

**BAYES CORRELATED EQUILIBRIUM AND THE COMPARISON
OF INFORMATION STRUCTURES IN GAMES**

By

Dirk Bergemann and Stephen Morris

September 2013

Revised April 2015

COWLES FOUNDATION DISCUSSION PAPER NO. 1909RRR



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY**

Box 208281

New Haven, Connecticut 06520-8281

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

Bayes Correlated Equilibrium

and

the Comparison of Information Structures in Games*

Dirk Bergemann[†]

Stephen Morris[‡]

April 15, 2015

Abstract

A game of incomplete information can be decomposed into a basic game and an information structure. The basic game defines the set of actions, the set of payoff states the payoff functions and the common prior over the payoff states. The information structure refers to the signals that the players receive in the game.

We characterize the set of outcomes that can arise in Bayes Nash equilibrium if players observe the given information structure but may also observe additional signals. The characterization corresponds to the set of (a version of) incomplete information correlated equilibria which we dub Bayes correlated equilibria.

We identify a partial order on many player information structures (*individual sufficiency*) under which more information shrinks the set of Bayes correlated equilibria. This order captures the role of information in imposing (incentive) constraints on behavior.

KEYWORDS: Correlated equilibrium, incomplete information, Bayes Nash equilibrium, Bayes correlated equilibrium, robust predictions, information structure, sufficiency, Blackwell ordering.

*We acknowledge financial support through NSF Grants SES 0851200 and ICES 1215808. Earlier versions (Bergemann and Morris [2011]) were circulated under the title "Correlated Equilibrium in Games of Incomplete Information." We are grateful for discussions with Ben Brooks, with whom we are working on related projects, and to Hanming Fang and Ilya Segal, with whom we discussed related ideas for possible joint work in the past. We have benefited from comments from seminar participants and also from Josh Cherry, Mark Dean, Françoise Forges, Matthew Gentzkow, Olivier Gossner, Atsushi Kajii, Emir Kamenica, Qingmin Liu, Daniel Martin, Lones Smith, Eran Shmaya, Bruno Strulovici, Jonathan Weinstein, Alex Wolitzky and especially Marcin Peski. We acknowledge valuable research assistance from Constantinos Kalfarentzos, Alex Smolin and Áron Tóbiás.

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

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

1 Introduction

1.1 Motivation and Results

We investigate behavior in a given game of incomplete information, where the latter is described by a “basic game” and by an “information structure”. The basic game refers to the set of actions, the set of payoff states, the utility functions of the players, and the common prior over the payoff states. The information structure refers to the type space of the game, which is generated by a mapping from the payoff states to a probability distribution over types, or signals. We ask what might happen in equilibrium if players may have access to additional signals beyond the given “information structure”? We show in Theorem 1 that behavior corresponds to a Bayes Nash equilibrium for some extra information that the players might observe if and only if it is an incomplete information version of correlated equilibrium that we dub *Bayes correlated equilibrium*.

A *decision rule* specifies a distribution over actions for each type profile and payoff state. A decision rule is a *Bayes correlated equilibrium* if it satisfies an obedience condition: a player does not have an incentive to deviate from the action recommended by the decision rule if he knows only his type and the action recommendation. There are a number of reasons why the notion of Bayes correlated equilibrium and its characterization result are of interest. *First*, it allows the analyst to identify properties of equilibrium outcomes that are going to hold independently of features of the information structure that the analyst does not know; in this sense, properties that hold in all Bayes correlated equilibria of a given incomplete information game constitute robust predictions. *Second*, it provides a way to partially identify parameters of the underlying economic environment independently of knowledge of the information structure. *Third*, it provides an indirect method of identifying socially or privately optimal information structures without explicitly working with a space of all information structures. In Bergemann and Morris [2013a], we illustrate these uses of the characterization result in a particular class of continuum player, linear best response games, focussing on normal distributions of types and actions and symmetric information structures and outcomes. In this paper, we focus on game theoretic foundations.¹

The separation between the basic game and the information structure enables us to ask how changes in the information structure affect the equilibrium set for a fixed basic game. A second contribution of the paper is that (i) we introduce a statistical partial order on information structures - called *individual*

¹We report an example in the Appendix that illustrates both Bayes correlated equilibrium and these applications in the context of a finite game and thus the setting of this paper.

sufficiency - that captures intuitively when one information structure contains more information than another; and (ii) we show that the set of Bayes correlated equilibria *shrinks* in all games if and only if the informativeness of the information structure increases. Thus, if the information structure of the players contains more information, then a smaller set of outcomes is incentive compatible.

To describe the order on information structures, it is useful to note that a one-player version of an information structure is an "experiment" in the sense studied by Blackwell [1951], [1953]. An experiment consists of a set of signals and a mapping from states to probability distributions over signals. Suppose that we are interested in comparing a pair of experiments. A *combination* of the two experiments is a new experiment where a pair of signals - one from each experiment - is observed and the marginal probability over signals from each of the original experiments corresponds to the original distribution over signals for that experiment. One way of characterizing the classic sufficiency condition of Blackwell [1951] is the following: one experiment is *sufficient* for another if it is possible to construct a combined experiment such that the former experiment is a sufficient statistic for the combined experiment.

Our partial order on (many-player) information structures is a player-by-player generalization of sufficiency. One information structure is *individually sufficient* for another if there exists a combined information structure where each player's signal from the former information structure is a sufficient statistic for the state and other players' signals in the latter information structure. This partial order has a couple of key properties - each generalizing well known properties in the one-player case - that suggest that it is the "right" ordering on (many-player) information structures. First, two information structures are individually sufficient for each other if and only if they have the same canonical representation, where signals are identified with higher-order beliefs about states. Second, one information structure is individually sufficient for another if and only if it is possible to start with the latter information structure and then have each player observe an extra signal, so that the expanded information structure has the same canonical representation as the former information structure.

We analyze an "incentive ordering" on information structures: an information structure is more *incentive constrained* than another if it gives rise to a smaller set of Bayes correlated equilibria. Our main result, Theorem 2, is that one information structure is more incentive constrained than another if and only if the former is individually sufficient for the latter. Thus we show the equivalence between a statistical ordering and an incentive ordering.

Blackwell's theorem showed that if one experiment was sufficient for another, then making decisions based on the former experiment allows a decision maker to attain a richer set of outcomes, and thus higher ex ante utility. Thus Blackwell's theorem showed the equivalence of a "statistical ordering" on experiments (sufficiency) and a "feasibility ordering" (more valuable than). Our main result, restricted to the one-

player case, has a natural interpretation and shows an equivalence between a statistical ordering and an incentive ordering, and thus can be seen as an extension of Blackwell's theorem. To further understand the connection to Blackwell's theorem, we also describe a feasibility ordering on many-player information structures and establish in Theorem 3 that is equivalent to individual sufficiency and "more incentive constrained than".

Taken together, our main result and discussion of the relation to Blackwell's theorem highlight the dual role of information. By making more outcomes feasible, more information allows more outcomes that can occur. By adding incentive constraints, more information restricts the set of *equilibrium* outcomes that do occur. The same partial order—individual sufficiency, reducing to sufficiency in the one-player case—captures both roles of information simultaneously.

1.2 Related Literature

Hirshleifer [1971] showed how information might be damaging in a many-player context because it removed options to insure ex ante. In mechanism design, it is well understood how more information may reduce the set of attainable outcomes by adding incentive constraints. Our result on the incentive constrained ordering can be seen as a formalization of the idea behind the observation of Hirshleifer [1971]: we give a general statement of how more information creates more incentive constraints and thus reduces the set of incentive compatible outcomes.

Aumann [1974], [1987] introduced the notion of correlated equilibrium in games with complete information and a number of definitions of correlated equilibrium in games with incomplete information have been suggested, notably in Forges [1993], [2006]. A maintained assumption in that literature - which we dub "join feasibility" - is that play can only depend on the combined information of all the players. This restriction is natural under the maintained assumption that correlated equilibrium is intended to capture the role of correlation of the players' actions but not unexplained correlation with the state of nature. Our different motivation leads us to allow such unexplained correlation. Liu [2015] also relaxes the join feasibility assumption, but imposes a belief invariance assumption (introduced and studied in combination with join feasibility in Forges [1993], [2006]), requiring that, from each player's point of view, the action recommendation that he receives from the mediator not change his beliefs about others' types and the state. Intuitively, the belief invariant Bayes correlated equilibria of Liu [2015] capture the implications of common certainty of rationality and a fixed information structure, while our Bayes correlated equilibria capture the implications of common certainty of rationality and the fact that the player have observed at least the signals in the information structure, and possibly additional information.

Gossner [2000] introduced a partial order on information structures, expressed in statistical terms, that

characterized when the set of outcomes that can arise in Bayes Nash equilibrium shrinks going from one information structure to another. We perform the analogous exercise for Bayes correlated equilibrium. The analysis of the Bayes Nash equilibrium inevitably conflated issues of incentives - more information imposes more incentive constraints - and feasibility - more information allows more things to happen. As a result, the statistical partial order of Gossner [2000] *never* ranks information structures corresponding to different beliefs and higher-order beliefs about the state. It simply characterizes when one information structure permits more correlation than another.² By contrast, the notion of Bayes correlated equilibrium abstracts from feasibility considerations by construction. In fact, we show that mere correlation possibilities are irrelevant in our partial order and information structures are ranked based *only* on beliefs and higher-order beliefs about the state. Nonetheless, our arguments are closest to those of Gossner [2000], as our main result can be seen as removing feasibility considerations from the main argument of Gossner [2000]. Lehrer et al. [2013] study solution concepts which are intermediate between Bayes correlated equilibrium and Bayes Nash equilibrium, and provide partial characterizations of how the set of equilibrium outcomes vary with the information structure.

Our characterization result also has an important one-player analogue. Consider a decision maker who has access to an experiment, but may have access to more information. The joint distribution of actions, signals and states that might result in a given decision problem is equal to the set of one-player Bayes correlated equilibria. Such one-player Bayes correlated equilibria have already arisen in a variety of contexts. Kamenica and Gentzkow [2011] consider the problem of cheap talk with commitment ("Bayesian persuasion"). In order to understand the behavior that a sender/speaker can induce a receiver/decision maker to choose, one must first characterize all outcomes that can arise for some committed cheap talk (independent of the objectives of the speaker). This, in our language, is the set of one-player Bayes correlated equilibria in the case of a null experiment (where the information structure of the receiver is simply the common prior over the payoff states and hence contains null additional information). In this sense, our work provides an approach for studying a many receiver version of Kamenica and Gentzkow [2011] where receivers have prior information. Kamenica and Gentzkow [2011] is based on a concavification argument introduced in the study of repeated games by Aumann and Maschler [1995].³ Thus our work can be seen as an extension of Aumann [1987] to environments with incomplete information by extending the

²The main result in Gossner [2000] is about complete information games, but our discussion of Gossner [2000] here and in the rest of the paper refers to Section 6 and Theorem 17 which briefly reports the extension to incomplete information. See Cherry and Smith [2014] for an alternative approach to Gossner's question in the complete information case.

³Aumann and Maschler [1995] showed that in infinitely repeated zero sum games with one sided uncertainty and without discounting, the outcome of the repeated game is as if the informed player can commit to reveal only certain information about the state in the corresponding static game. They then showed that a concavification of the complete information payoff function yields the complete characterization of the set of feasible payoffs in the one player game of private information.

analysis of Aumann and Maschler [1995] to many players and general, many-player, information structures.

Our main result concerns an ordering on information structures based on the idea that more information reduces the set of outcomes by imposing more incentive constraints, i.e., an *incentive ordering*. By contrast, for the one-player case, Blackwell [1951] characterized an order on information structures based on the idea that more information increases the set of feasible outcomes, and thus increases the set of attainable payoffs; i.e., a *feasibility ordering*. Lehrer et al. [2010] propose a natural way of studying feasibility orderings in the many-player case: see what can happen in equilibria in common interest games under different solution concepts. If we look for the best (common) payoff under feasible strategy profiles (under a given solution concept), then more information, by making more outcomes feasible, will lead to a higher maximum common payoff. They characterize the ordering on information structures that increases the maximum payoff in all common interest games, for different solution concepts. The relevant ordering on information structures varies with the feasibility constraints built into the solution concept. It is an easy corollary of the results of Lehrer et al. [2010] that an information structure is individually sufficient for another if and only if, in any common interest game, the maximum payoff attainable in *belief invariant* Bayes correlated equilibrium (as defined above) is weakly higher under the former information structure than under the latter information structure. Thus our feasibility result, Theorem 3, follows Lehrer et al. [2010], [2013] in showing that the same ordering on information structures which is relevant for incentive orderings is also relevant for feasibility orderings.⁴

The structure of the remainder of the paper is as follows. In Section 2, we define the notion of Bayes correlated equilibrium for a general finite game and establish the first result, Theorem 1, showing the relationship between Bayes correlated equilibria and Bayes Nash equilibria of games with more information. In Section 3, we describe a many-player generalization of the sufficiency ordering of information structures, dubbed *individual sufficiency*. We also relate individual sufficiency to beliefs and higher-order beliefs, and illustrate the different notions with binary information structures. In Section 4 we present the second result, Theorem 2, which establishes an equivalence between the incentive based ordering and the statistical ordering. In Section 5, we place "Bayes correlated equilibrium" in the context of the literature on incomplete information correlated equilibrium, discuss the relation to alternative orderings on information structures, including feasibility orderings, and Blackwell's Theorem, and show how our results can be used to give a many-player approach to "Bayesian persuasion".

⁴Gossner [2010] also highlights the dual role of information in a different analytic setting.

2 Bayes Correlated Equilibrium

2.1 Definition

There are I players, $1, 2, \dots, I$, and we write i for a typical player. There is a finite set of states, Θ , and we write θ for a typical state. A *basic game* G consists of (1) for each player i , a finite set of actions A_i , where we write $A = A_1 \times \dots \times A_I$, and a utility function $u_i : A \times \Theta \rightarrow \mathbb{R}$; and (2) a full support common prior $\psi \in \Delta_{++}(\Theta)$. Thus $G = \left((A_i, u_i)_{i=1}^I, \psi \right)$. An *information structure* S consists of (1) for each player i , a finite set of signals (or types) T_i , where we write $T = T_1 \times \dots \times T_I$; and (2) a signal distribution $\pi : \Theta \rightarrow \Delta(T)$. Thus $S = \left((T_i)_{i=1}^I, \pi \right)$. A possible (and natural) information structure is the *null* information structure in which each player's set of signals T_i is a singleton, $T_i = \{t_i\}$. This corresponds to the situation in which each player has no information over and above the common prior ψ .

Together, the basic game G and the information structure S define a standard incomplete information game. This division of an incomplete information game into the "basic game" and the "information structure" has now been widely used (see, for example, Gossner [2000]).

A *decision rule* in the incomplete information game (G, S) is a mapping $\sigma :$

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

One way to mechanically understand the notion of the decision rule is to view σ as the strategy of an omniscient mediator who first observes the realization of $\theta \in \Theta$ chosen according to ψ and the realization of $t \in T$ chosen according to $\pi(\cdot|\theta)$; and then picks an action profile $a \in A$, and privately announces to each player i the draw of a_i . For players to have an incentive to follow the mediator's recommendation in this scenario, it would have to be the case that the recommended action a_i was always preferred to any other action a'_i conditional on the signal t_i that player i had received and his knowledge of the recommended action a_i . This is reflected in the following "obedience" condition.

Definition 1 (Obedience)

Decision rule σ 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

$$\begin{aligned} & \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i})|\theta) \sigma((a_i, a_{-i})|(t_i, t_{-i}), \theta) u_i((a_i, a_{-i}), \theta) \\ & \geq \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i})|\theta) \sigma((a_i, a_{-i})|(t_i, t_{-i}), \theta) u_i((a'_i, a_{-i}), \theta); \end{aligned} \quad (2)$$

for all $a'_i \in A_i$.

Our definition of Bayes correlated equilibrium requires obedience and nothing else.

Definition 2 (Bayes Correlated Equilibrium)

A decision rule σ is a Bayes correlated equilibrium (BCE) of (G, S) if it is obedient for (G, S) .

If there is complete information, i.e., if Θ is a singleton, then this definition reduces to the Aumann [1987] definition of correlated equilibrium for a complete information game. If S is the null information structure, then this is essentially the "universal Bayesian solution" of Forges [1993]. If, in addition, there is only one player then this definition reduces to behavior in the concavification problem of Aumann and Maschler [1995] and the Bayesian persuasion of Kamenica and Gentzkow [2011]. We postpone until Section 5 a discussion of these connections and how this definition relates to (and is in general weaker than) other definitions in the literature on incomplete information correlated equilibrium. We provide our motivation for studying this particular definition next.

Consider an analyst who knew that:

1. the basic game G describes actions, payoff functions depending on states, and a prior distribution on states,
2. the players observe at least information structure S , but may observe more,
3. the players' actions constitute a Bayes Nash equilibrium given the actual information structure.

What joint distributions of actions, signals (in the original information structure, S) and states can arise in such an equilibrium? We will formalize this question and show that the answer is the set of Bayes correlated equilibria of (G, S) .

We first note the standard definition of Bayes Nash equilibrium in this setting. A (behavioral) strategy for player i in the incomplete information game (G, S) is $\beta_i : T_i \rightarrow \Delta(A_i)$.

Definition 3 (Bayes Nash Equilibrium)

A strategy profile β is a Bayes Nash equilibrium (BNE) of (G, S) if for each $i = 1, \dots, I$, $t_i \in T_i$ and $a_i \in A_i$ with $\beta_i(a_i|t_i) > 0$, we have

$$\begin{aligned} & \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i}) | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j) \right) u_i((a_i, a_{-i}), \theta) \\ & \geq \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i}) | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j) \right) u_i((a'_i, a_{-i}), \theta), \end{aligned} \tag{3}$$

for each $a'_i \in A_i$.

2.2 Foundations

We want to discuss situations where players observe more information than that contained in a given information structure. To formalize this, we use the concept of *combinations* of information structures. If we have two information structures $S^1 = (T^1, \pi^1)$ and $S^2 = (T^2, \pi^2)$, we will say that information structure $S^* = (T^*, \pi^*)$ is a *combination* of information structures S^1 and S^2 if the *combined* information structure $S^* = (T^*, \pi^*)$ is obtained by forming a product space of the signals, $T_i^* = T_i^1 \times T_i^2$ for each i , and a signal distribution $\pi^* : \Theta \rightarrow \Delta(T^1 \times T^2)$ that preserves the marginal distribution of its constituent information structures.

Definition 4 (Combination)

The information structure $S^* = (T^*, \pi^*)$ is a combination of information structures $S^1 = (T^1, \pi^1)$ and $S^2 = (T^2, \pi^2)$ if

$$T_i^* = T_i^1 \times T_i^2 \text{ for each } i; \quad (4)$$

and

$$\begin{aligned} \sum_{t^2 \in T^2} \pi^*(t^1, t^2 | \theta) &= \pi^1(t^1 | \theta) \text{ for each } t^1 \in T^1 \text{ and } \theta \in \Theta, \\ \sum_{t^1 \in T^1} \pi^*(t^1, t^2 | \theta) &= \pi^2(t^2 | \theta) \text{ for each } t^2 \in T^2 \text{ and } \theta \in \Theta. \end{aligned} \quad (5)$$

Note that the above definition places no restrictions on whether signals $t^1 \in T^1$ and $t^2 \in T^2$ are independent or correlated, conditional on θ , under π^* . Thus any pair of information structures S^1 and S^2 will have many combined information structures.

Definition 5 (Expansion)

An information structure S^* is an expansion of S^1 if S^* is a combination of S^1 and some other information structure S^2 .

Suppose strategy profile β was played in (G, S^*) , where S^* is a combination of two information structures S^1 and S^2 . Now, if the analyst did not observe the signals of the combined information structure S^* , but only the signals of S^1 , then the behavior under the strategy profile β would induce a decision rule for (G, S^1) . Formally, the strategy profile β for (G, S^*) induces the decision rule σ for (G, S^1) :

$$\sigma(a | t^1, \theta) \triangleq \frac{\sum_{t^2 \in T^2} \pi^*(t^1, t^2 | \theta) \prod_{j=1}^I \beta_j(a_j | t_j^1, t_j^2)}{\pi^1(t^1 | \theta)},$$

for each $a \in A$ whenever $\pi^1(t^1 | \theta) > 0$.

Theorem 1

A decision rule σ is a Bayes correlated equilibrium of (G, S) if and only if, for some expansion S^* of S , there is a Bayes Nash equilibrium of (G, S^*) which induces σ .

Thus this is an incomplete information analogue of the Aumann [1987] characterization of correlated equilibrium for complete information games. An alternative interpretation of this result – following Aumann [1987] – would be to say that BCE captures the implications of common certainty of rationality (and the common prior assumption) in the game G when players have at least information S , since requiring BNE in some game with expanded information is equivalent to describing a belief closed subset where the game G is being played, players have access to (at least) information S and there is common certainty of rationality.⁵

In the Appendix, in Section 6, we provide an example to illustrate the Theorem. The example also demonstrates the usefulness of the characterization for identifying which expansion of the information structure is most desirable for the players of the game. In particular, public disclosure is optimal in games with strategic complementarities while private disclosure is optimal in games with strategic substitutes.

The proof follows the logic of the classic result of Aumann [1987] for complete information and that of Forges [1993] for the Bayesian solution for incomplete information games (discussed in Section 5).

Proof. Suppose that σ is a Bayes correlated equilibrium of (G, S) . Thus

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

for each i , $t_i \in T_i$, $a_i \in A_i$ and $a'_i \in A_i$. Let $S^* = \left((T_i^*)_{i=1}^I, \pi^* \right)$ be an expansion of S , and, in particular,

a combination of $S = \left((T_i)_{i=1}^I, \pi \right)$ and $S' = \left((T'_i)_{i=1}^I, \pi' \right)$, where $T'_i = A_i$ for each i and π^* satisfies

$$\pi^* \left((t_i, a_i)_{i=1}^I \middle| \theta \right) = \pi(t | \theta) \sigma(a | t, \theta), \tag{6}$$

for each $t \in T$, $a \in A$ and $\theta \in \Theta$. Now, in the game (G, S^*) , consider the "truthful" strategy β_j^* for player j , with

$$\beta_j^* (a'_j | (t_j, a_j)) = \begin{cases} 1, & \text{if } a'_j = a_j, \\ 0, & \text{if } a'_j \neq a_j, \end{cases} \tag{7}$$

⁵See Bergemann and Morris [2014b] for more on the relationship between Theorem 1 and the epistemic foundations of game theory literature.

for all $t_j \in T_j$ and $a_j \in A_j$. Now the interim payoff to player i observing signal (t_i, a_i) and choosing action a'_i in (G, S^*) if he anticipates that each opponent will follow strategy β_j^* is proportional to

$$\begin{aligned}
& \sum_{a'_i, a_i, t_{-i}, \theta} \psi(\theta) \pi^*((t_i, t_{-i}), (a_i, a_{-i}) | \theta) \left(\prod_{j \neq i} \beta_j^*(a'_j | t_j, a_j) \right) u_i((a'_i, a'_{-i}), \theta) \\
&= \sum_{a_i, t_{-i}, \theta} \psi(\theta) \pi^*((t_i, t_{-i}), (a_i, a_{-i}) | \theta) u_i((a'_i, a_{-i}), \theta), \text{ by (7)} \\
&= \sum_{a_i, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i}) | \theta) \sigma((a_i, a_{-i}) | (t_i, t_{-i}), \theta) u_i((a'_i, a_{-i}), \theta), \text{ by (6)}
\end{aligned}$$

and thus Bayes Nash equilibrium optimality conditions for the truth telling strategy profile β^* are implied by the obedience conditions on σ . By construction, β induces σ .

Conversely, suppose that β is a Bayes Nash equilibrium of (G, S^*) , where S^* is a combined information structure for S and S' . Write $\sigma : T \times \Theta \rightarrow \Delta(A)$ for the decision rule for (G, S) induced by β , so that

$$\pi(t | \theta) \sigma(a | t, \theta) = \sum_{t' \in T'} \pi^*((t_i, t'_i)_{i=1}^I | \theta) \prod_{j=1}^I \beta_j(a_j | t_j, t'_j)$$

for each $t \in T$, $a \in A$ and $\theta \in \Theta$. Now $\beta_i(a_i | (t_i, t'_i)) > 0$ implies

$$\begin{aligned}
& \sum_{a_i, t_{-i}, t'_{-i}, \theta} \psi(\theta) \pi^*((t_i, t'_i)_{i=1}^I | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j, t'_j) \right) u_i((a_i, a_{-i}), \theta) \\
& \geq \sum_{a_i, t_{-i}, t'_{-i}, \theta} \psi(\theta) \pi^*((t_i, t'_i)_{i=1}^I | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j, t'_j) \right) u_i((a'_i, a_{-i}), \theta),
\end{aligned}$$

for each i , $t_i \in T_i$, $t'_i \in T'_i$ and $a'_i \in A_i$. Thus

$$\begin{aligned}
& \sum_{t'_i} \beta_i(a_i | (t_i, t'_i)) \sum_{a_{-i}, t_{-i}, t'_{-i}, \theta} \psi(\theta) \pi^*((t_i, t'_i)_{i=1}^I | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j, t'_j) \right) u_i((a_i, a_{-i}), \theta) \\
& \geq \sum_{t'_i} \beta_i(a_i | (t_i, t'_i)) \sum_{a_{-i}, t_{-i}, t'_{-i}, \theta} \psi(\theta) \pi^*((t_i, t'_i)_{i=1}^I | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j, t'_j) \right) u_i((a'_i, a_{-i}), \theta),
\end{aligned}$$

for each i , $t_i \in T_i$ and $a'_i \in A_i$. But

$$\begin{aligned}
& \sum_{t'_i} \beta_i(a_i | (t_i, t'_i)) \sum_{a_{-i}, t_{-i}, t'_{-i}, \theta} \psi(\theta) \pi^*((t_i, t'_i)_{i=1}^I | \theta) \left(\prod_{j \neq i} \beta_j(a_j | t_j, t'_j) \right) u_i((a'_i, a_{-i}), \theta) \\
&= \sum_{a_i, t_{-i}, \theta} \psi(\theta) \pi((t_i, t_{-i}) | \theta) \sigma((a_i, a_{-i}) | (t_i, t_{-i}), \theta) u_i((a'_i, a_{-i}), \theta),
\end{aligned}$$

and thus BNE equilibrium conditions in (G, S^*) imply the obedience conditions of σ in (G, S) . ■

3 A Partial Order on Information Structures

We will study the following partial order on information structures. An information structure S is *individually sufficient* for S' if each player's probability of his signal under S' conditional on his signal under S is independent of the state and others' signals in S . To be more precise, we require that these player-by-player conditional independence properties hold in some "combined information structure" - a probability space which embeds both information structures; we formally defined a combined information structure in Definition 4 in the previous Section. Thus we have:

Definition 6 (Individual Sufficiency)

Information structure $S = (T, \pi)$ is individually sufficient for information structure $S' = (T', \pi')$ if there exists a combined information structure $S^* = (T^*, \pi^*)$ such that, for each i ,

$$\Pr(t'_i | t_i, t_{-i}, \theta) \triangleq \frac{\sum_{t'_{-i}} \pi^*((t_i, t_{-i}), (t'_i, t'_{-i}) | \theta)}{\sum_{\tilde{t}'_i, t'_{-i}} \pi^*((t_i, t_{-i}), (\tilde{t}'_i, t'_{-i}) | \theta)} \text{ is independent of } t_{-i} \text{ and } \theta \quad (8)$$

whenever the denominator is non-zero.

Thus, for each player i , the probability of t'_i conditional on t_i is independent of (t_{-i}, θ) . In the one-player special case, individual sufficiency reduces to the sufficiency ordering on experiments of Blackwell [1951], [1953]. An equivalent way of stating the condition, using the pairwise Markov property of conditional independence, is that, for each player i , the probability of (t_{-i}, θ) conditional on t_i is independent of t'_i .⁶ Thus the key property of this extension of Blackwell's order is that an agent's signal t'_i in the combined information structure S^* *neither* provides new information about θ relative to t_i *nor* provides new information about t_{-i} . Thus while the signals t'_i may be correlated across agent, i.e., with t'_{-i} , they do not convey any additional information about beliefs and higher order beliefs represented by (t_{-i}, θ) . Crucially, while we only require the conditional independence properties to hold player-by-player, we do require them to hold in the same combined information structure. We postpone until Section 5.2 a discussion of alternative extensions of the Blackwell order to the many-player case, and how they have been used in the existing literature.

The partial order of individual sufficiency inherits two key properties of Blackwell's order in the one-player case. First, for any given experiment (i.e., one-player information structure), we can define its canonical representation to be the one where we merge signals that give rise to the same posterior over

⁶In Bergemann and Morris [2014a] we give a third equivalent definition of individual sufficiency which extends the Markov kernel statement of sufficiency of Blackwell [1951], (1953) to the many-player version of individual sufficiency.

states and label signals according to their posteriors over states; two experiments are sufficient for each other if and only if they have the same canonical representation. Second, if you start with an experiment, and then the decision maker observes an additional signal, then the combined experiment is trivially sufficient for the original experiment. But a converse is also true. If an experiment is sufficient for another, then we can start with the latter experiment, provide an additional signal to the player, and show that the combined experiment has the same canonical representation as the former.

To state the many-player analogues of these two properties, recall that Mertens and Zamir [1985] defined a canonical representation of a many-player information structure, a universal type space, to be one where we merge types with the same beliefs and higher order beliefs about the state and label types according to their beliefs and higher-order beliefs about the state, and that we defined an "expansion" of an information structure S to be the combination of S with any other information structure. Now we have:

Claim 1

1. *Two information structures are individually sufficient for each other if and only if they have the same canonical representation.*⁷
2. *Information structure S is individually sufficient for S' if and only if there exists an expansion of S' which has the same canonical representation as S .*

We now report two examples which we will use to illustrate the ordering and these two properties.⁸ We will return to them when we discuss the relation to alternative orderings. For both examples, we assume that there are two possible states, θ_0 and θ_1 .

Example 1. The first comparison illustrates the irrelevance of access to correlating devices, i.e., information which is "redundant" in the sense of Mertens and Zamir [1985]. Examples such as this have been leading examples in the literature, see Dekel et al. [2007], Ely and Peski [2006] and Liu [2015]. Let S be a "null" information structure where each player has only one possible signal which is always observed.

⁷Showing this is equivalent to showing that if two information structures are individually sufficient for each other, then one can use the same combined information structure to verify this. This was shown by Liu [2015] and is an implication of arguments in Lehrer et al. [2013]. In an earlier version of this paper (Bergemann and Morris [2014a]), we reported a formal statement and proof of these claims in our language.

⁸Thus, strictly speaking, individual sufficiency per se forms a preorder and becomes a partial order after restricting attention to the equivalence classes suggested by the canonical representation above.

Let S' be given by

$\pi'(\cdot \theta_0)$	t'_0	t'_1
t'_0	$\frac{1}{2}$	0
t'_1	0	$\frac{1}{2}$

$\pi'(\cdot \theta_1)$	t'_0	t'_1
t'_0	0	$\frac{1}{2}$
t'_1	$\frac{1}{2}$	0

where each player observes one of two signals, t'_0 and t'_1 ; the above tables describe the probabilities of different signal profiles, where player 1's signal corresponds to the row, player 2's signal corresponds to the column, the left and right hand tables correspond to the distribution of signal profiles in states θ_0 and θ_1 respectively, and the table entries correspond to the conditional probabilities of those signal profiles.

In this case, there is a unique combined information structure (because signals in S are redundant). Each information structure is individually sufficient for the other, because each player's signal in S' gives no additional information about the state (and the redundant signal of the other player in S). Thus there is common certainty that each player assigns probability $\frac{1}{2}$ to each state. This example illustrates the first part of Claim 1.

This example may suggest that individual sufficiency can be checked by first removing redundancies and then checking "informativeness" player-by-player. The next example is the simplest possible to illustrate that this is not the case and that individual sufficiency is a more subtle relation.

Example 2. We now compare two new information structures with the same signal sets and labels that we used previously. Let information structure S be given by

$\pi(\cdot \theta_0)$	t_0	t_1
t_0	$\frac{1}{2}$	0
t_1	0	$\frac{1}{2}$

$\pi(\cdot \theta_1)$	t_0	t_1
t_0	0	0
t_1	0	1

(9)

Under information structure S , if the state is θ_0 , with probability $\frac{1}{2}$, it is common knowledge that the state is θ_0 (and both players observe signal t_0); otherwise, both players observe t_1 . Consider now a second binary information structure S' given by:

$\pi'(\cdot \theta_0)$	t'_0	t'_1
t'_0	$\frac{1}{2}$	$\frac{1}{6}$
t'_1	$\frac{1}{6}$	$\frac{1}{6}$

$\pi'(\cdot \theta_1)$	t'_0	t'_1
t'_0	$\frac{1}{3}$	0
t'_1	0	$\frac{2}{3}$

(10)

In the information structure S' , each player observes a signal with "accuracy" $\frac{2}{3}$ in either state: that is, if the state is θ_0 , then each player observes t'_0 with probability $\frac{2}{3}$; if the state is θ_1 , then each player observes t'_1 with probability $\frac{2}{3}$. But in state θ_1 , the signals are perfectly correlated across players, whereas in state θ_0 , the signals are less than perfectly correlated.

Neither of these two information structures has any redundancies, but S is individually sufficient for S' . We can illustrate our characterization by identifying a combined information structure S^* which can be used to establish individual sufficiency:

$\pi^*(\cdot \theta_0)$	$t'_0t'_0$	$t'_0t'_1$	$t'_1t'_0$	$t'_1t'_1$
t_0t_0	$\frac{1}{2}$	0	0	0
t_1t_1	0	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$

$\pi^*(\cdot \theta_1)$	$t'_0t'_0$	$t'_0t'_1$	$t'_1t'_0$	$t'_1t'_1$
t_1t_1	$\frac{1}{3}$	0	0	$\frac{2}{3}$

If we look at the marginal of π^* on signals from S , we obtain (9), whereas if we look at the marginal of π^* on signals from S' , we get (10). The following table reports the conditional probabilities necessary to verify individual sufficiency: the rows correspond to the three triples consisting of player 1's type in S , player 2's type in S and the state that arises with positive probability; the columns correspond (for either player, by symmetry) to the player's type in information structure S' and the table entries correspond to the probability of the latter conditional on the former:

	t'_0	t'_1
$t_0t_0\theta_0$	1	0
$t_1t_1\theta_0$	$\frac{1}{3}$	$\frac{2}{3}$
$t_1t_1\theta_1$	$\frac{1}{3}$	$\frac{2}{3}$

Since, for any player, the probability of his signal under S' conditional on his signal under S is independent of the other player's signal under S and the state, the conditional independence property (8) is verified.

We can also illustrate the second part of Claim 1, by verifying that the combined information structure has the same canonical representation as S . Suppose that the two states are equally likely and we start out with information structure S' . Suppose that players are then given an additional public signal: if the true state is θ_0 and both observed signal t'_0 , then both will observe signal t_0 ; otherwise both will observe signal t_1 . This expansion of information structure S' corresponds to the combined information structure S^* . If players observe t_0 , then there is common certainty that the state is θ_0 . If players observe t_1 , then one can verify that there is common certainty that both players assign probability $\frac{1}{3}$ to state θ_0 (independent of what signals in S' they started out with).

To formally verify this, note that in the combined information structure S^* , each player has three possible types (or *combined types*) which are a combination of types in S and S' : $t^* \in \{t_0t'_0, t_1t'_0, t_1t'_1\}$. We can re-arrange the representation of the combined information structure S^* in the following table: rows correspond to combined types t^* for S^* (of either player), columns correspond to the possible profiles of the other player's combined type and state θ , and entries correspond to the conditional probabilities of the

latter given the former:

	$\theta_0 t_0 t'_0$	$\theta_0 t_1 t'_0$	$\theta_0 t_1 t'_1$	$\theta_1 t_1 t'_0$	$\theta_1 t_1 t'_1$
$t_0 t'_0$	1	0	0	0	0
$t_1 t'_0$	0	0	$\frac{1}{3}$	$\frac{2}{3}$	0
$t_1 t'_1$	0	$\frac{1}{6}$	$\frac{1}{6}$	0	$\frac{2}{3}$

Now from the above table, we can see that combined types $t_1 t'_0$ and $t_1 t'_1$ both assign probability $\frac{1}{3}$ to state θ_0 (and $\frac{2}{3}$ to state θ_1), and thus cannot be distinguished on the basis of their first order beliefs. But we also see that combined types $t_1 t'_0$ and $t_1 t'_1$ both assign probability $\frac{1}{3}$ to the event that $\theta = \theta_0$ and the other player assigning probability $\frac{1}{3}$ to state 0. Thus combined types $t_1 t'_0$ and $t_1 t'_1$ cannot be distinguished on the basis of their second order beliefs. And so on.

4 Comparing Information Structures

Giving players more information will generate more obedience constraints and thus reduce in size the set of Bayes correlated equilibria. If "giving players more information" is interpreted to mean that we expand their information structures, allowing them to keep their previous signals and observe more, then this claim follows trivially from the definition and characterization of Bayes correlated equilibria in Section 2. In this Section, we strengthen this observation by showing that it is also true if by "giving player more information," we mean that we replace their information structure with one that is individually sufficient for it. And we prove a converse, showing that if an information structure, S , is not individually sufficient for another, S' , then there exists a basic game G such that (G, S) has a Bayes correlated equilibrium that generates outcomes that could not arise under a Bayes correlated equilibrium of (G, S') .

In order to compare outcomes across information structures, we will be interested in what can be said about actions and states if signals are not observed. We will call a mapping

$$\nu : \Theta \rightarrow \Delta(A), \tag{11}$$

an *outcome*, and say ν is *induced* by decision rule σ if it is the marginal of σ on A , so that

$$\nu(a|\theta) \triangleq \sum_{t \in T} \sigma(a|t, \theta) \pi(t|\theta), \tag{12}$$

for each $a \in A$ and $\theta \in \Theta$. Outcome ν is a Bayes correlated equilibrium outcome of (G, S) if it is induced by a Bayes correlated equilibrium decision rule σ of (G, S) .

We now define a partial order on information structures that corresponds to shrinking the set of BCE outcomes in all basic games. Writing $BCE(G, S)$ for the set of BCE outcomes of (G, S) , we say:

Definition 7 (Incentive Constrained)

Information structure S is more incentive constrained than information structure S' if, for all basic games G :

$$BCE(G, S) \subseteq BCE(G, S').$$

We call this partial order "more incentive constrained than" because, given our definition of Bayes correlated equilibrium, it captures exactly the role of information in imposing more incentive constraints. Thus an information structure giving rise to a smaller set of Bayes correlated equilibria in all games corresponds to a *more* informed information structure. By contrast, if we replaced Bayes correlated equilibrium in the above definition with Bayes Nash equilibrium - which corresponds to the problem studied by Gossner [2000] - a smaller set of Bayes Nash equilibria corresponds to a *less* informed information structure as the notion of Bayes Nash equilibrium imposes strong feasibility constraints on the outcome mapping.

Theorem 2

Information structure S is individually sufficient for information structure S' if and only if S is more incentive constrained than S' .

We report an example illustrating the Theorem in the Appendix.

To prove the result, we first show constructively that if S is individually sufficient for S' and ν is a BCE outcome of (G, S) , then we can use the BCE decision rule inducing ν and the combined information structure establishing individual sufficiency to construct a decision rule of (G, S') which induces ν . The incentive constraints under S' are averages of the incentive constraints under S , and therefore the incentive compatibility of the original decision rule for (G, S) is preserved for (G, S') . Versions of this argument have been used by Gossner [2000], Lehrer et al. [2013] and Liu [2015] to prove similar claims working with different solution concepts and orderings on information structures.

To prove the converse, we consider, for any information structure S , a particular basic game G and a particular BCE outcome ν of (G, S) . If S is more incentive constrained than S' , that particular ν must also be a BCE outcome of (G, S') . We then show that our choice of G and ν imply that, if ν is a BCE outcome of (G, S') , there must exist a combined information structure establishing that S is individually sufficient for S' . To show this, we use the basic game G where each player i is asked to report *either* a type in T_i (which is associated under S with a belief over $T_{-i} \times \Theta$) *or* an arbitrary belief over $T_{-i} \times \Theta$ (which does not in general correspond to an element of T_i). Players are then given an incentive to truthfully report their beliefs over $T_{-i} \times \Theta$ (which may or may not correspond to a type in T_i) using a quadratic scoring rule.

There is a BCE of (G, S) where players "truthfully" report their types in S . This BCE thus induces the outcome $\pi : \Theta \rightarrow \Delta(T)$. Now consider any decision rule σ' for (G, S') which induces the same outcome π . Combining $\pi' : \Theta \rightarrow \Delta(T')$ and $\sigma' : T' \times \Theta \rightarrow \Delta(T)$ gives a combined information structure for S and S' with $\pi^*(t, t'|\theta) = \pi'(t'|\theta) \sigma'(t|t', \theta)$. Obedience of σ' in the game (G, S') now implies that, under the combined information structure, the beliefs of type t'_i about (t_{-i}, θ) when recommended to take action t_i must equal the beliefs of t_i about (t_{-i}, θ) under information structure S alone. But now we have a combined information structure establishing individual sufficiency. This heuristic argument uses an infinite action basic game, and we are restricted to finite games. In the formal proof, we use finite approximations of this infinite action game and a continuity argument to establish our result.

This step has parallels in an similar argument in Gossner [2000].⁹ There are two differences. First, and less substantively, Gossner [2000] allows general games (not just finite games) which changes technical aspects of the argument. More importantly, because Gossner [2000] works with the solution concept of Bayes Nash equilibrium, feasibility constraints matter in the argument and so, in addition to establishing that there is a combined experiment establishing individual sufficiency, the combined experiment must satisfy additional properties reflecting feasibility restrictions not present in our analysis, giving rise to a very different statistical ordering that we will discuss in Section 5.2.

Proof. Suppose that S is individually sufficient for S' . Take any basic game G and any BCE σ of (G, S) . We will construct $\sigma' : T' \times \Theta \rightarrow \Delta(A)$ which is a BCE of (G, S') which gives rise to the same outcome as σ . Write $V_i(a_i, a'_i, t_i)$ for the expected utility for player i under decision rule σ if he is type t_i , receives recommendation a_i but chooses action a'_i , so that

$$V_i(a_i, a'_i, t_i) \triangleq \sum_{a_{-i} \in A_{-i}, t_{-i} \in T_{-i}, \theta \in \Theta} \psi(\theta) \pi((t_i, t_{-i})|\theta) \sigma((a_i, a_{-i})|(t_i, t_{-i}), \theta) u_i((a'_i, a_{-i}), \theta).$$

Now – by Definition 1 – for each $i = 1, \dots, I$, $t_i \in T_i$ and $a_i \in A_i$, we have

$$V_i(a_i, a_i, t_i) \geq V_i(a_i, a'_i, t_i), \tag{13}$$

for each $a'_i \in A_i$. Since S is individually sufficient for S' , there exists a combined information structure satisfying (8). Define $\sigma' : T' \times \Theta \rightarrow \Delta(A)$ by

$$\sigma'(a|t', \theta) = \frac{\sum_{t \in T} \pi^*(t, t'|\theta) \sigma(a|t, \theta)}{\pi'(t'|\theta)}, \tag{14}$$

⁹We are grateful to Marcin Peski for clarifying the connection to Gossner [2000], which also suggested a simplification of the proof of Theorem 2. In private communication, Peski has suggested how our proof could be unified with (a finite version of) Gossner [2000].

for all $(a, t', \theta) \in A \times T' \times \Theta$ whenever $\pi(t'|\theta) > 0$ (and if $\pi(t'|\theta) = 0$, we are free to choose an arbitrary probability distribution $\sigma'(a|t', \theta)$). By construction, decision rules $\sigma(a|t, \theta)$ and $\sigma'(a|t', \theta)$ induce the same outcome function $\nu : \Theta \rightarrow \Delta(A)$. Write $V'_i(a_i, a'_i, t'_i)$ for the expected utility for player i under decision rule σ' if he is type t'_i , receives recommendation a_i but chooses action a'_i , so that

$$V'_i(a_i, a'_i, t'_i) \triangleq \sum_{a_{-i} \in A_{-i}, t'_{-i} \in T'_{-i}, \theta \in \Theta} \psi(\theta) \pi'((t'_i, t'_{-i})|\theta) \sigma'((a_i, a_{-i})|(t'_i, t'_{-i}), \theta) u_i((a'_i, a_{-i}), \theta).$$

Now σ' satisfies the obedience condition (Definition 1) to be a correlated equilibrium of (G, S') if for each $i = 1, \dots, I$, $t'_i \in T'_i$ and $a_i \in A_i$,

$$V'_i(a_i, a_i, t'_i) \geq V'_i(a_i, a'_i, t'_i),$$

for all $a'_i \in A_i$. Condition (8) in the definition of individual sufficiency implies the existence of $\phi_i : T_i \rightarrow \Delta(T'_i)$ such that

$$\phi_i(t'_i|t_i) \pi((t_i, t_{-i})|\theta) = \sum_{t'_{-i}} \pi^*((t_i, t_{-i}), (t'_i, t'_{-i})|\theta) \quad (15)$$

for each t'_i, t_i, t_{-i} and θ . Now

$$\begin{aligned} & V'_i(a_i, a'_i, t'_i) \\ = & \sum_{a_{-i} \in A_{-i}, t'_{-i} \in T'_{-i}, \theta \in \Theta} \psi(\theta) \pi'((t'_i, t'_{-i})|\theta) \sigma'((a_i, a_{-i})|(t'_i, t'_{-i}), \theta) u_i((a'_i, a_{-i}), \theta) \\ = & \sum_{a_{-i} \in A_{-i}, t'_{-i} \in T'_{-i}, \theta \in \Theta, t \in T} \psi(\theta) \pi^*(t, t'|\theta) \sigma((a_i, a_{-i})|t, \theta) u_i((a'_i, a_{-i}), \theta) \\ & \text{by the definition of } \sigma', \text{ see (14),} \\ = & \sum_{a_{-i} \in A_{-i}, \theta \in \Theta, t \in T} \psi(\theta) \sigma((a_i, a_{-i})|t, \theta) u_i((a'_i, a_{-i}), \theta) \sum_{t'_{-i} \in T'_{-i}} \pi^*(t, (t'_i, t'_{-i})|\theta) \\ = & \sum_{a_{-i} \in A_{-i}, \theta \in \Theta, t \in T} \psi(\theta) \sigma((a_i, a_{-i})|t, \theta) u_i((a'_i, a_{-i}), \theta) \pi((t_i, t_{-i})|\theta) \phi_i(t'_i|t_i), \text{ by (15),} \\ = & \sum_{t_i \in T_i} \phi_i(t'_i|t_i) \left[\sum_{a_{-i} \in A_{-i}, \theta \in \Theta, t_{-i} \in T_{-i}} \psi(\theta) \pi((t_i, t_{-i})|\theta) \sigma((a_i, a_{-i})|(t_i, t_{-i}), \theta) u_i((a'_i, a_{-i}), \theta) \right] \\ = & \sum_{t_i \in T_i} \phi_i(t'_i|t_i) V_i(a_i, a'_i, t_i). \end{aligned} \quad (16)$$

Now for each $i = 1, \dots, I$, $t'_i \in T'_i$ and $a_i \in A_i$,

$$\begin{aligned} V'_i(a_i, a_i, t'_i) &= \sum_{t_i \in T_i} \phi_i(t'_i|t_i) V_i(a_i, a_i, t_i), \text{ by (16)} \\ &\geq \sum_{t_i \in T_i} \phi_i(t'_i|t_i) V_i(a_i, a'_i, t_i), \text{ by (13) for each } t_i \in T_i \\ &= V'_i(a_i, a'_i, t'_i), \text{ by (16)} \end{aligned}$$

for each $a'_i \in A_i$. Thus σ' is a BCE of (G, S') . By construction σ' and σ induce the outcome $\nu : \Theta \rightarrow \Delta(A)$. Since this argument started with an arbitrary BCE outcome ν of (G, S) and an arbitrary G , we have $BCE(G, S) \subseteq BCE(G, S')$ for all basic games G .

We now show the converse. We first introduce a notion of approximate individual sufficiency. Fix a full support prior $\psi \in \Delta_{++}(\Theta)$. Let $\lambda_i : T_i \rightarrow \Delta(T_{-i} \times \Theta)$ denote the induced belief of type t_i about (t_{-i}, θ) under S :

$$\lambda_i(t_{-i}, \theta | t_i) \triangleq \frac{\psi(\theta) \pi((t_i, t_{-i}) | \theta)}{\sum_{\tilde{t}_{-i}, \tilde{\theta}} \psi(\tilde{\theta}) \pi((t_i, \tilde{t}_{-i}) | \tilde{\theta})}$$

For any combined information structure $S^* = (T \times T', \pi^*)$ for S and S' , write $\lambda_i^{\pi^*}(t_{-i}, \theta | t_i, t'_i)$ for the induced beliefs of player i about (t_{-i}, θ) given a combined type (t_i, t'_i) :

$$\lambda_i^{\pi^*}(t_{-i}, \theta | t_i, t'_i) \triangleq \frac{\sum_{t'_{-i}} \psi(\theta) \pi^*((t_i, t_{-i}), (t'_i, t'_{-i}) | \theta)}{\sum_{\tilde{t}_{-i}, \tilde{\theta}} \psi(\tilde{\theta}) \pi^*((t_i, \tilde{t}_{-i}), (t'_i, \tilde{t}_{-i}) | \tilde{\theta})}, \quad (17)$$

provided that the denominator does not vanish. We say that S is ε -individually sufficient for S' if there exists a combined information structure $S^* = (T \times T', \pi^*)$ with

$$\lambda_i^{\pi^*}(t_{-i}, \theta | t_i, t'_i) - \lambda_i(t_{-i}, \theta | t_i) \leq \varepsilon$$

for all such t'_i, t_i, t_{-i} and θ that (17) are well-defined.

We will now construct a finite basic game such that $G_\varepsilon = ((A_i, u_i)_{i=1}^I, \psi)$ and an outcome $\nu^* : \Theta \rightarrow \Delta(A)$ such that (i) $\nu^* \in BCE(G, S)$ and (ii) $\nu^* \in BCE(G, S')$ imply that S is ε -individually sufficient for S' . Let Ξ_i be any ε -grid of $\Delta(T_{-i} \times \Theta)$, i.e., a finite subset of $\Delta(T_{-i} \times \Theta)$ satisfying the property that, for all $\xi_i \in \Delta(T_{-i} \times \Theta)$, there exists $\xi'_i \in \Xi_i$ with $\|\xi_i - \xi'_i\| \leq \varepsilon$. Now for every player i , let the set of actions be $A_i \triangleq \Xi_i \cup T_i$. Write $\chi_i(a_i)$ for the belief over $T_{-i} \times \Theta$ naturally associated with a_i , so $\chi_i : A_i \rightarrow \Delta(T_{-i} \times \Theta)$ is defined by

$$\chi_i(a_i) \triangleq \begin{cases} \lambda_i(\cdot | t_i), & \text{if } a_i = t_i \in T_i; \\ \xi_i, & \text{if } a_i = \xi_i \in \Xi_i. \end{cases}$$

Now let the payoff function of each player i be:

$$u_i((a_i, a_{-i}), \theta) \triangleq \begin{cases} 2\chi_i(t_{-i}, \theta | a_i) - \sum_{\tilde{t}_{-i} \in T_{-i}, \tilde{\theta} \in \Theta} \left(\chi_i(\tilde{t}_{-i}, \tilde{\theta} | a_i) \right)^2, & \text{if } a_{-i} = t_{-i} \in T_{-i}; \\ 0, & \text{if } a_{-i} \notin T_{-i}. \end{cases}$$

Thus if others' actions are within T_{-i} , utility function u_i gives player i an incentive to choose an action associated with his true beliefs via a quadratic scoring rule. More precisely, suppose player i assigns

probability 1 to his opponents choosing $a_{-i} \in T_{-i}$ and, in particular, for some $\xi_i \in \Delta(T_{-i} \times \Theta)$, assigns probability $\xi_i(t_{-i}, \theta)$ to his opponents choosing $a_{-i} = t_{-i} \in T_{-i}$ and the state being θ . The expected payoff to player i with this belief over $A_{-i} \times \Theta$ of choosing an action a_i with $\chi_i(a_i) = \xi'_i$ is

$$\begin{aligned}
& \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} 2\xi'_i(t_{-i}, \theta) \left(\xi'_i(t_{-i}, \theta) - \sum_{\tilde{t}_{-i} \in T_{-i}, \tilde{\theta} \in \Theta} \left(\xi'_i(\tilde{t}_{-i}, \tilde{\theta}) \right)^2 \right) \\
&= 2 \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \xi_i(t_{-i}, \theta) \xi'_i(t_{-i}, \theta) - \sum_{\tilde{t}_{-i} \in T_{-i}, \tilde{\theta} \in \Theta} \left(\xi'_i(\tilde{t}_{-i}, \tilde{\theta}) \right)^2 \\
&= 2 \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \xi_i(t_{-i}, \theta) \xi'_i(t_{-i}, \theta) - \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \left(\xi'_i(t_{-i}, \theta) \right)^2 \\
&= \left(\|\xi_i\|^2 - \|\xi'_i - \xi_i\|^2 \right).
\end{aligned}$$

Now the game (G, S) has - by construction - a "truth-telling" *BCE* where each type t_i always chooses action t_i . This gives rise to an outcome ν^* where

$$\nu^*(a|\theta) = \begin{cases} \pi(a|\theta), & \text{if } a = t \in T; \\ 0, & \text{if otherwise.} \end{cases}$$

So ν^* is a BCE outcome of (G, S) . For ν^* to be BCE outcome of (G, S') , there must exist a BCE of (G, S') , $\sigma' : \Theta \times T' \rightarrow \Delta(T)$, inducing ν^* . Note that formally, the BCE of (G, S') is specified by $\sigma' : \Theta \times T' \rightarrow \Delta(A)$, but if $a \in A$ is such that $a \notin T$, then the fact that σ' induces ν^* implies that

$$\sum_{t' \in T'} \pi'(t'|\theta) \sigma'(a|t'(a|\theta)) = 0.$$

Therefore, if $\pi'(t'|\theta) > 0$, then $\sigma'(a|t', \theta) = 0$. Now setting

$$\pi^*(t, t'|\theta) = \pi'(t'|\theta) \sigma'(t|t', \theta),$$

information structure $S^* = (T \times T', \pi^*)$ is a combined information structure for S and S' . Obedience constraints imply that

$$\left\| \lambda_i^{\pi^*}(\cdot|t_i, t'_i) - \lambda_i(\cdot|t_i) \right\| \leq \varepsilon.$$

Thus S is ε^2 -individually sufficient for S' .

But now S being more incentive constrained than S' requires that $BCE(G_\varepsilon, S) \subseteq BCE(G_\varepsilon, S')$ for all such games G_ε , and thus that S be ε^2 -individually sufficient for S' for all $\varepsilon > 0$. But because the set of mappings of combined information structures, $\pi^* : \Theta \rightarrow \Delta(T \times T')$, is a compact set, if S is ε^2 -individually sufficient for S' for each $\varepsilon > 0$, S is individually sufficient for S' . ■

5 Discussion

5.1 Obedience and Incomplete Information Correlated Equilibrium

Aumann [1974], [1987] introduced the notion of correlated equilibrium for complete information games. A correlated equilibrium is a joint distribution over actions such that each player's action is optimal for that player if all the player knew is the action he is playing and the joint distribution over actions. Bayes correlated equilibrium is the natural incomplete information generalization where we (i) add incomplete information; and (ii) require that players' actions be optimal when they condition on their type as well as their equilibrium action. This is the obedience condition. Thus Bayes correlated equilibrium captures only the role of information in tightening obedience constraints. Theorem 1 formalizes this motivation for studying Bayes correlated equilibrium: the solution concept captures rational behavior given that players have access to the signals in the information structure, but may have additional information.

The existing literature on incomplete information correlated equilibrium has focussed on additional restrictions on behavior that capture the idea that players are constrained by what information is available to them. To put our solution concept in context, we report some key feasibility restrictions imposed in the literature. A decision rule σ is *belief invariant* if, for each player i , the probability distribution over player i 's actions that it induces depends only on player i 's type, and is independent of other players' types and the state. Writing $\sigma_i : T \times \Theta \rightarrow \Delta(A_i)$ for the probability distribution over player i 's actions induced by σ ,

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

decision rule σ is belief invariant for (G, S) if, for each player i , $\sigma_i(a_i | (t_i, t_{-i}), \theta)$ is independent of t_{-i} and θ . An equivalent statement is that player i 's beliefs about (t_{-i}, θ) conditional on t_i do not depend on a_i . In the language of mediation it says that the mediator's recommendation does not give a player any additional information about the state and other players' types. The condition of belief invariance was introduced in this form and so named by Forges [2006]. If a decision rule σ is belief invariant for (G, S) , then players have no less but also no more information under σ and S than under information structure S . If we impose belief invariance as well as obedience on a decision rule, we get a solution concept that was introduced in Liu [2015].

Definition 8 (Belief Invariant BCE)

Decision rule σ is a belief invariant Bayes correlated equilibrium of (G, S) if it is obedient and belief invariant for (G, S) .

It captures the implications of common knowledge of rationality and that players know exactly the

information contained in S (and no more) if the common prior assumption is maintained.¹⁰ The set of Bayes correlated equilibria of (G, S) is the union of all belief invariant BCE of (G, S') for all information structures S' which are individually sufficient for S . Liu [2015] showed that if two information structures have the same canonical representation, then they have the same set of belief invariant Bayes correlated equilibria. This in turn implies that they have the same set of Bayes correlated equilibria.

Much of the literature on incomplete information correlated equilibrium started from the premise that an incomplete information definition of correlated equilibrium should capture what could happen if players had access to a correlation device / mediator under the maintained assumption that the correlation device/mediator did not have access to information that was not available to the players. We can describe the assumption formally as:

Definition 9 (Join Feasible)

Decision rule σ is join feasible for (G, S) if $\sigma(a|t, \theta)$ is independent of θ .

Thus the probability of a profile of action recommendations conditional on the players' type profile is independent of the state. If join feasibility but not belief invariance is assumed, we get another solution concept:

Definition 10 (Bayesian Solution)

Decision rule σ is a Bayesian solution of (G, S) if it is obedient and join feasible.

Join feasibility was implicitly assumed in Forges [1993] and other works, because it was assumed that type profiles exhausted payoff relevant information;¹¹ Lehrer et al. [2010], [2013] explicitly impose this assumption. The Bayesian solution was named by Forges [1993] and it is the weakest version of incomplete information correlated equilibrium she studies. Imposing both join feasibility and belief invariance, we get a solution concept that has been an important benchmark in the literature.

Definition 11 (Belief Invariant Bayesian Solution)

Decision rule σ is a belief invariant Bayesian solution of (G, S) if it is obedient, belief invariant and join feasible.

¹⁰Liu [2015] notes that this solution concept can be seen as the common prior analogue of the solution concept of interim correlated rationalizability discussed by Dekel et al. [2007]; see Bergemann and Morris [2014b] for a discussion of the rationalizability notion which is the non-common prior analogue of Bayes correlated equilibrium.

¹¹The issue is discussed in Section 4.5 of Forges [1993], where she notes how analyzing a "reduced form" game is not innocuous in general. But in many natural economic settings, type profiles do exhaust payoff relevant information and in those cases, there is an equivalence between Bayes correlated equilibria and Bayesian solutions. See Bergemann and Morris [2014b] for a discussion of this issue and economic settings (such as private values environments) where join feasibility is a natural maintained assumption.

Forges [2006] introduced this name. The other incomplete information correlated equilibrium solution concepts for an incomplete information game in Forges [1993], [2006] - communication equilibrium, agent normal form correlated equilibrium and strategic form correlated equilibrium - are all strictly stronger than the belief invariant Bayesian solution, by imposing additional "truth-telling" constraints (for communication equilibrium), feasible correlation structure constraints (for agent normal form correlated equilibrium) and a combination of the two (for strategic form correlated equilibrium). Forges [1993] also discusses a "universal Bayesian solution" which corresponds to Bayes correlated equilibrium in the case where S is degenerate, i.e., there is no prior information structure (beyond the common prior over payoff states).

5.2 Alternative Orderings on Many Player Information Structures and Their Uses

If we fix a pair of information structures S and S' , a combined information structure for these two information structures, and a prior on states, we generate a probability distribution on the space $T \times T' \times \Theta$. We can identify a variety of conditional independence properties that we might be interested in on that space:

- (i) the distribution of t'_i conditional on t_i is independent of θ for each i ;
- (ii) the distribution of t'_i conditional on t_i is independent of (t_{-i}, θ) for each i ;
- (iii) the distribution of t'_i conditional on t_i is independent of $(t_{-i}, t'_{-i}, \theta)$ for each i ;
- (iv) the distribution of t' conditional on t is independent of θ .

In the one-player case, these four conditions are all equivalent to each other (and to Blackwell's order). In the many-player case, they are all different from each other. Intuitively, condition (i) requires only that information structure S' conveys no new information to any player about the state; condition (ii) requires that information structure S' conveys no new information to any player about the state and higher order beliefs about the state; condition (iii) requires that information structure S' conveys no new information to any player about the state, higher order beliefs about the state, and redundant signals that other players may be observing; condition (iv) requires that information structure S' conveys no new information about the state to the players collectively (combining their information) that they did not collectively possess before. The exact relation between them is subtle: one can verify that (iii) \Rightarrow (ii) \Rightarrow (i) and (iii) \Rightarrow (iv) but there are no further implications relating these conditional independence properties. In particular, in Example 2 in Section 3, information structure S was individually sufficient for S' , and thus conditional independence (ii) was satisfied, but one can verify that (iv) fails not only in the particular combined

information structure used to establish individual sufficiency but also in any other combined information structure.

We can understand the related literature by comparing *which* conditional independence properties are required and in *which* combined experiments. We showed that information structure S gives rise to fewer Bayes correlated equilibrium outcomes than information structure S' if and only if there exists a combined information structure such that (ii) holds. Gossner [2000] asked when information structure S gives rise to fewer Bayes Nash equilibrium outcomes than information structure S' .¹² He showed that this is true if and only if there exists a combined information structure such that two conditions hold, namely (ii) t'_i conditional on t_i is independent of (t_{-i}, θ) for each player i) and the stronger condition (iii) in its reverse form, i.e., t_i conditional on t'_i is independent of $(t_{-i}, t'_{-i}, \theta)$ for each i .¹³ This combination of conditions implies that S and S' must have the same canonical representation. Intuitively, this is because feasibility considerations (implicit in the definition of Bayes Nash equilibrium) require that information structure S' contain at least as much information about beliefs and higher order beliefs as S and incentive considerations require that S contain at least as much information about beliefs and higher order beliefs as S' . However, Gossner's characterization also requires that S' has more information about redundant information than S , i.e., more correlation possibilities. Thus Gossner does not order information structures with distinct canonical representations and shows how more redundant information, i.e., correlation possibilities, must lead to a larger set of Bayes Nash equilibrium outcomes. We show that redundant information does not affect the set of Bayes correlated equilibrium outcomes and more payoff relevant information must lead to a smaller set of Bayes correlated equilibrium outcomes.

It is an implication of Gossner [2000] that two information structures give rise to the same set of Bayes Nash equilibrium outcomes if and only if there is a single combined information structure where (iii) holds (t'_i conditional on t_i is independent of $(t_{-i}, t'_{-i}, \theta)$ for each i) and its reverse holds, i.e., t_i conditional on t'_i is independent of $(t_{-i}, t'_{-i}, \theta)$ for each i . Lehrer et al. [2013] show that this result remains true if the conditional independence properties hold in distinct combined information structures, i.e., there exists one combined information structure where t'_i conditional on t_i is independent of $(t_{-i}, t'_{-i}, \theta)$ for each i ,¹⁴ and another combined information structure where t_i conditional on t'_i is independent of $(t_{-i}, t'_{-i}, \theta)$ for each i . Thus Lehrer et al. [2013] show that it is without loss of generality to require that the conditional independence properties hold in the same combined information structure.¹⁵ Lehrer et al. [2013] also

¹²Gossner and Mertens [2001] and Peski [2008] characterize the value of information in zero sum games.

¹³In this case, Gossner [2000] says that "there is a faithful and compatible interpretation from S to S' ".

¹⁴In this case, Lehrer et al. [2013] say that "there is an independent garbling from S to S' ".

¹⁵As we noted in footnote 7, this argument can be adapted to show that if S is individually sufficient for S' and S' is individually sufficient for S , we can without loss of generality establish both directions of individual sufficiency using the same combined information structure, and thus the two information structures have the same canonical representation.

establish analogous results for solution concepts that are intermediate between Bayes Nash equilibrium and Bayes correlated equilibrium. Thus they show that two information structures give rise to the same set of belief invariant Bayesian solution outcomes if and only if there exists a combined information structure where (ii) and (iv) hold,¹⁶ and another combined information structure where the reverse properties hold.

5.3 The One-Player Special Case and Many-Player Bayesian Persuasion

Our results apply to the case of one player. In the one-player case, a basic game reduces to a *decision problem*, mapping actions and states to a payoff of the *decision maker*. An information structure corresponds to an *experiment* in the sense of Blackwell [1951], [1953]. A decision rule is now a mapping from state and signals to probability distributions over actions. A decision rule is a Bayes correlated equilibrium if it is obedient. To interpret obedience, consider a decision maker who observed a signal under the experiment and received an action recommendation chosen according to the decision rule. The decision rule is obedient if he would have an incentive to follow the recommendation. Theorem 1 states that the set of Bayes correlated equilibria for a fixed decision problem and experiment equals the set of decision rules from a decision maker choosing an optimal action with access to that experiment and possibly more information (an expanded experiment). Thus Bayes correlated equilibria capture all possible optimal behavior if the decision maker had access to the fixed experiment and perhaps some additional information.

Now consider the case where the original information structure is degenerate (there is only one signal which represents the prior over the states of the world). In this case, the set of Bayes correlated equilibria correspond to joint distributions of actions and states that could arise under rational choice by a decision maker with any information structure. Kamenica and Gentzkow [2011] consider a problem of "Bayesian persuasion". Suppose a "sender" could pick the experiment that the decision maker, the "receiver", could observe. Kamenica and Gentzkow [2011] characterize the set of joint distributions over states and actions that the sender could induce through picking an experiment and having the decision maker choose optimally. This set is exactly what we label Bayes correlated equilibria. They can then analyze which (in our language) Bayes correlated equilibrium the sender would prefer to induce in a variety of applications.

Thus if we want extend Bayesian persuasion to the case of many receivers who have some prior information, the set of Bayes correlated equilibria is the set of outcomes that can be induced. In the Appendix, we report an example to illustrate optimal multi-player Bayesian persuasion use of our characterization.¹⁷

¹⁶In this case, Lehrer et al. [2013] say that "there is a non-communicating garbling from S to S' ".

¹⁷Caplin and Martin [2013] introduce a theoretical and experimental framework for analyzing imperfect perception. The set of joint distributions over states and actions that arise in their framework also correspond to one player Bayes correlated equilibria where S is null.

5.4 Feasibility and Blackwell's Theorem

Our Theorem 2 relates a *statistical* ordering (individual sufficiency) to an *incentive ordering* (more incentive constrained). More information leads to a *smaller* set of Bayes correlated equilibria because it adds incentive constraints. Information is unambiguously "bad" in the sense of reducing the set of possible outcomes. Lack of information is never a constraint on what is feasible for players because the solution concept of Bayes correlated equilibrium imposes no feasibility constraints on players' behavior.

On the other hand, Blackwell's Theorem relates a statistical ordering to a *feasibility* ordering. In the one-player case, more information is "good" in the sense of leading to more feasible joint distributions of actions and states and thus (in the one-player case) to higher ex-ante utility. Incentive constraints do not bind, because there is a single decision maker. In this Section, we will report a result which relates our statistical ordering to a feasibility ordering in the many-player case. The approach and result is a straightforward variation on the work of Lehrer et al. [2010], so we report the result without formal proof.

Say that basic game G has common interests if $u_1 = u_2 = \dots = u_I = u^*$. Fix a common interest basic game G and an information structure S . Recall from Section 5.1 that a decision rule is belief invariant for (G, S) if, for each player, the distribution of his action depends only on his type and is independent of others' types and the state. Let $v(G, S)$ be the highest possible (common) ex-ante utility that is attained by players under a belief invariant decision rule:

$$v(G, S) \triangleq \max_{\sigma \in \Sigma_B} \sum_{a, t, \theta} \psi(\theta) \pi(t|\theta) \sigma(a|t, \theta) u^*(a, \theta), \quad (18)$$

where

$$\Sigma_B \triangleq \{ \sigma : T \times \Theta \rightarrow \Delta(A) \mid \sigma \text{ is belief invariant for } (G, S) \}.$$

Thus we are asking what is the highest (common) payoff that players could obtain if they were able to correlate their behavior but could only do so using correlation devices in the sense of Liu [2015] under which a player's action recommendation gives him no additional information about others' types and the state. Here the information structure is constraining (through belief invariance) the set of joint distributions over actions and states that can arise. Say that an information structure S is more valuable than S' if, in every common interest basic game G , there is a belief invariant decision rule for (G, S) that gives a higher common ex-ante payoff than any belief invariant decision rule for (G, S') .

Definition 12 (More Valuable)

Information structure S is more valuable than information structure S' if, for every common interest basic game G , $v(G, S) \geq v(G, S')$.

Now we have:

Theorem 3

Information structure S is individually sufficient for information structure S' if and only if S is more valuable than S' .

Notice that obedience constraints do not arise in any of the properties used to state this theorem. In that sense, the Theorem relates a statistical ordering to a feasibility ordering and does not make reference to incentive compatibility constraints. But also notice that, since the game has common interests, the belief invariant decision rule that is the argmax of expression (18) will automatically satisfy obedience. Recall from Definition 8 that a decision rule is a belief invariant Bayes correlated equilibrium of (G, S) if it satisfies belief invariance and obedience. Thus $v(G, S)$ is also the ex-ante highest common payoff that can be obtained in a belief invariant Bayes correlated equilibrium.

In the special case of one player, Theorem 3 clearly reduces to the classic statement of Blackwell's theorem favored by economists. In the many-player case, it follows from the arguments of Lehrer et al. [2010]. We can sketch a direct proof of the harder direction of Theorem 3. Manipulations of definitions shows that S is individually sufficient for information structure S' if and only if the set of outcomes induced by belief invariant decision rules for (G, S) is larger than that set for (G, S') . In other words, for any action sets for the players, S supports a larger set of feasible outcomes than S' . Since these sets are compact and convex, the separating hyperplane theorem implies we can choose a common utility function such that ex-ante expected utility is higher under S than under S' .

6 Appendix: Applications and a Binary Action Game

We have used Bayes correlated equilibria and the interpretation suggested by Theorem 1 in a variety of applications. In Bergemann and Morris [2013a] and Bergemann et al. [2015c], we have considered games played by a continuum of players, with symmetric payoffs and linear best responses, and focussed on symmetric equilibria. In our work on third degree price discrimination, Bergemann et al. [2015a], we exploit the fact that the outcomes of third degree price discrimination correspond to one-player Bayes correlated equilibria. In Bergemann et al. [2015b] we use the results to look at all outcomes that could arise for different information structures that players have in a first price auction.

Rather than trying to review this work, we instead present a " $2 \times 2 \times 2$ " basic game, where there are two players, two actions for each player and two states, to illustrate the structure of Bayes correlated equilibria and Theorem 1 and the orderings of information structures in Theorem 2. In particular, we identify, in this class of games, the expanded information structures that support or "decentralize" welfare maximizing Bayes correlated equilibria as Bayes Nash equilibria. The role of strategic substitutes and strategic complements in these results complements welfare results in Bergemann and Morris [2013a]. Even in this simple class of games, the analysis becomes algebraically quite involved. This is not surprising, given the demonstration in Calvó-Armengol [2006] that – even in complete information games – characterizing and visualizing all correlated equilibria of all two player two actions games are not easy. We restrict attention to a two dimensional class of (symmetric) decision rules. We emphasize that we are using this class of examples to illustrate results that apply to general, asymmetric, information structures and general, asymmetric, decision rules. We analyzed a related class of such $2 \times 2 \times 2$ basic games in Bergemann and Morris [2013b], as does Taneva [2014].

A Binary Investment Game Each player can either invest, $a = I$ or not invest, $a = N$ and the payoffs are given in the bad state θ_B and the good state θ_G by the following matrices:

$$\begin{array}{|c|c|c|} \hline \theta_B & I & N \\ \hline I & z - 1 + y_B, z - 1 + y_B & -1, z \\ \hline N & z, -1 & 0, 0 \\ \hline \end{array} \quad \begin{array}{|c|c|c|} \hline \theta_G & I & N \\ \hline I & z + 1 + y_G, z + 1 + y_G & 1, z \\ \hline N & z, 1 & 0, 0 \\ \hline \end{array} . \quad (19)$$

The payoffs are symmetric across players and have three components: (i) there is a payoff 1 to invest in the good state θ_G and a payoff -1 to invest in the bad state θ_B ; (ii) there is always an externality $z > 0$ if the other player invests, and (iii) there is an additional, possibly state dependent payoff y_j , $j = B, G$, to investing if the other player invests as well. The payoff y_j can be positive or negative, but of uniform sign across states, leading to a game with strategic complements or substitutes, respectively. We will focus

on the case where $z \gg 1$ and $y_j \approx 0$.¹⁸ Thus, if the players were to know the state, i.e. under complete information, then each player would have a strict dominant strategy to invest in θ_G and not to invest in θ_B . Importantly, given that the externality z is assumed to be large, i.e. $z \gg 1$, the sum of the payoffs is maximized if both players invest in both states, θ_B and θ_G . Notice that z is a pure externality that influences players' utilities but not the best responses and a fortiori not the set of BCE. Finally, we assume that state θ_G occurs with probability ψ , while state θ_B occurs with probability $1 - \psi$.

A Binary Information Structure We consider a binary information structure S where, if the state is bad, each player observes a signal t_b , saying that the state is bad, with probability q . If a player doesn't receive the signal t_b , then he receives the signal t_g , and thus the signal t_g is always observed in the good state. The signal distribution $\pi : \Theta \rightarrow \Delta(T)$ is given by:

$$\begin{array}{|c|c|c|} \hline \pi(\cdot | \theta_B) & t_b & t_g \\ \hline t_b & q & 0 \\ \hline t_g & 0 & 1 - q \\ \hline \end{array} \quad \begin{array}{|c|c|c|} \hline \pi(\cdot | \theta_G) & t_b & t_g \\ \hline t_b & 0 & 0 \\ \hline t_g & 0 & 1 \\ \hline \end{array}. \tag{20}$$

Thus signal realizations are perfectly correlated. The information structure is symmetric across players, but not across states. In particular, the conditional probability q is a measure of the accuracy of the information structure. An increase in q leads, after a realization of t_g to a strict increase in the posterior probability that the state is θ_G , and after a realization of t_b the posterior probability that the state is θ_B is always 1 (and thus is *weakly* increasing in q).

We restrict attention to decision rules σ , as defined earlier in (1), that are symmetric across players. Accordingly, we must specify the action profile for each state-signal profile (θ, t) . After observing the negative signal t_b , each player knows that the state is θ_B and has a strictly dominant strategy to choose N , so we will take this behavior as given. We can parameterize the symmetric (across players) decision rule σ conditional on the positive signals t_g and the state θ_j , for $j = B, G$, by:

$$\begin{array}{|c|c|c|} \hline \sigma(\theta_j, t_g) & \text{I} & \text{N} \\ \hline \text{I} & \gamma_j & \alpha_j - \gamma_j \\ \hline \text{N} & \alpha_j - \gamma_j & \gamma_j + 1 - 2\alpha_j \\ \hline \end{array}. \tag{21}$$

We thus have four parameters, $\alpha_B, \alpha_G, \gamma_B, \gamma_G$. The probability that any one player invests in state θ_j is α_j and γ_j is the probability that both players invest. These parameters are subject to the nonnegativity restrictions:

$$\alpha_j \geq 0, \gamma_j \geq 0, \text{ and } 2\alpha_j - 1 \leq \gamma_j \leq \alpha_j, \text{ for } j = B, G. \tag{22}$$

¹⁸Formally, we require that $z > 1$ and that $z > 1 - 2y_j$ for $j = B, G$.

The set of parameters $(\alpha_B, \alpha_G, \gamma_B, \gamma_G)$ which form a Bayes correlated equilibrium are completely characterized by the obedience conditions for $a = I, N$. Thus, explicitly, if a player is observing signal t_g and is advised to invest, then he will invest if:

$$\psi(\alpha_G + \gamma_G y_G) + (1 - \psi)(1 - q)(-\alpha_B + \gamma_B y_B) \geq 0; \quad (23)$$

and a player advised to not invest will not invest if:

$$\psi(1 - \alpha_G + (\alpha_G - \gamma_G) y_G) + (1 - \psi)(1 - q)(-(1 - \alpha_B) + (\alpha_B - \gamma_B) y_B) \leq 0. \quad (24)$$

Welfare Maximizing Bayes Correlated Equilibria We will first focus on the characterization of the "second-best BCE" which maximizes the sum of players' utility subject to being a BCE and then describe the expanded information structures that can achieve the Bayes correlated equilibrium as a Bayes Nash equilibrium. To this end, it suffices to identify the parameters $(\alpha_B, \alpha_G, \gamma_B, \gamma_G)$ that maximize the expected utility of a (representative) player:

$$\psi(\alpha_G(z + 1) + \gamma_G y_G) + (1 - \psi)(1 - q)(\alpha_B(z - 1) + \gamma_B y_B) \quad (25)$$

subject to the obedience conditions (23) and (24) and the nonnegativity restrictions (22). In the analysis it will prove useful to distinguish between the strategic complements, $y_j \geq 0$, and strategic substitutes, $y_j \leq 0$.

Strategic Complements We begin with strategic complements. As a player never invests after observing the negative signal t_b , after correctly inferring that the state is θ_B , we immediately ask under what conditions investment can occur after the realization of the positive signal t_g . If investment could always be achieved, independent of the true state, then the resulting decision rule σ would have $\alpha_G = \gamma_G = \alpha_B = \gamma_B = 1$, and inserting these values in the obedience constraint for investing, see (23), yields:

$$\psi(1 + y_G) + (1 - \psi)(1 - q)(y_B - 1) \geq 0 \quad \Leftrightarrow \quad q \geq 1 - \frac{\psi}{1 - \psi} \frac{1 + y_G}{1 - y_B}. \quad (26)$$

Thus, if the information structure S , as represented by q , is sufficiently accurate, then investment following the realization of the signal t_g can be achieved with probability one. In fact, the above condition (26) is a necessary and sufficient condition for a Bayes Nash equilibrium with investment after the signal t_g to exist. Hence, we know that this decision rule can be informationally decentralized without any additional information if q is sufficiently large.

By contrast, if q fails to satisfy the condition (26), then the second-best BCE is to maintain investment in the good state: $\alpha_G = \gamma_G = 1$, while maximizing the probability of investment α_B in the bad state

subject to the obedience constraint (23).¹⁹ In a game with strategic complements, this is achieved by coordinating investments, i.e. setting $\alpha_B = \gamma_B$ and satisfying (23) as an equality:

$$\psi(1 + y_G) - (1 - \psi)(1 - q)\alpha_B(1 - y_B) = 0 \Leftrightarrow \alpha_B = \gamma_B = \frac{\psi(1 + y_G)}{(1 - \psi)(1 - q)(1 - y_B)}. \quad (27)$$

Now, we observe that the solution (27) requires the probabilities to differ across the states, or $\alpha_B = \gamma_B < \alpha_G = \gamma_G = 1$. It follows that this decision rule requires additional information, and hence an expansion of the information structure S for it to be decentralized as a Bayes Nash equilibrium. The necessary expansion is achieved by two additional signals t'_b, t'_g which lead to an expansion S^* and an associated signal distribution $\pi^*(t, t' | \theta)$ as displayed below:

$\pi^*(\cdot \theta_B)$	t_b, t'_b	t_g, t'_b	t_g, t'_g
t_b, t'_b	q	0	0
t_g, t'_b	0	r	0
t_g, t'_g	0	0	$1 - q - r$

$\pi^*(\cdot \theta_G)$	t_b, t'_b	t_g, t'_g
t_b, t'_b	0	0
t_g, t'_g	0	1

We observe that the expansion preserves the public nature of the signals, in that the realizations remain perfectly correlated across the players. The additional signals confirm the original signals everywhere except for the pair (t_g, t'_b) which changes the posterior of each player to a probability one belief that the state is θ_B . In other words, the additional signals t'_b, t'_g “split” the posterior conditional on receiving t_g in the information structure S . We can readily compute the minimal probability that the public signal (t_g, t'_b) has to have so that in the associated BNE the players invest with probability one after receiving the signal (t_g, t'_g) , namely by requiring that the best response for investment is met as an equality in the BNE:

$$\psi(1 + y_G) - (1 - \psi)(1 - q - r)(1 - y_B) = 0 \Leftrightarrow r = 1 - q - \frac{\psi(1 + y_G)}{(1 - \psi)(1 - y_B)}.$$

Strategic Substitutes Next, we discuss the game with strategic substitutes, $y_j \leq 0$. While the basic equilibrium conditions remain unchanged, the information structures that decentralize the second-best BCE have very different properties with strategic substitutes. In particular, private rather than public signals become necessary to decentralize the decision rule σ as a Bayes Nash equilibrium.

To begin with, just as in the case of strategic complements, if the information structure S , as represented by q , is sufficiently accurate, then investment following the realization of the signal t_g can be achieved with probability one, this is the earlier condition (26). Similarly, if q fails to satisfy the condition (26) then the second-best BCE is to maintain investment in the good state: $\alpha_G = \gamma_G = 1$, while maximizing the probability of investment α_B in the bad state subject to the obedience constraint (23). But importantly, in

¹⁹The no investment constraint (24) will automatically be satisfied.

a game with strategic substitutes, the obedience constraint is maintained by minimizing the probability of joint investments, hence *minimizing* γ_B . In terms of the decision rule $\sigma(\cdot, t_g)$ as represented in the matrix (21), we seek to place most probability *off* the diagonal, in which only one, but not both players, invest. If there is substantial slack in the obedience constraint (23) when evaluated at zero entries *on* the diagonal of (21), then the residual probability can lead to investment by both players, but if there is little slack, then it will require that the residual probability leads to no investment by either player, which suggests a second threshold for q , beyond the one established in (26).

Thus if condition (26) fails, then it is optimal to maximize α_B and minimize γ_B , where the latter is constrained by the nonnegativity restrictions of (22): $\gamma_B = \max\{0, 2\alpha_B - 1\}$. Thus we want α_B to solve the obedience constraint for investment, (23), as an equality:

$$\psi(1 + y_G) + (1 - \psi)(1 - q)(-\alpha_B + \max(2\alpha_B - 1, 0)y_B) = 0.$$

This leads to a strictly positive solution of γ_B , the probability of joint investment, as long as the probability q is not too low, or

$$1 - \frac{2\psi}{1 - \psi}(1 + y_G) \leq q \leq 1 - \frac{\psi}{1 - \psi} \frac{1 + y_G}{1 - y_B}, \quad (28)$$

and the second-best decision rule given by:

$$\alpha_G = \gamma_G = 1, \alpha_B = \frac{1}{1 - 2y_B} \left(\frac{\psi(1 + y_G)}{(1 - \psi)(1 - q)} - y_B \right), \gamma_B = 2\alpha_B - 1.$$

Finally, if q falls below the lower threshold established in (28), then the second-best decision rule σ prescribes investment only by one player, but never by both players simultaneously:

$$\alpha_G = \gamma_G = 1, \quad \alpha_B = \frac{\psi(1 + y_G)}{(1 - \psi)(1 - q)}, \quad \gamma_B = 0. \quad (29)$$

As expected, we find that both the probability of investment by a player, given by α_B , as well as the probability of a joint investment, γ_B , are increasing in the accuracy q .

We ask again which expanded information structures decentralize these second-best decision rules. As $\gamma_B < \alpha_B$, the decision rule σ requires with positive probability investment by one player only. This can only be achieved by private signals that lead to distinct choices by the players with positive probability. The expansion can still be achieved with two additional signals, t'_b, t'_g , and as before the additional signals refine or split the posterior that each player held at t_g in the information structure S . But importantly, now the signal realizations cannot be perfectly correlated across the players anymore. Thus if q is not too low, i.e. condition (28) prevails, then the following information structure decentralizes the second-best

BCE:

$\pi^*(\cdot \theta_B)$	t_b, t'_b	t_g, t'_b	t_g, t'_g
t_b, t'_b	q	0	0
t_g, t'_b	0	0	r
t_g, t'_g	0	r	$1 - q - 2r$

$\pi^*(\cdot \theta_G)$	t_b, t'_b	t_g, t'_g
t_b, t'_b	0	0
t_g, t'_g	0	1

and by contrast if q is sufficiently low, i.e. below the lower bound of (28), then the expanded information structure below decentralizes the BCE:

$\pi^*(\cdot \theta_B)$	t_b, t'_b	t_g, t'_b	t_g, t'_g
t_b, t'_b	q	0	0
t_g, t'_b	0	$1 - q - 2r$	r
t_g, t'_g	0	r	0

$\pi^*(\cdot \theta_G)$	t_b, t'_b	t_g, t'_g
t_b, t'_b	0	0
t_g, t'_g	0	1

In either case, the expansion requires private signals in the sense that conditional on receiving a given signal, either (t_g, t'_g) or (t_g, t'_b) respectively, each player remains uncertain as to the signal received by the other player, i.e. either (t_g, t'_b) or (t_g, t'_g) . As required, the expanded information structure S^* preserves the likelihood distribution ψ of the initial information structure S .²⁰

The Set of all Symmetric Bayes Correlated Equilibria The above analysis focussed on second-best Bayes correlated equilibria that maximize welfare. We now visualize all symmetric Bayes correlated equilibria in a special case. We stay with a game of strategic substitutes, $y_j \leq 0$ and the illustrations below are computed for the prior probability of the good state, $\psi = 1/3$ and $z = 2, y_G = 0, y_B = -1/6$. Because there is never investment conditional on bad signals, it is enough the focus on the probabilities α_G and α_B that any player invests, conditional on good and bad states respectively, after observing the positive signal t_g . Figures 1 through 3 show the set of all values of α_G and α_B corresponding to symmetric BCE for low, intermediate and high levels of accuracy q , namely $q = 1/5, 11/20$ and $4/5$, respectively.

The set of Bayes correlated equilibria for the binary games is completely characterized by the obedience constraints (23) and (24), given the parameterized decision rule σ , see (21). The detailed computations for the present example are recorded in Appendix B of Bergemann and Morris [2014a].

For all values of $q \in [0, 1]$, the action profile that maximizes the sum of the payoffs is $\alpha_B = \alpha_G = 1$, the first-best action profile. Every Bayes Nash equilibrium under the given information structure S has to be located on the 45° line, as each player cannot distinguish between the states θ_B and θ_G conditional on

²⁰An interesting question that we do not explore in any systematic manner in this paper is what we can say about the relation between Bayes correlated equilibria and the expansions that are needed to support them as Bayes Nash equilibria. Milchtaich [2012] examines properties of devices needed to implement correlated equilibria, and tools developed in his paper might be useful for this task.

t_g . In fact, the symmetric Bayes Nash equilibrium in the game with strategic substitutes is unique for all levels of q , and depending on the accuracy q , it is either a pure strategy equilibrium with no investment as in Figure 1, a mixed strategy equilibrium with positive probability of investment as in Figure 2, or a pure strategy equilibrium with investment as in Figure 3, respectively. By contrast, the second-best BCE, as computed by (29), always yields a strictly positive level of investment in the bad state θ_B , and one that is strictly higher than in any BNE, unless the BNE itself is a pure strategy equilibrium with investment (following t_g), see Figure 3.




Figure 1: BNE and set of BCE with low accuracy: $q = 1/5$.

If we consider an intermediate level of accuracy q , rather than a low level of accuracy q , as in Figure 2, then we find that there is unique mixed BNE which provides investment with positive probability following t_g . The BNE is therefore in the interior of the unit square of conditional investment probabilities (α_G, α_B) . By contrast, the second-best BCE remains at the boundary of the unit square, and yields a strictly higher probability of investment in the bad state than the corresponding Bayes Nash equilibrium. Interestingly, the BNE is in the interior of the set of BCE, when expressed in the space of investment probabilities rather than an extreme point of the set of BCE. If the accuracy of the information structure increases even further, see Figure 3, then conditional on receiving the positive signal t_g , it is sufficiently likely that the state is θ_G , and investment occurs with probability one even in the Bayes Nash equilibrium. Essentially, the high probability of θ_G (and resulting high payoffs from investment) more than offset the low probability of θ_B (and resulting low payoffs from investment).

This first set of illustrations depicts the probabilities of investment conditional on the realization of the positive signal t_g and the state $\theta_j, j = B, G$. But as we vary the accuracy q , we are changing the probability of the signal t_g , and hence the above figures do not directly represent the probabilities of investment β_j conditional on the state θ_j only, which are simply given by $\beta_B = (1 - q) \alpha_B$ and $\beta_G = \alpha_G$. The resulting




Figure 2: BNE and set of BCE with intermediate accuracy: $q = 11/20$.




Figure 3: BNE and set of BCE with high accuracy: $q = 4/5$.

sets of investment probabilities are depicted in Figure 4, for all three levels of q . The set of BCE is shrinking as the information structure S , as represented by q , becomes more accurate. This comparative static illustrates Theorem 2. Because the set of BCE is shrinking, the best achievable BCE welfare is necessarily getting weakly lower with more information and, in this example, is strictly lower. On the other hand, as q increases, welfare in BNE will increase over some ranges.

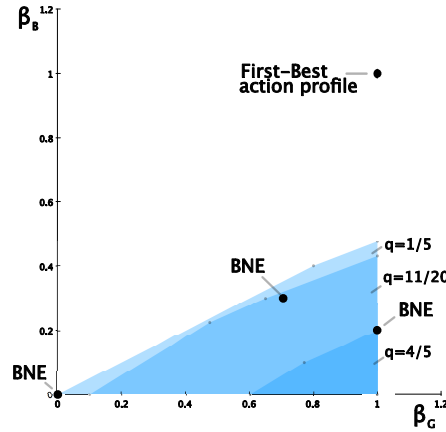


Figure 4: Set inclusion of BCE with increasing information

References

- Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- Robert J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55: 1–18, 1987.
- Robert J. Aumann and Michael Maschler. *Repeated Games with Incomplete Information*. MIT Press, 1995.
- Dirk Bergemann and Stephen Morris. Correlated equilibrium in games of incomplete information. Technical Report 1822, Cowles Foundation for Research in Economics, Yale University and Princeton University, 2011.
- Dirk Bergemann and Stephen Morris. The comparison of information structures in games: Bayes correlated equilibrium and individual sufficiency. Technical Report 1909R, Cowles Foundation for Research in Economics, Yale University, 2014a.
- Dirk Bergemann and Stephen Morris. Robust predictions in games with incomplete information. *Econometrica*, 81:1251–1308, 2013a.
- Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and comparison of information structures. Technical Report 1822R, Cowles Foundation for Research in Economics, Yale University, 2013b.
- Dirk Bergemann and Stephen Morris. Informational robustness and solution concepts. Technical Report 1973, Cowles Foundation of Research, Yale University and Princeton University, 2014b.

- Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105:921–957, 2015a.
- Dirk Bergemann, Benjamin Brooks, and Stephen Morris. Information in first price auctions. Technical Report 1926, Yale University, University of Chicago and Princeton University, 2015b.
- Dirk Bergemann, Tibor Heumann, and Stephen Morris. Information and volatility. *Journal of Economic Theory*, forthcoming, 2015c.
- David H. Blackwell. Comparison of experiments. In *Proc. Second Berkeley Symp. Math. Statist. Probab.*, pages 93–102. University of California Press, Berkeley, 1951.
- David H. Blackwell. Equivalent comparison of experiments. *Annals of Mathematics and Statistics*, 24: 265–272, 1953.
- Antonio Calvó-Armengol. The set of correlated equilibria of 2x2 games. Technical report, Universitat Autònoma de Barcelona, 2006.
- Andrew Caplin and Dean Martin. A testable theory of imperfect perception. New York University, 2013.
- Josh Cherry and Lones Smith. Strategically valuable information. Technical report, Northwestern University MEDS and University of Wisconsin, 2014.
- Eddie Dekel, Drew Fudenberg, and Stephen Morris. Interim correlated rationalizability. *Theoretical Economics*, 2:15–40, 2007.
- Jeffrey C. Ely and Marcin Peski. Hierarchies of belief and interim rationalizability. *Theoretical Economics*, 1:19–65, 2006.
- Francoise Forges. Correlated equilibrium in games with incomplete information revisited. *Theory and Decision*, 61:329–344, 2006.
- Francoise Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35:277–310, 1993.
- Olivier Gossner. Comparison of information structures. *Games and Economic Behavior*, 30:44–63, 2000.
- Olivier Gossner. Ability and knowledge. *Games and Economic Behavior*, 69:95–106, 2010.
- Olivier Gossner and Jean-Francois Mertens. The value of information in zero-sum games. Université Paris-Nanterre and CORE, Université Catholique de Louvain, 2001.

- Jack Hirshleifer. The private and social value of information and the reward to inventive activity. *American Economic Review*, 61:561–574, 1971.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101:2590–2615, 2011.
- Ehud Lehrer, Dina Rosenberg, and Eran Shmaya. Signaling and mediation in games with common interest. *Games and Economic Behavior*, 68:670–682, 2010.
- Ehud Lehrer, Dina Rosenberg, and Eran Shmaya. Garbling of signals and outcome equivalence. *Games and Economic Behavior*, 81:179–191, 2013.
- Qingmin Liu. Correlation and common priors in games with incomplete information. *Journal of Economic Theory*, 157:49–75, 2015.
- Jean-Francois Mertens and Shmuel Zamir. Formalization of bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29, 1985.
- Igal Milchtaich. Implementability of correlated and communication equilibrium outcomes in incomplete information games. part i: Definitions and results. Technical report, Bar Ilan University, 2012.
- Marcin Peski. Comparison of information structures in zero-sum games. *Games and Economic Behavior*, 62:732–735, 2008.
- Ina Taneva. Information design. Technical report, University of Texas at Austin, 2014.