q)(1 - p_{Gb})y - \frac{q}{2}(1 - p_{Bb}) \leq 0. \quad (24d)$$

We 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. An information designer with elicitation faces four additional constraints ruling out non-truthful deviations:

$$qp_{Gg}x - (1 - q)p_{Bg} \geq qp_{Gb}x - (1 - q)p_{Bb}, \quad (25a)$$

$$qp_{Gg}x - (1 - q)p_{Bg} \geq q(p_{Gb}x + (1 - p_{Gb})y) - (1 - q)\frac{1 + p_{Bb}}{2}, \quad (25b)$$

$$(1 - q)p_{Gb}x - qp_{Bb} \geq (1 - q)p_{Gg}x - qp_{Bg}, \quad (25c)$$

$$(1 - q)p_{Gb}x - qp_{Bb} \geq (1 - q)p_{Gg}y - \frac{q}{2}p_{Bg}. \quad (25d)$$

Again, it is sufficient to consider a smaller class of deviations because the high type finds investment more attractive than the low type does. Formally, $E(G, S)$ is the set of $p \in [0, 1]^4$ satisfying (24a)-(25d).

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

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

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

$$q(p_{Gg} - p_{Gb})x - (1 - q)(p_{Bg} - p_{Bb}) \geq q(p_{Gg} - p_{Gb})y - \frac{1 - q}{2}(p_{Bg} - p_{Bb}), \quad (26a)$$

$$(1 - q)(p_{Gg} - p_{Gb})y - \frac{q}{2}(p_{Bg} - p_{Bb}) \leq 0. \quad (26b)$$

Formally, $NE(G, S)$ is the set of decision rules in $E(G, S)$ satisfying (26a) and (26b).

After some algebra, we can see that (26a) is equivalent to

$$\Delta_g^1 \geq \Delta_b^1 + p_{Gb}(2q - 1)(x - y) + p_{Bb}(2q - 1)/2. \quad (27)$$

When p puts positive probability on investing after a bad signal, (27) eliminates decision rules for which (24a) has little slack. Similarly, (26b) is equivalent to

$$\Delta_b^0 \leq \Delta_g^0 - (1 - p_{Gg})(2q - 1)y - (1 - p_{Bg})(2q - 1)/2. \quad (28)$$

When p puts positive probability on not investing after a good signal, (28) eliminates decision rules for which (24d) has little slack.

Example 1 (Two Agents) *Suppose $I = 2$, $\Theta = \{B, G\}$, $A_1 = A_2 = \{\text{invest, not invest}\}$, and $u_i(a_i, a_{-i}, \theta) = u(a_i, \theta)$ with u as in the opening example given by (3). 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 strategy 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 (3), but now split the bad state into two bad states B_1 and B_2 , each with prior probability $1/4$ and the same payoffs as in state B of the original example. Suppose the agents receive a completely uninformative binary signal t taking values t_1 and t_2 with equal probability. Consider the following decision rule: type t_i invests precisely in states G and B_i . For $x \in (1/2, 1)$, this decision rule is incentive compatible. It is uniquely induced by recommending $(b(t_1), b(t_2)) = (\text{invest, invest})$ in state G ; $(b(t_1), b(t_2)) = (\text{invest, not invest})$ in state B_1 ; and $(b(t_1), b(t_2)) = (\text{not invest, invest})$ in state B_2 . However, this strategy 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

- ARIELI, I., AND Y. BABICHENKO (2016): “Private Bayesian Persuasion,” Discussion paper, Technion.
- AUMANN, R., AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*. MIT Press.
- BASSAN, B., O. GOSSNER, M. SCARSINI, AND S. ZAMIR (2003): “Positive value of information in games,” *International Journal of Game Theory*, 32(1), 17–31.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2015): “Designing and Pricing Information,” Discussion paper, Yale University and MIT and University of Bonn.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–957.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2016a): “Informationally Robust Auction Design,” Discussion paper, Yale University, University of Chicago and Princeton University.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (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 (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): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.

- (2016c): “Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium,” *American Economic Review Papers and Proceedings*, 106(5), 586–591.
- 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.
- 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.
- DEKEL, E., D. FUDENBERG, AND S. MORRIS (2007): “Interim Correlated Rationalizability,” *Theoretical Economics*, 2, 15–40.
- 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.

- 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(1), 215–260.
- 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,” *The Review of Economic Studies*, 84(1), 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.
- 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, Penn 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., M. LI, T. MYLOVANOV, AND A. ZAPECHELNYUK (2015): “Persuasion of Privately Informed Receiver,” Discussion paper.
- 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.
- 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 (2016): “Information Design: The Belief-Hierarchy Approach,” 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.
- 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(1), 45–78.
- MYERSON, R. (1982): “Optimal Coordination Mechanism in Generalized Principal-Agent Problems,” *Journal of Mathematical Economics*, 10, 67–81.
- MYERSON, R. (1983): “Mechanism design by an informed principal,” *Econometrica*, 51, 1767–1797.
- MYERSON, R. (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.
- MYERSON, R. (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.
- 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. SONNENSCHNEIN (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.
- RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior under ‘Almost Common Knowledge’,” *American Economic Review*, 79, 385–391.
- TANEVA, I. (2015): “Information Design,” Discussion paper, University of Edinburgh.
- UI, T. (2001): “Robust equilibria of potential games,” *Econometrica*, 69(5), 1373–1380.

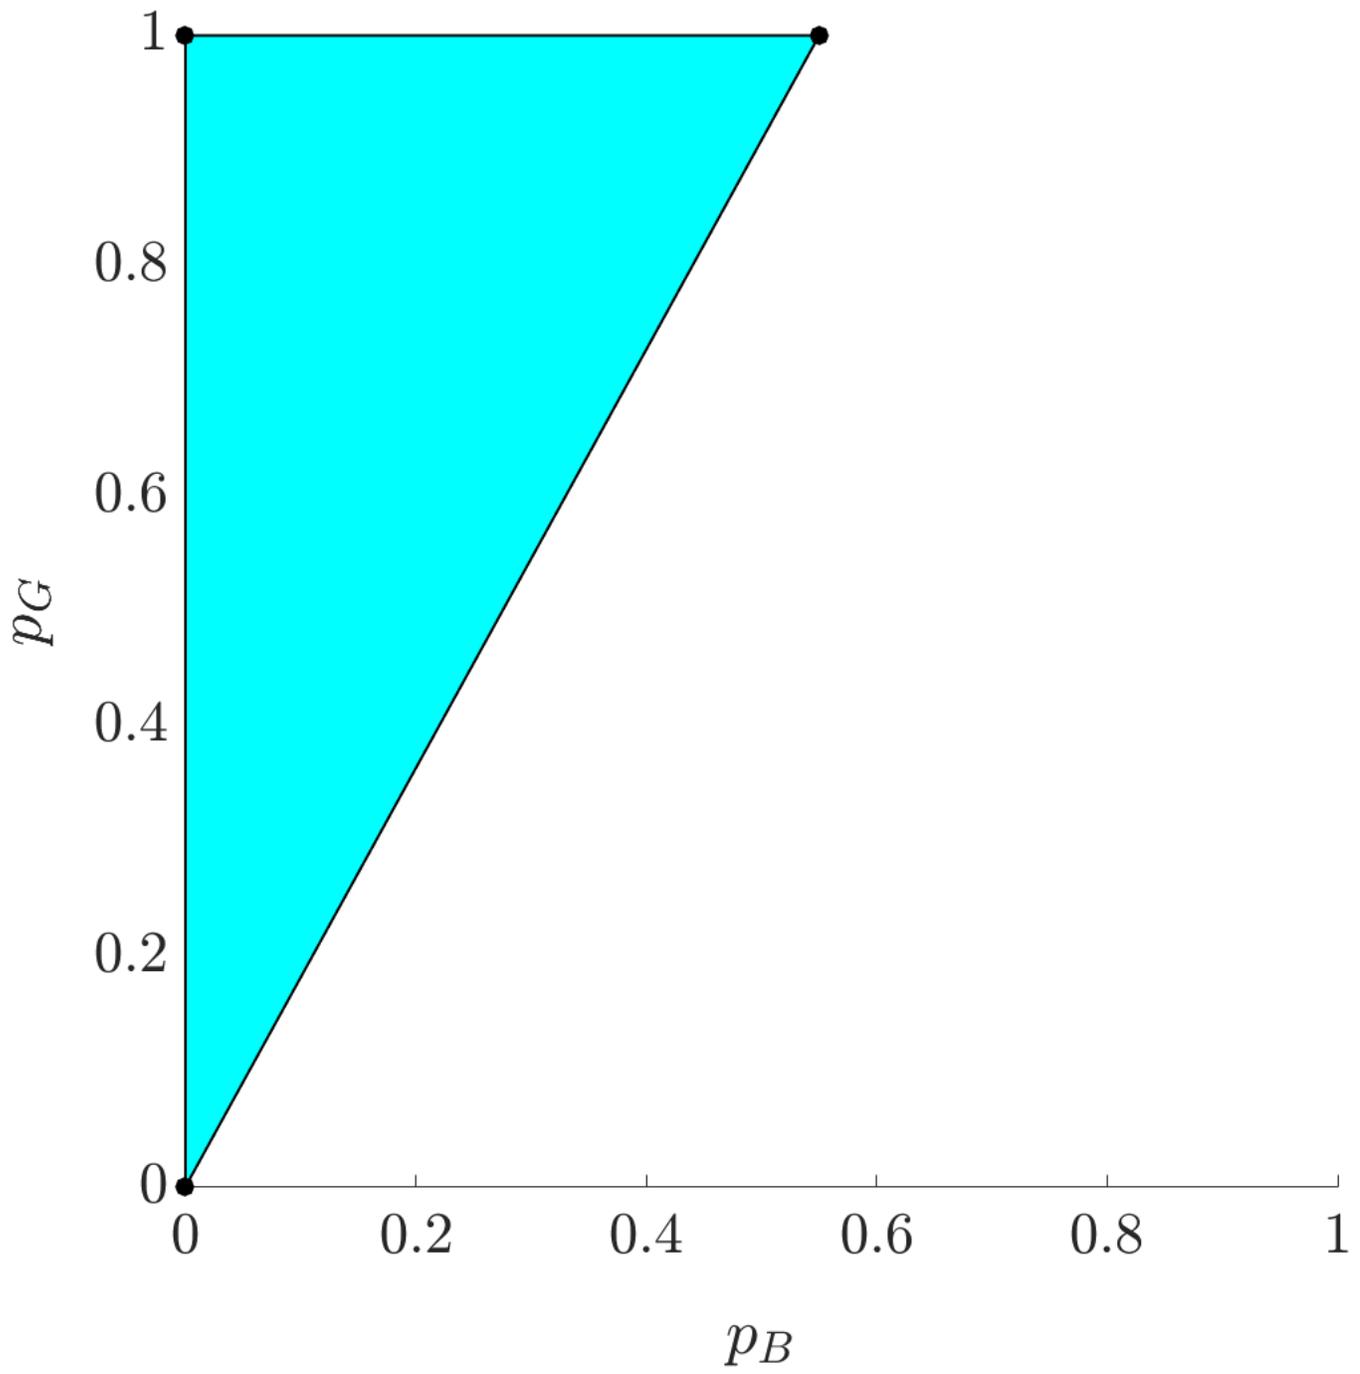


Figure 1: Investment Probabilities with Uninformed Player: $x=55/100$

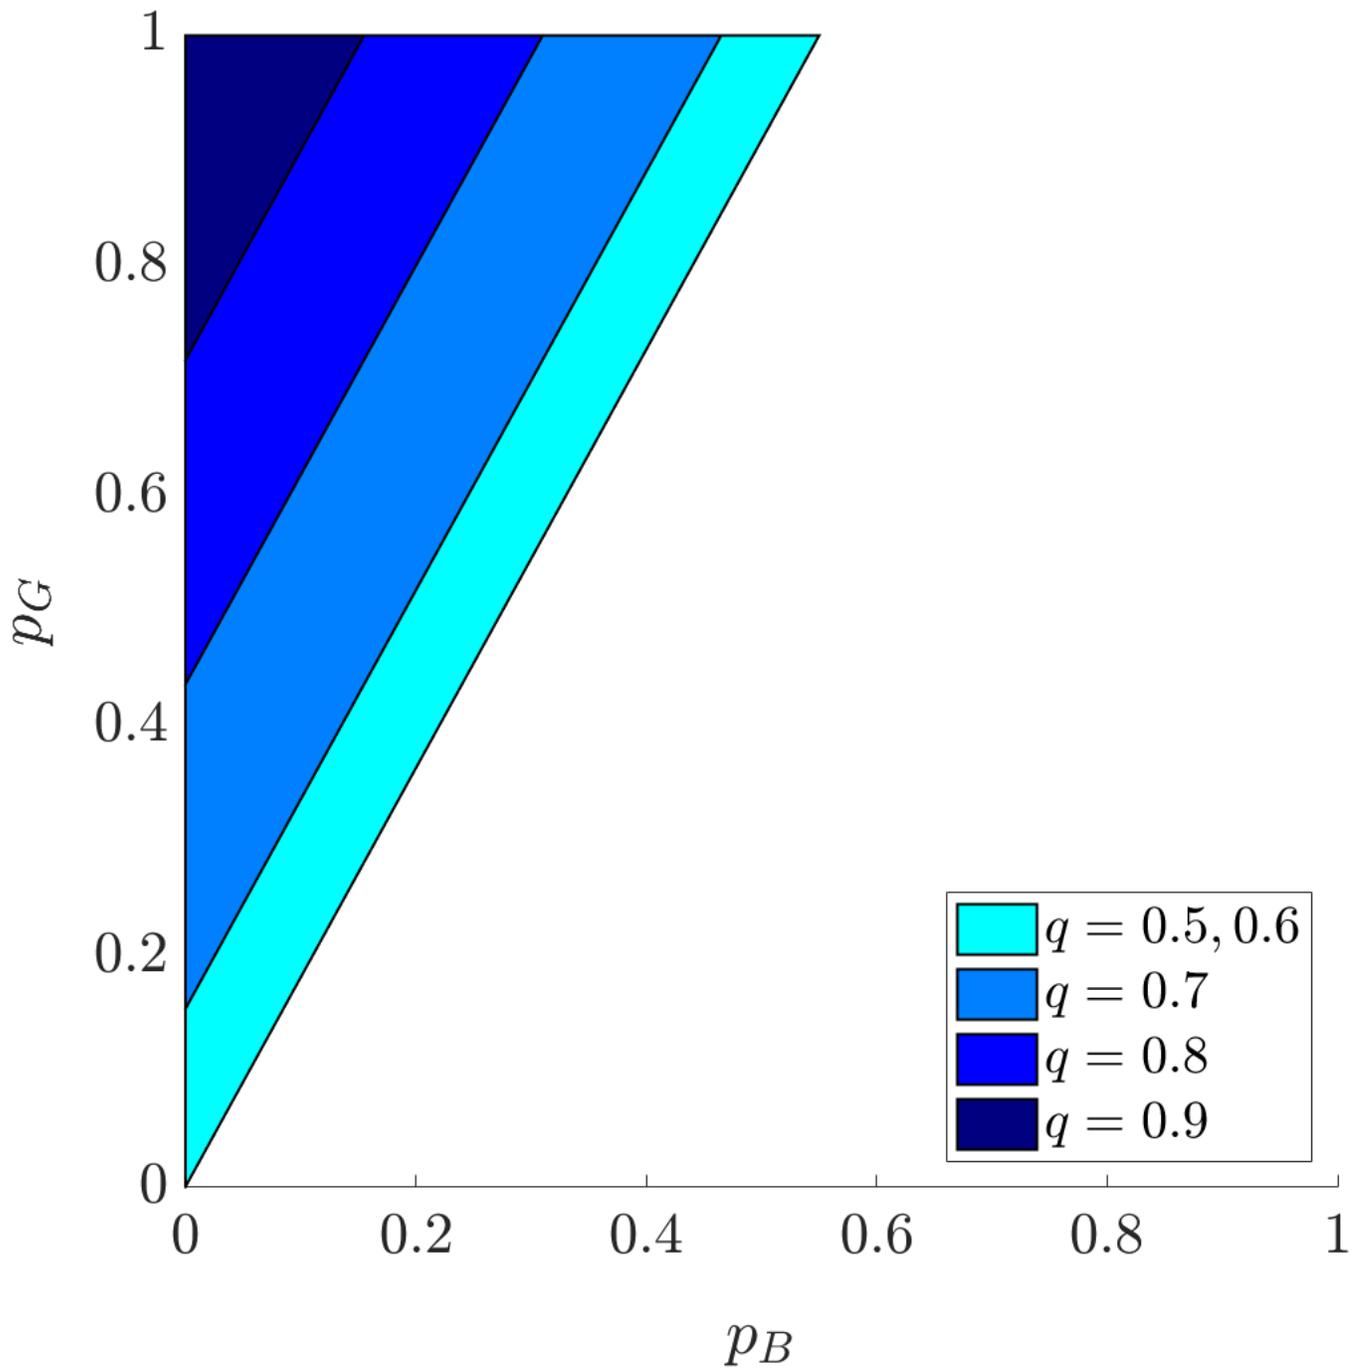


Figure 2: Investment Probability with Informed Player: $x=55/100$.

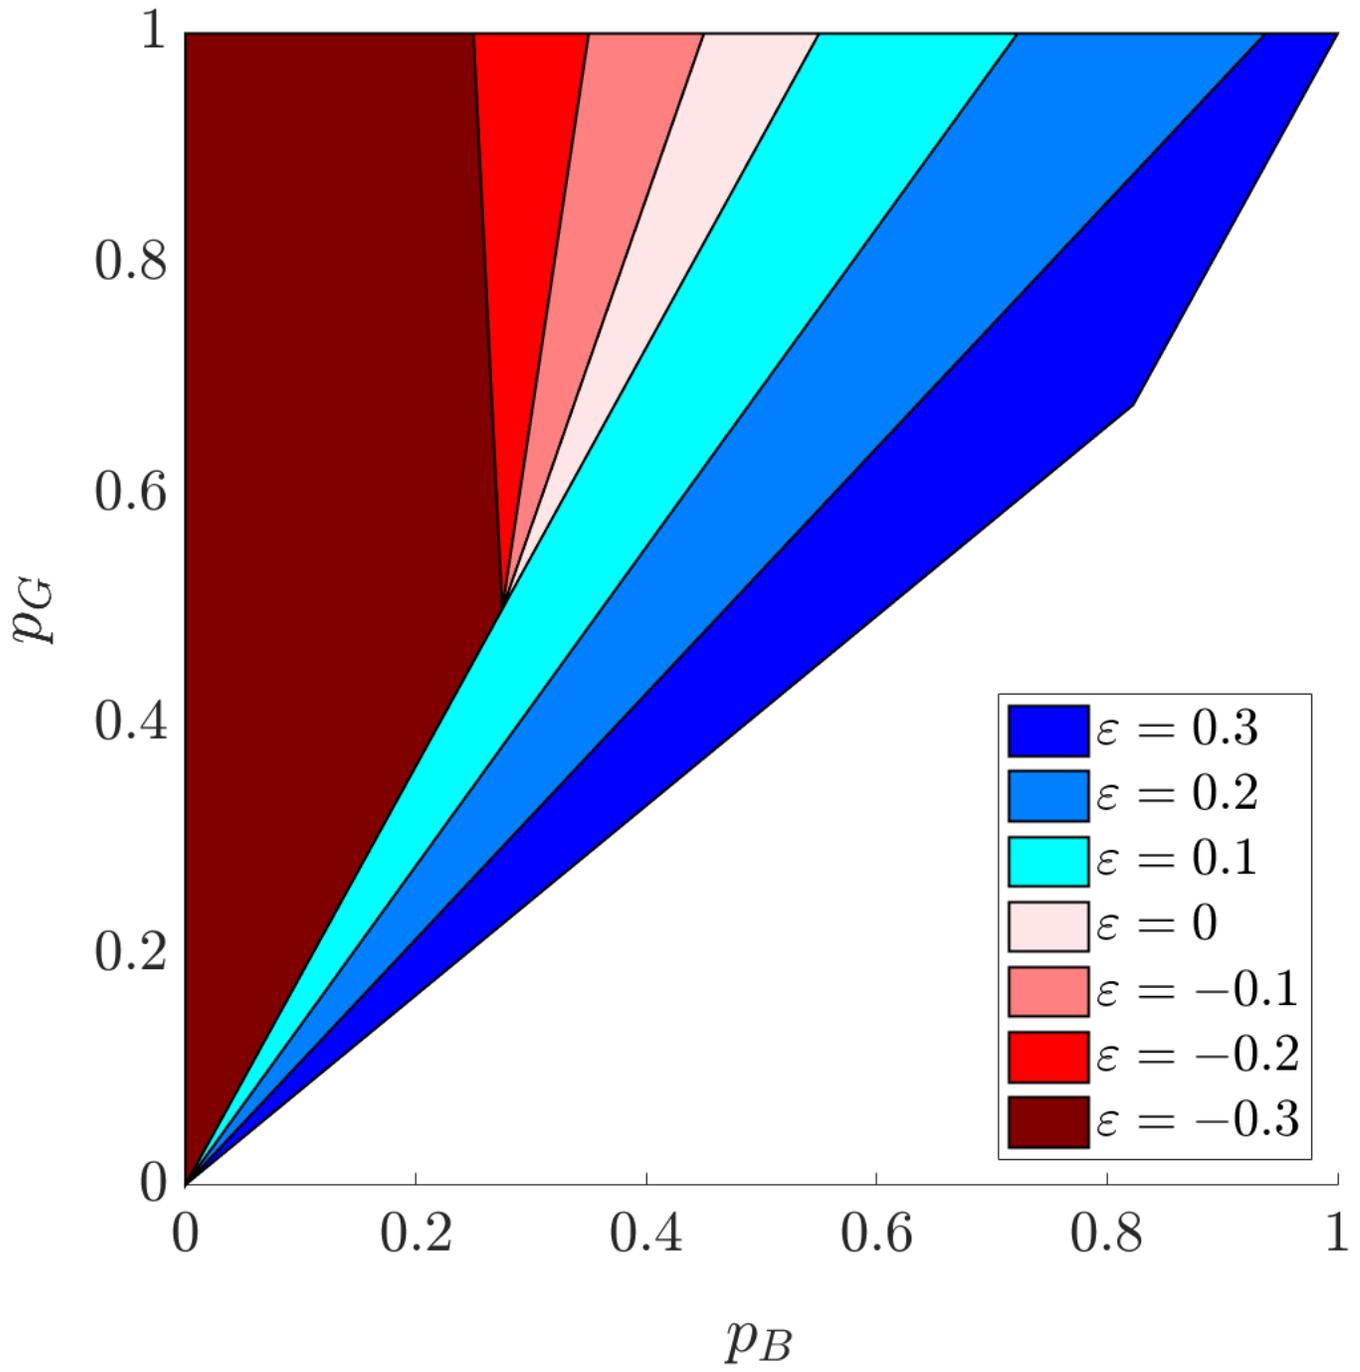


Figure 3: Investment Probability with Negative or Positive Strategic Term ϵ .

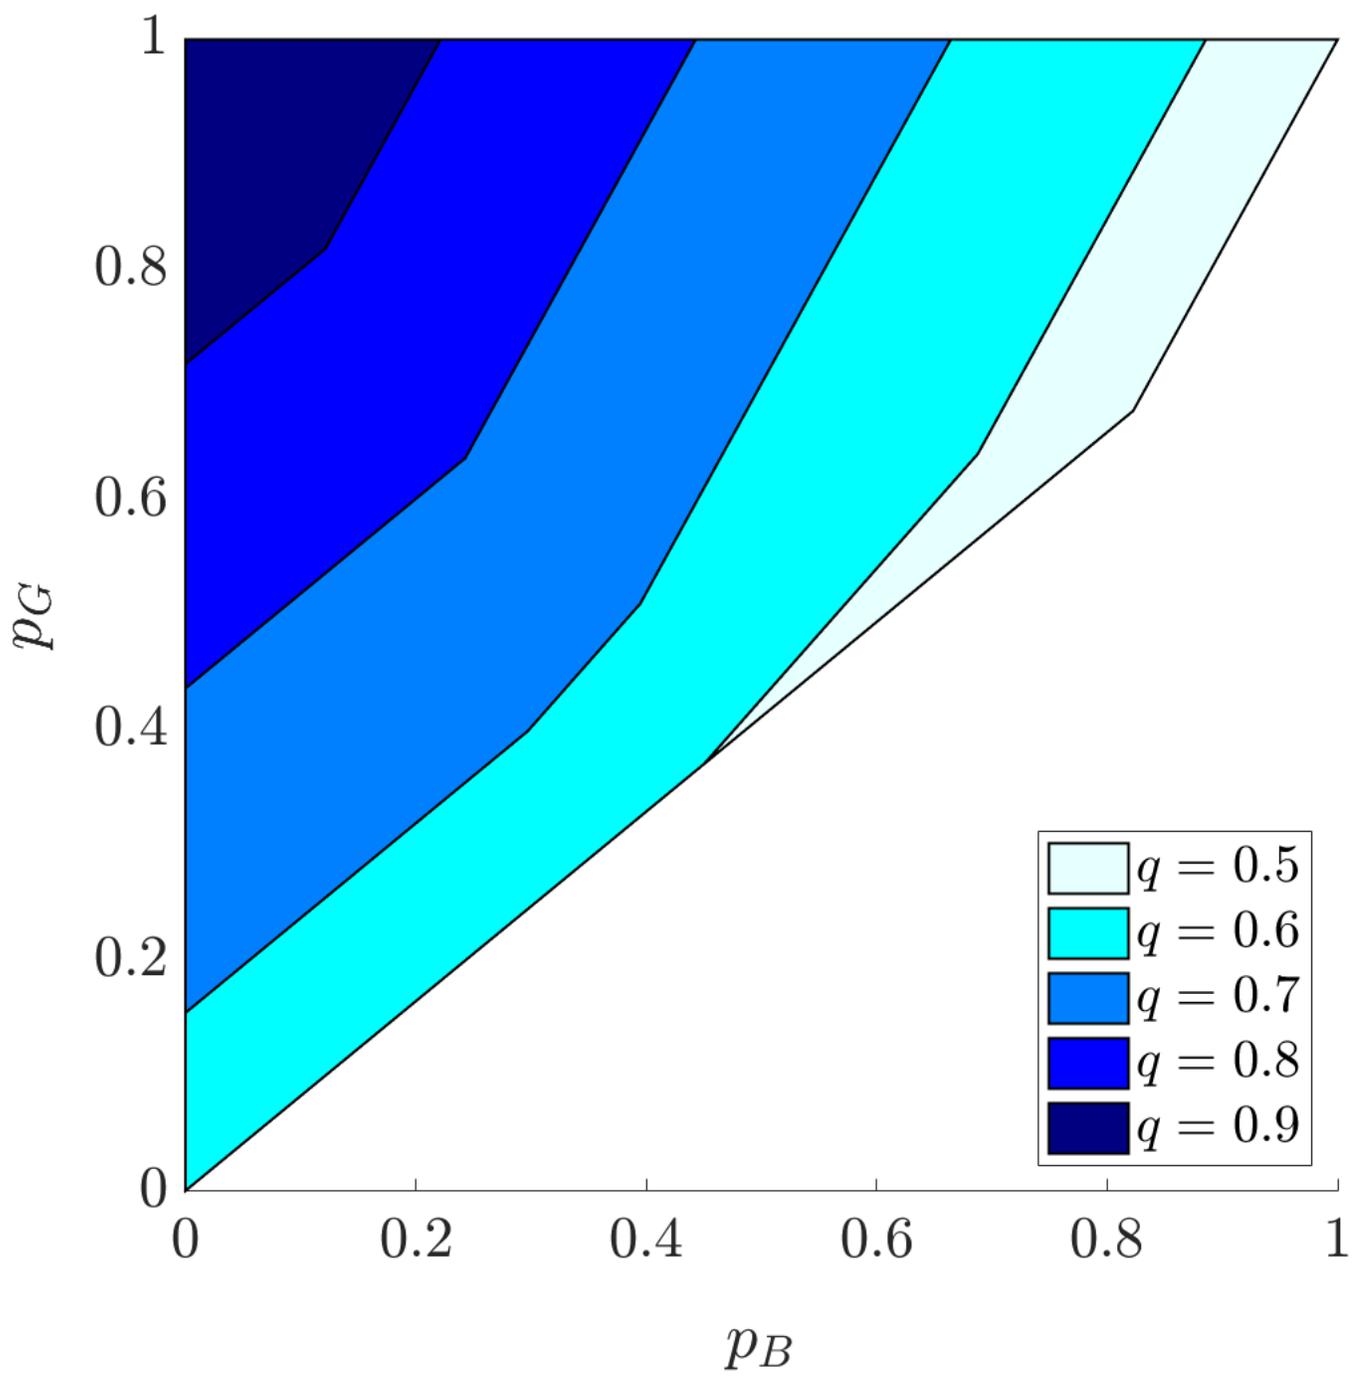


Figure 4: Investment Probability with Two Players with Prior Information, with Strategic Term $\varepsilon=3/10$.

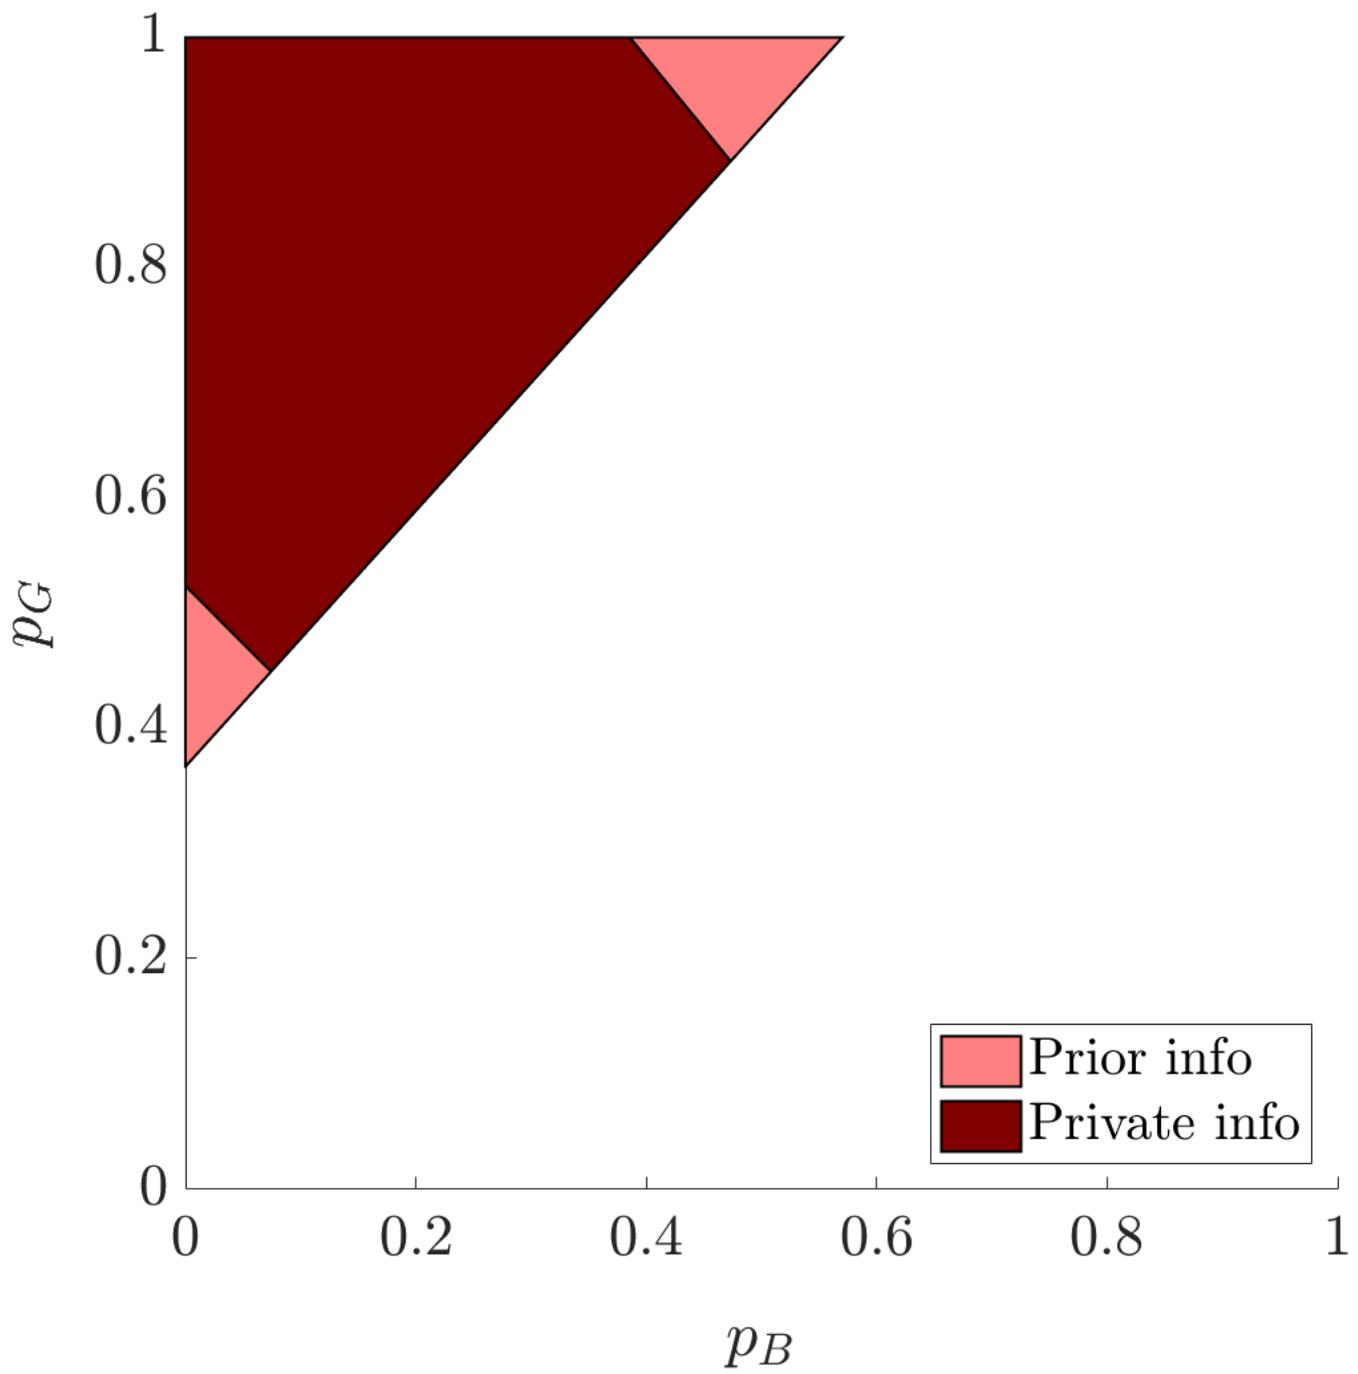


Figure 5: Investment Probability with Private Information

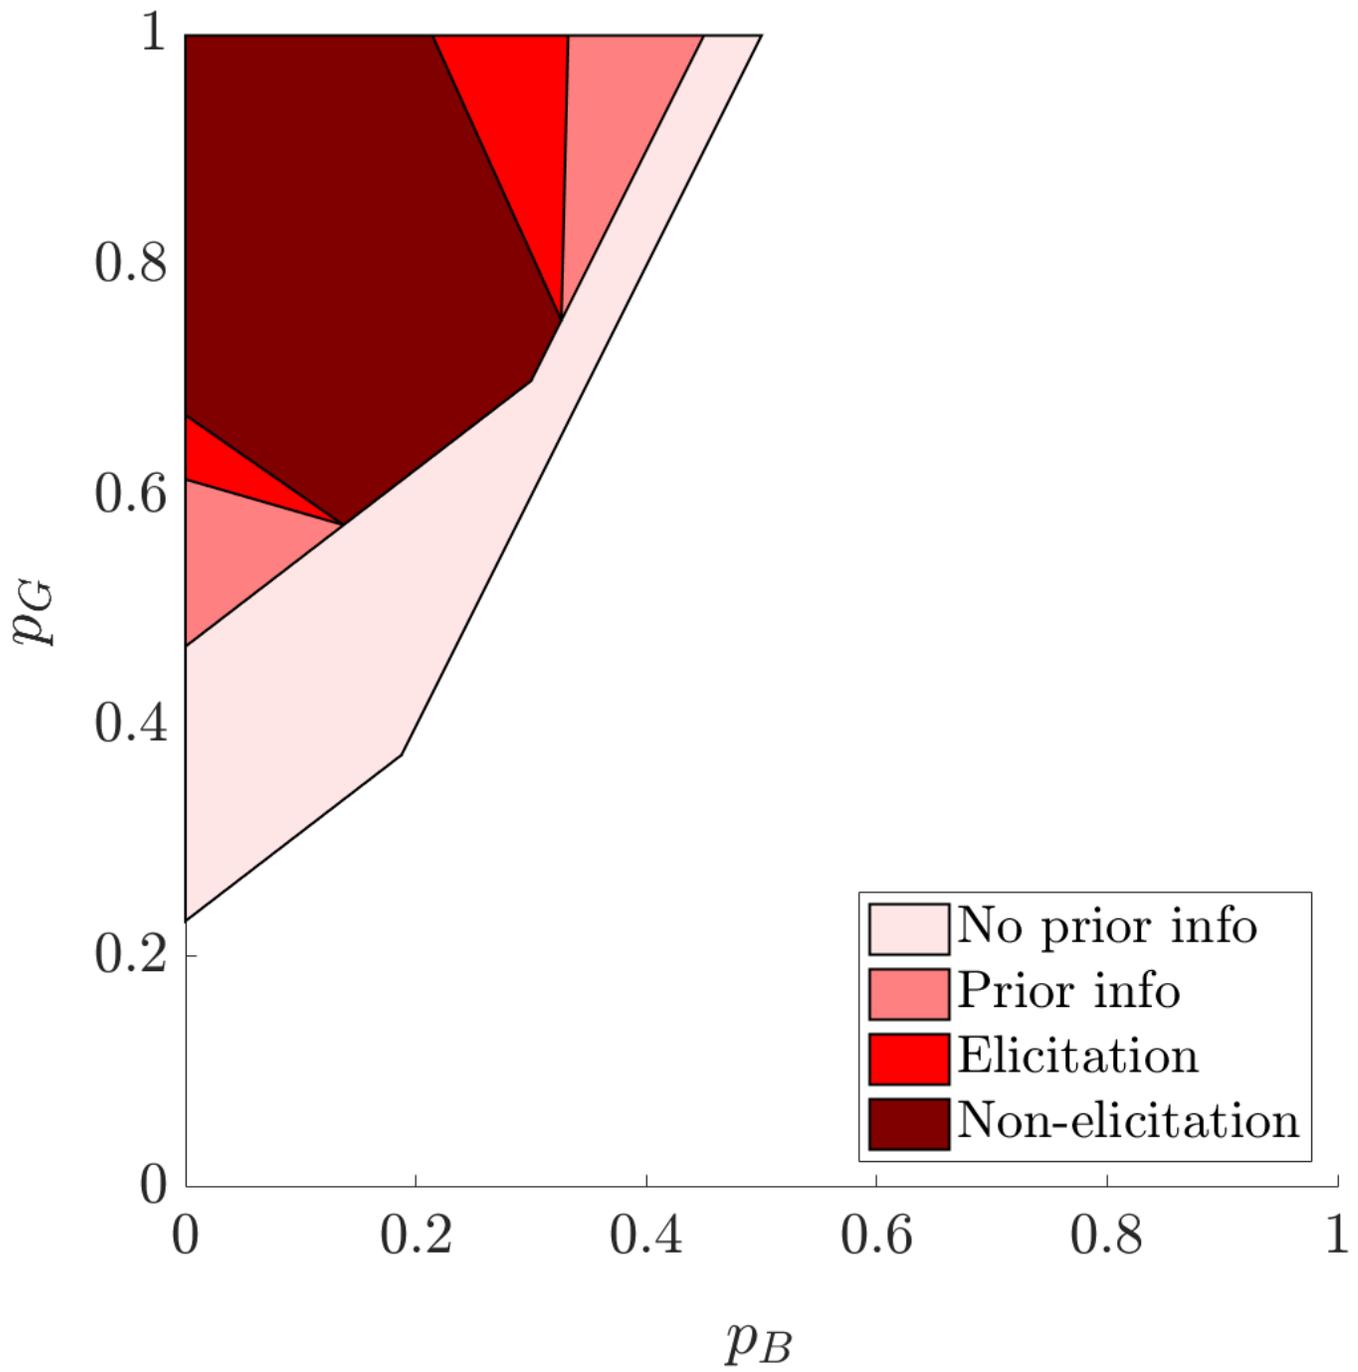


Figure 6: Investment Probability under Different Information Design Scenarios.