

GARBLING OF SIGNALS AND OUTCOME EQUIVALENCE

EHUD LEHRER, DINAH ROSENBERG, AND ERAN SHMAYA

ABSTRACT. In a game with incomplete information players receive stochastic signals about the state of nature. The distribution of the signals given the state of nature is determined by the information structure. Different information structures may induce different equilibria.

Two information structures are *equivalent* from the perspective of a modeler, if they induce the same equilibrium outcomes. We characterize the situations in which two information structures are equivalent in terms of natural transformations, called *garblings*, from one structure to another. We study the notion of ‘being equivalent to’ in relation with three equilibrium concepts: Nash equilibrium, agent normal-form correlated equilibrium and the belief invariant Bayesian solution.

Date: May 2, 2013.

Key words and phrases. Incomplete information; Information structure; Correlated equilibrium; Garbling; Robustness.

We thank Christoph Kuzmics and Francoise Forges for very helpful discussions. Two anonymous referees and an Associate Editor of *Games and Economic Behavior* are particularly acknowledged for their comments, which greatly contributed to the improvement of the paper.

Lehrer and Shmaya acknowledge the support of BSF Grant 2010253.

The first draft of this paper had the title “Signaling and mediation in Bayesian games”. Since 2006 it has been circulated and reviewed.

1. INTRODUCTION

Players are often engaged in a strategic interaction that involves some unknown parameters. Following Harsanyi's seminal work, it became standard to model such situations as games with incomplete information. In particular, information about an unknown parameter is modeled as a stochastic signal, whose distribution depends on this parameter. Thus, beyond the action sets and the payoff functions, the description of the game must include an information structure, which determines the distribution of the signals received by the players.

We say that two information structures are outcome-equivalent with respect to an equilibrium concept, if both yield the same set of outcomes. The phrase 'outcome of the interaction' refers to a pair consisting of a state and an action-profile, which ultimately determine the payoffs. Thus, two information structures are outcome-equivalent if they induce the same set of distributions over state-action pairs. Whether the modeler chooses one information structure or another, equivalent to it, the predictions he makes about the outcomes are the same.

We examine when two information structures are equivalent with respect to Nash equilibrium and to two variations of correlated equilibrium that have been studied in a seminal paper of Forges [12]: agent normal-form correlated equilibrium and Bayesian solution. As we shall see, these solution concepts rely on the same incentive compatibility property, and differ only in the joint distributions of signals and actions they allow.

To motivate our study consider the following scenario. An economist wishes to analyze a situation in which two players jointly have sufficient information to deduce the state of the economy but neither knows anything about it individually. This situation lends itself to various ways of modeling that use distributions over the states and private signals. The economist, the modeler, selects one. This selection, together with the other components of the game (actions, utility functions, etc.) and the equilibrium concept he employs, determines the modeler's prediction about the outcome of the interaction. The question arises as to how sensitive the outcomes are to the modeling selection. Stated differently, what are the implications on the modeler's predictions of his (possibly arbitrary) choices when he explicitly delineates the tiny details of the information structure?

We view Harsanyi's stochastic structure as a modeling tool. It does not necessarily correspond to a distribution of 'real world' signals. We focus only on informational issues and assume that all other components of the games (e.g., actions, payoffs, space of states of nature, and even the identity of the players) are unambiguously known, and do not require a modeling selection.

This paper is motivated also by Blackwell's study of comparison of information structures in single-player decision problems under uncertainty [4]. Blackwell introduces two partial orders over the set of information structures. The first is operational: it compares two structures with respect to attainable state-contingent payoffs in different decision

problems. The second order is purely probabilistic. It is formulated in terms of transformations from one structure to another that involve adding noise to the signals. Blackwell proved the equivalence of these two orders. Following Marschak and Miyasawa [18], a transformation between information structures is called *garbling*. This term alludes to the fact that such a transformation typically entails a loss of information. In the single player setup Marschak and Miyasawa [18] have shown that Blackwell's operational order can also be formulated in terms of the optimal payoff a decision maker can guarantee in different problems. They show that one structure yields a higher optimal payoff than another in every decision problem if and only if the latter structure is a garbled version of the former.

The simplest example of a garbling is changing the names of the signals. Recall that as part of choosing the information structure, the modeler chooses the set of signals that the player receives. Clearly the names of the signals convey no information and changing them has no informational implication. Indeed, changing the names that the modeler chooses for the signals is an operation that any player can do on her own. While changing names is a deterministic transformation, general garblings allow also stochastic operations. But they are motivated by the same idea: these are operations that the players can perform on their own on the signals and thus rendering the modeler's choice inconsequential.

The multiplayer case is more intricate. The reason is that additional information may eliminate equilibria and, as a result, be detrimental for the players. This phenomenon has been famously demonstrated by Hirshleifer in the example of revelation of information in insurance markets [16]. While Blackwell's single-player case uses one type of garbling, the intricacy of multiplayer case gives rise to several natural types of garblings. Each type is described by a particular set of restrictions imposed over the operations performed on signals. Our results associate various solution concepts with different classes of garblings. That is, for any solution concept we specify a type of garbling that makes two information structures outcome-equivalent with respect to this concept.

Outcome-equivalence of information structures reflects the modeler's perspective: the two structures are equivalent ways of modeling. On the other hand, garblings reflects the players' perspective: the players can transform the situation represented by one structure to a situation represented by another. Thus, the results of this paper can be interpreted as establishing a similarity between the perspectives of the modeler and that of the players.

In an earlier paper [17] we studied games with common interests, where more information is always advantageous. In this paper, we introduced the definition of garbling in multi-player setup and proved an analogue of Blackwell's theorem for games with common interests. We showed that an information structure supports equilibria that induce lower payoffs than another structure if and only if it is a garbled version of the latter. In the current paper we show that in a multiplayer setup, even without the restriction

of common interests, one corollary of Blackwell’s argument can still be recovered: two information structures are outcome-equivalent if and only if they give the players the same information about the state of nature.

In a related recent paper Bergemann and Morris [5] study a solution concept for which more information can eliminate equilibria, but not create new ones. With this notion, they prove that if one structure is a garbled version of another then the latter has a smaller set of equilibrium outcomes than the former. In another close paper, Bergemann and Morris [6] characterize the set of Bayes correlated equilibria in a class of continuum player games with quadratic payoffs and normally distributed uncertainty.

In addition to games with common interests, there is another family of games that exhibit the effect that more information is the same effect on payoffs: zero-sum games. In this case more information to a player is always advantageous to the player and disadvantageous to the opponent. Pęski [20] proved an analogous result to Blackwell’s Theorem in zero-sum games. Gossner and Mertens [15] proved that two information structures are equivalent in zero-sum games if and only if they induce the same distribution over hierarchies of beliefs.

The paper is structured as follows. Section 2 presents the formal framework, the various equilibrium concepts, the classes of garblings, and the main results. Section 3 contains equivalent definitions of the equilibrium concepts that highlights the fact that they all rely on the same incentive compatibility conditions. Section 4 contains the proofs.

2. THE MODEL

2.1. Games with incomplete information. Throughout the paper, we fix a finite set $I = \{1, \dots, n\}$ of n *players* and a finite set K of *states of nature*. We divide the components of the game into “information structure,” which captures the information available to the players about the realized state of nature, and “basic game,” which captures the available actions and the payoffs. This division is by now standard. We take the terminology from Bergemann and Morris [5].

A *basic game* is given by a probability distribution p over K called the *common prior*; and, for every player $i \in I$, a finite set A_i of *actions* and a *payoff function* $u^i : K \times A \rightarrow \mathbb{R}$ where $A = \prod_{i \in I} A_i$ is the set of *action profiles*. Elements of $K \times A$ are called *plays* or *outcomes*.

An *information structure* consists of finite sets $\{S_i\}_{i \in I}$ of *signals*, and a function $\sigma : K \rightarrow \Delta(S)$ where $S = \prod_{i \in I} S_i$ is the set of *joint signals*: for every state of nature k and every $s = (s_1, \dots, s_n) \in S$, the probability that the players receive the joint signal s (i.e., that player i receives signal s_i) is given by $\sigma(k)[s]$, which we also denote by $\sigma(s|k)$.

Given the basic game and information structure, the game is played as follows: First, a state of nature $k \in K$ is randomized according to p . Then signals $s = (s_1, \dots, s_n) \in S$ is randomized according to $\sigma(k)$ and player i is informed of s_i . Then an action profile

$a = (a_1, \dots, a_n) \in A$ is chosen by the players, where a_i is the action of player i . Each player i then receive payoff $u^i(k, a)$.

2.2. Equilibrium concepts. In this section, we formalize the equilibrium concepts we are working with. Throughout the section, we fix a basic game with prior p , action sets A_i and payoff functions $u^i : K \times A \rightarrow \mathbb{R}$ where $A = \prod_i A_i$ is the set of action profiles, and fix an information structure with signal sets S_i and $\sigma : K \rightarrow \Delta(S)$ where $S = \prod_i S_i$ is the set of joint signals.

2.2.1. Nash equilibrium. A *strategy* of player i is a function $x_i : S_i \rightarrow \Delta(A_i)$: when player i receives the signal $s \in S_i$, he randomizes an action from A_i according to $x_i(s)$. We denote by $x_i(a|s) = x_i(s)[a]$ the probability that player i who receives signal s plays action a when playing strategy x_i . When the *strategy profile* $x = (x_1, \dots, x_n)$ is played, the induced distribution D_x on plays is given by

$$D_x[k, a] = p(k) \sum_{s \in S} \sigma(s|k) \prod_i x_i(a_i|s_i)$$

for every state of nature k and action profile $a = (a_1, \dots, a_n)$, where the sum ranges over all joint signals $s = (s_1, \dots, s_n)$. The expected payoff of player i under the strategy profile x is given by $u^i(x) = \sum_{k,a} u^i(k, a) D_x[k, a]$. A profile x^* is a *Bayes Nash equilibrium profile* (henceforth, Nash equilibrium) if $u^i(x_{-i}^*, x_i) \leq u^i(x^*)$ for every player i and every strategy x_i of player i , where as usual (x_{-i}^*, x_i) is the strategy profile in which player i plays x_i and each player $j \neq i$ plays x_j^* . A distribution $\pi \in \Delta(K \times A)$ is called a *Nash equilibrium* if $\pi = D_{x^*}$ for some Nash equilibrium profile x^* .

2.2.2. Correlated equilibrium. Aumann [1] introduced the notion of correlated equilibria in games with complete information. In this equilibrium players are allowed to condition their actions on different messages, produced by a *correlation device* which is independent of the primitives of the game. Thus, players' actions are typically not independent. A well-known special case is a sunspot equilibria, in which players condition their actions on a publicly known signal.

Correlated equilibrium can be generalized to incomplete information games in various ways. The definition of agent normal-form correlated equilibrium, introduced by Forges [11] and Samuelson and Zhang [21], is based on the *agent normal-form* or *Selten normal-form* of the game. In the Selten normal-form, each player is replaced by a number of agents, each of which is responsible for taking an action on behalf of the player in one information set. Thus, elements of the set S_i are viewed as agents of player i . Agent $s_i \in S_i$ becomes active in the Selten game when player i observes the signal s_i in the original incomplete-information game. In this case, agent s_i gets the expected payoff corresponding to player i in the incomplete information game. Formally, the Selten game is a complete-information where the set of players is the disjoint union $\biguplus_i S_i$ (so the number of players is $\sum_i |S_i|$). The players in S_i are called *agents of player i* . The action set of the agents of player i is A_i . Thus, the set of action profiles in the Selten

game equals the set $R = \prod_i S_i^{A_i}$ of pure strategy profiles in the game with incomplete information: Under profile $\bar{f} = (f_1, \dots, f_n)$ such that $f_i \in S_i^{A_i}$, agent $s_i \in S_i$ of player i plays $f_i(s_i) \in A_i$. The payoff of agent $s_i \in S_i$ of player i corresponding to this profile is given by

$$\sum_{k \in K} p(k) \sum_{s_{-i} \in S_{-i}} \sigma(s_i, s_{-i} | k) u^1(k, f_1(s_1), f_2(s_2), \dots, f_n(s_n)).$$

An *agent normal-form correlated equilibrium* of a game with incomplete information is a correlated equilibrium (a la Aumann) in the corresponding complete information Selten game. Note that in a correlated equilibrium of the Selten game, different agents of the same player may observe different messages from the correlation device. This fact reflects the idea that the original player interprets the message from the correlation device differently depending on his signal. See Cotter [8] for an elaboration on this point, and Forges [12] and Milchtaich [19] for discussions about implementing agent normal-form correlated equilibria using a mediator.

2.2.3. Bayesian solution. The Bayesian approach is a general approach to expressing rationality under uncertainty. We follow Forges [13]. According to this approach, the uncertainty about the state of nature and about the players' signals and actions is represented by a probability measure $\mathbb{P}$ over a set Ω of *states of the world*, and the players' information is represented by partitions $\mathfrak{A}_1, \dots, \mathfrak{A}_n$ over Ω . Thus, an epistemic model for the game is given by $(\Omega, \mathbb{P}, \mathfrak{A}_1, \dots, \mathfrak{A}_n)$, a random variable $\kappa : \Omega \rightarrow K$ representing the state of nature, and, for every player i , random variables $\varsigma_i : \Omega \rightarrow S_i$ and $\alpha_i : \Omega \rightarrow A_i$, representing player i 's signal and action. In addition, we assume that ς_i and α_i are $\mathfrak{A}_i$ -measurable.

The random variable $\varsigma = (\varsigma_1, \dots, \varsigma_n)$ represent any information that the players may possess about the parameters of the game (i.e., the state and the other player's signal). The partitions $\mathfrak{A}_1, \dots, \mathfrak{A}_n$ represent the entire knowledge of the players. This includes, among other things, any game-relevant information such as the player's signal and action and possibly other information which is independent of the game and which may serve as a correlation device between the players. The following definition captures the idea that the partitions do not give the players more information than what is dictated by σ .

Definition 2.1. An epistemic model is *belief invariant* if it satisfies the following conditions:

- (1) Consistency: The joint distribution of κ, ς is given by the common prior p on the state of nature and the information structure σ :

$$\mathbb{P}(\kappa = k, \varsigma = s) = p[k] \cdot \sigma(s|k)$$

for every $k \in K$ and $s \in S$.

- (2) The joint signal of the players completely describes their joint information about the state of nature (i.e., ς is a sufficient statistic for κ given $\vee_i \mathfrak{A}_i$):

$$\mathbb{P}(\kappa = k | \vee_i \mathfrak{A}_i) = \mathbb{P}(\kappa = k | \varsigma).$$

for every $k \in K$.¹

- (3) The signal ς_i of player i completely describes his information about the state of nature and the other players' signals (i.e., ς_i is a sufficient statistic for the pair (κ, ς_{-i}) given $\mathfrak{A}_i$):

$$\mathbb{P}(\kappa = k, \varsigma_{-i} = s_{-i} | \mathfrak{A}_i) = \mathbb{P}(\kappa = k, \varsigma_{-i} = s_{-i} | \varsigma_i),$$

for every $k \in K$ and $s_{-i} \in S_{-i}$.²

The following definition captures the idea of rationality of the players.

Definition 2.2. An epistemic model is *incentive compatible* if no deviation by player i from playing α_i is profitable:

$$\mathbb{E} u^i(\kappa, a, \alpha_{-i} | \mathfrak{A}_i) \leq \mathbb{E} u^i(\kappa, \alpha | \mathfrak{A}_i)$$

for every action $a \in A_i$, where $\mathbb{E}(\cdot | \mathfrak{A}_i)$ is the conditional expectation over $(\Omega, \mathbb{P})$ given $\mathfrak{A}_i$.

Definition 2.3. A distribution $\pi \in \Delta(K \times A)$ over plays is a *belief invariant Bayesian solution* if π is the joint distribution of (κ, α) in some belief invariant, incentive compatible epistemic model.

Remark 2.4. If one dispenses with condition (3) in the definition of belief invariant model, one gets what Forges [13] calls a *Bayesian solution*. If one also dispenses with Condition (2), one gets what Bergemann and Morris [5] call a *correlated Bayesian equilibrium* (See also Remark 2.11 below).

2.3. Stochastic maps. The following notation will be used extensively throughout the rest of the paper. Let X, X' be two finite sets. A *stochastic map* from X to X' is a function from X to $\Delta(X')$. We denote the set of such maps by $\mathcal{S}(X, X')$. If $\varphi \in \mathcal{S}(X, X')$ we sometimes denote $\varphi(x'|x)$ for $\varphi(x)[x']$. Note that if K is the set of states and S_i are sets of players' signals, then $\mathcal{S}(K, \prod_i S_i)$ is the set of information structures with signal sets S_i .

Every stochastic map from X to X' extends to a linear function from $\Delta(X)$ to $\Delta(X')$. We do not distinguish between a stochastic map from X to X' and its linear extension.

Let X, X', Y, Y' be finite sets and let $\varphi \in \mathcal{S}(X, X')$ and $\psi \in \mathcal{S}(Y, Y')$. Their *product*, denoted by $\varphi \otimes \psi$, is a stochastic map from $X \times Y$ to $X' \times Y'$, defined by

$$(1) \quad \varphi \otimes \psi : (x, y) \mapsto \varphi(x) \otimes \psi(y)$$

¹Equivalently, κ and $\vee_i \mathfrak{A}_i$ are conditionally independent given ς .

²Equivalently, (κ, ς_{-i}) and $\mathfrak{A}_i$ are conditionally independent given ς_i .

for every $x, y \in X \times Y$, where $\varphi(x) \otimes \psi(y)$ is the product distribution of $\varphi(x)$ and $\psi(y)$.

Let X, X', X'' be finite sets and let $\varphi \in \mathcal{S}(X, X')$ and $\psi \in \mathcal{S}(X', X'')$. The composition of φ and ψ induces a stochastic map $\psi \circ \varphi \in \mathcal{S}(X, X'')$ such that $\psi \circ \varphi(x''|x) = \sum_{x' \in X'} \psi(x''|x')\varphi(x'|x)$ for every $x \in X$ and $x'' \in X''$. The composition notation is consistent with the convention not to distinguish between ψ and its linear extension from $\Delta(X')$ to $\Delta(X'')$.

2.4. Garbling of information. By *garbling* of information, we mean applying a stochastic map on a signal to create a new one. For every player i fix two finite sets of signals, S_i and S'_i , and let $S = \prod_i S_i$ and $S' = \prod_i S'_i$ be the corresponding sets of joint signals. We say that an information structure $\sigma' : K \rightarrow \Delta(S')$ is a *garbled version of the information structure* $\sigma : K \rightarrow \Delta(S)$ if $\sigma' = q \circ \sigma$ for some stochastic map $q \in \mathcal{S}(S, S')$. We use the term ‘garbling’ both for the process and for the map q itself. In the case of one-player decision problems, Blackwell [4] defined one structure σ to be sufficient to another structure σ' , if σ' is a garbled version of σ . In this case one may simulate a σ' -signal by garbling (or *manipulating*) a σ -signal. In a multiplayer environment, players may be unable or unwilling to join forces to produce new signals from old. This means that players will not be able to employ every garbling. The following definition will be used to restrict the garblings available to the players.

Definition 2.5. Let X_i, Y_i for every $i \in I$ and let $\varphi \in \mathcal{S}(X, Y)$ be a stochastic map, where $X = \prod_i X_i$ and $Y = \prod_i Y_i$.

- (1) φ is *independent* if $\varphi = \varphi_1 \otimes \cdots \otimes \varphi_n$ for some $\varphi_i \in \mathcal{S}(X_i, X'_i)$.
- (2) φ is *coordinated* if it is a convex combination of independent stochastic maps, i.e., if $\varphi = \sum_{\ell=1}^L \lambda^\ell \prod_i \varphi_i^\ell$ where $\lambda^\ell \geq 0$, $\sum_{\ell=1}^L \lambda^\ell = 1$, and $\varphi_i^\ell \in \mathcal{S}(X_i, X'_i)$ for every $\ell = 1, \dots, L$ and $i \in I$.
- (3) φ is *noncommunicating* if there exist stochastic maps $\varphi_i \in \mathcal{S}(X_i, X'_i)$ such that

$$\text{Marg}_{X'_i}(\varphi(x)) = \varphi_i(x_i)$$

for every player i and every $x = (x_1, \dots, x_n) \in X$, where $\text{Marg}_{X'_i}(\xi)$ is the marginal of ξ over X'_i for every $\xi \in \Delta(X')$.

Every independent stochastic map is coordinated and every coordinated stochastic map is non-communicating (see also our earlier paper [17] for more on these types of stochastic maps).

We now discuss the implication of these conditions on the implementation of garblings. Let S_i, S'_i be finite sets of player i 's signals, and let $S = \prod_i S_i$, $S' = \prod_i S'_i$ be the corresponding sets of joint signals.

Assume that $q \in \mathcal{S}(S, S')$ is an independent garbling, so that $q = \prod_i q_i$ for some $q_i \in \mathcal{S}(S_i, S'_i)$. More explicitly, the condition means that $q(s'|s) = \prod_i q_i(s'_i|s_i)$ for every $s = (s_1, \dots, s_n) \in S$ and for every $s' = (s'_1, \dots, s'_n) \in S'$. Thus, independent garbling

can be implemented without communication: player i manipulates his signal on his own using q_i , independently of other players.

A coordinated garbling can be implemented using a public message that is independent of the players' signals. All players observe message $\ell \in \{1, \dots, L\}$ with probability λ^ℓ , and then jointly implement an independent garbling (which may depend on ℓ).

In order to explain the meaning of noncommunicating garbling, let $q \in \mathcal{S}(S, S')$. Let s be a random joint signal generated according to some distribution $\xi \in \Delta(S)$, and let s' be its random garbling according to q . If q is noncommunicating, then player i 's posterior distribution of s_{-i} given s_i, s'_i , coincides with his posterior distribution of s_{-i} given s_i alone. In other words, if the garbling is performed by a mediator who knows the joint signal s and uses it to produce the new joint signals s' , then, while the mediator uses the information that all players have, she does not provide a player with more information about his opponents' signals than what he had previously possessed. The term "non-communicating" reflects the idea that the garbling mechanism does not communicate information between the players.

Remark 2.6. Gossner [14] defined another property of stochastic maps, which he calls faithfulness. Faithfulness is different in nature from the properties considered in this paper because it is not a property of the stochastic map, but rather a property of the map combined with the information structure on which it acts: the same map might be faithful when acting on one structure and not faithful when acting on another.

2.5. Example. Let $K = \{0, 1, 2, 3\}$. Consider a situation in which each of two players can infer from his signal nothing about the state of nature, but the joint signal of the players fully reveals the state of Nature.

One way to model this situation is using an information structure σ with signal spaces $S_1 = S_2 = \{0, 1, 2, 3\}$ such that $\sigma(k)$ is the uniform distribution over the four pairs (s_1, s_2) and $k = s_1 + s_2$, where $+$ is addition modulo 4. Then, regardless of the state of nature, the distribution of each player's signal is uniform on $\{0, 1, 2, 3\}$, so that the player does not learn anything about the state of nature from his signal. However, an observer who knows both signals s_1, s_2 can deduce k as $k = s_1 + s_2$.

Another possible modeling of the same situation is using an information structure σ' with signal sets $S'_1 = S'_2 = \{0, 1, 2, 3\}$ and such that $\sigma'(k)$ is the uniform distribution over the four pairs (s_1, s_2) and $k = s_1 \oplus s_2$, where $\oplus$ is the group operation on $\{0, 1, 2, 3\}$ with the following multiplication table

$\oplus$	0	1	2	3
0	0	1	2	3
1	1	0	3	2
2	2	3	0	1
3	3	2	1	0

(This operation corresponds to the bitwise addition modulo 2 of numbers in their binary representation.)

The similarity between σ and σ' in terms of the information available to the players is reflected by the fact that there are noncommunicating garblings that transform σ to σ' and vice versa. Indeed, let $q \in \mathcal{S}(S, S')$ be such that $q(s)$ is the uniform distribution over the pairs (s'_1, s'_2) that satisfy $s'_1 \oplus s'_2 = s_1 + s_2$. Then $\sigma' = q \circ \sigma$. The garbling q can be interpreted as a mediator who, after observing $s = (s_1, s_2)$ and deducing the state k , first produces new signals $s' = (s'_1, s'_2)$ according to σ' given k , and then sends s'_i to player i . Note that player 1 does not gain any new information in the process about player 2's signal s_2 since his new signal s'_1 is uniformly distributed regardless of s_2 . Thus, q is noncommunicating. Similarly, if $q' \in \mathcal{S}(S', S)$ is such that $q'(s'_1, s'_2)$ is the uniform distribution over the pairs $s = (s_1, s_2)$ that satisfy $s_1 + s_2 = s'_1 \oplus s'_2$, then q' is noncommunicating and $\sigma = q' \circ \sigma'$.

One may conclude from this discussion that for all practical purposes, σ and σ' represent the same situation. However, the validity of this statement depends on what is meant by “practical purposes.” Consider, for instance, a game with a uniform prior over K , action sets $A_1 = A_2 = \{0, 1, 2, 3\}$, and the payoff functions

$$u^1(k, a_1, a_2) = u^2(k, a_1, a_2) = \begin{cases} 1, & \text{if } k = a_1 + a_2 \pmod{4} \\ 0, & \text{otherwise.} \end{cases},$$

The players would like to coordinate their actions, trying to produce actions a_1, a_2 such that $k = a_1 + a_2$.

Under σ , if every player plays his signal, the payoff is 1. This is a Nash equilibrium. However, under σ' , no Nash or correlated equilibrium yields a payoff of 1. To see this, it is sufficient to note that no pair of pure strategies for the players give payoff 1. This fact, which can be checked manually by going over all pairs of strategies, boils down to the fact that the group operations $+$ and $\oplus$ are not isomorphic. Thus, σ and σ' are totally different as far as Nash equilibria and correlated equilibria are concerned. We return to this example in Section 2.6.

Remark 2.7. The example is related to the examples of Dekel, Fudenberg, and Morris [9] and Ely and Peşki [10] with the twist that in our case the Δ -hierarchies, as defined by Ely and Peşki, are the same in both structures. Indeed, there is a close relationship between Δ -hierarchies and the Bayesian solution, but we do not pursue it in this paper.

2.6. Equivalent information structures. Fix finite sets S_i, S'_i of players' signals with the corresponding sets $S = \prod_i S_i$ and $S' = \prod_i S'_i$ of joint signals, and let $\sigma : K \rightarrow \Delta(S)$ and $\sigma' : K \rightarrow \Delta(S')$ be two information structures. Our purpose is to compare the equilibria induced by these structures. The following terminology will be useful: given a basic game we call the game induced by σ and the basic game the σ -game, and we call the game induced by σ' and the basic game the σ' -game.

The following definition captures the idea that two structures are equivalent if they generate the same outcomes.

Definition 2.8. The information structures σ and σ' are *equivalent* with respect to some equilibrium concept if, for every basic game, the sets of equilibria of the σ -game and the σ' -game are identical.

The following definition formalizes the equivalence of information structures in terms of garblings. It is the analogue of Blackwell's idea of simulating one structure from another.

Definition 2.9. The information structures σ and σ' are *equivalent* with respect to some class of garblings iff there exist garblings $q \in \mathcal{S}(S, S')$ and $q' \in \mathcal{S}(S', S)$ from this class such that $\sigma' = q \circ \sigma$ and $\sigma = q' \circ \sigma'$.

We are now in a position to state our main theorem:

Theorem 2.10. *Let σ and σ' be two information structures.*

- (1) *σ and σ' are equivalent with respect to Nash equilibria if and only if they are equivalent with respect to independent garblings.*
- (2) *σ and σ' are equivalent with respect to agent normal-form correlated equilibria if and only if they are equivalent with respect to coordinated garblings.*
- (3) *σ and σ' are equivalent with respect to belief invariant Bayesian solutions if and only if they are equivalent with respect to noncommunicating garblings.*

Consider again the example in Section 2.5. As we have seen, the structures σ and σ' are equivalent with respect to noncommunicating garblings. By Theorem 2.10 it follows that they are equivalent with respect to Bayesian solutions. However, as we have seen, they are not equivalent with respect to Nash and agent normal-form correlated equilibria. Therefore, by Theorem 2.10 they are not equivalent with respect to coordinated garblings. By the way, these observations show in particular that the noncommunicating garblings that transform σ to σ' and back are not coordinated.

Remark 2.11. Bergemann and Morris [5] define a more general form of garbling: a map in $\mathcal{S}(K \times S, S')$, implying that that distribution of the new signals (i.e., those in S') depends not only on the original signals (i.e., those in S) but also on the state of nature. They show that an information structure σ' is a garbled version of σ with respect to a noncommunicating garbling of this form if and only if it generates a larger set of Bayesian correlated equilibrium outcomes (See Remark 2.4). One corollary of their result is similar in spirit to Theorem 2.10: two information structures are equivalent with respect to correlated Bayesian equilibria if and only if each is a garbled version of the other with respect to a noncommunicating garblings of this form.

3. GLOBAL STRATEGIES

In this section we give equivalent formulations of the equilibrium concepts in Section 2.2 using a concept we call “global strategy”, that corresponds to the mixed distribution over action profiles as a function of the joint signals. This formulation shows that the solution concepts we study rely on the same incentive compatibility condition, but they differ in the joint distributions of signals and actions they allow.

Fix a basic game with prior p , action sets A_i and payoff functions $u^i : K \times A \rightarrow \mathbb{R}$ where $A = \prod_i A_i$ is the set of action profiles, and fix an information structure with signal sets S_i and $\sigma : K \rightarrow \Delta(S)$ where $S = \prod_i S_i$ is the set of joint signals.

A *global strategy* τ is a member of $\mathcal{S}(S, A)$: if s is the joint signals of the players, then $\tau(a|s)$ is the probability that the action profile a is played under τ .

Every global strategy $\tau \in \mathcal{S}(S, A)$ induces a probability distribution $D(\sigma, \tau)$ over plays given by

$$D(\sigma, \tau) = \sum_k p(k) \delta_k \otimes \tau(\sigma(k)).$$

Here, δ_k is an atomic measure of mass 1 at k . Equivalently,

$$D(\sigma, \tau)[k, a] = p(k) \sum_{s \in S} \sigma(s|k) \tau(a|s).$$

Let $U^i(\sigma, \tau) = \sum_{k,a} u^i(k, a) D(\sigma, \tau)[k, a]$ be the expected payoff of player i if the global strategy τ is played.

The new formulation expresses each equilibrium concept using two separate condition on the global strategy. The first condition, incentive compatibility, is the same for all equilibrium concepts; the second condition is an restriction on the global strategy which characterizes each equilibrium concept and is unrelated to the game itself.

Definition 3.1 (incentive compatibility). A global strategy τ is *incentive-compatible* if $U^1(\sigma; \tau_v) \leq U^1(\sigma; \tau)$ for every player i and for every *deviation function* $v : S_i \times A_i \rightarrow A_i$ of player i , where τ_v is the global strategy that is given by

$$\tau_v(a|s) = \sum_{b \in A_i | v_i(s, b) = a_i} \tau(b, a_{-i}|s)$$

for every action profile $a \in A$ and joint signal $s \in S$.

The following proposition gives the equivalent formulation of our solution concepts using the notion of global strategy. The proof essentially appears in our earlier paper [17, Lemma 5.1, Lemma 5.2], but is also sketched below for completeness.

Proposition 3.2. *Let π be a distribution over plays. Then*

- (1) *π is a Nash equilibrium if and only if $\pi = D(\sigma, \tau)$ for some incentive-compatible, independent global strategy $\tau \in \mathcal{S}(S, A)$.*
- (2) *π is an agent normal-form correlated equilibrium if and only if $\pi = D(\sigma, \tau)$ for some incentive-compatible, coordinated global strategy $\tau \in \mathcal{S}(S, A)$.*

- (3) π is a belief invariant Bayesian solution if and only if $\pi = D(\sigma, \tau)$ for some incentive-compatible, noncommunicating global strategy $\tau \in \mathcal{S}(S, A)$.

Sketch of proof. (1) The assertion follows from the fact that a strategy profile $x = (x_1, \dots, x_n)$ is Nash equilibrium if and only if the global strategy $\tau = x_1 \otimes \dots \otimes x_n$ is incentive compatible.

- (2) Let $R = \prod_i A_i^{S_i}$ be the set of action profiles in the Selten game. Recall that elements $\bar{f} = (f_1, \dots, f_n)$ of R can be viewed as a pure strategy profiles in the game with incomplete information. The map ρ from R to $\mathcal{S}(S, A)$ given by

$$\rho(f_1, \dots, f_n) : (s_1, \dots, s_n) \mapsto \delta_{(f_1(s_1), \dots, f_n(s_n))}$$

has a linear extension to a map $\tilde{\rho}$ from $\Delta(R)$ onto the set of coordinated global strategies, such that $\tilde{\rho}(\lambda) = \sum_{\bar{f}} \lambda[\bar{f}] \rho(\bar{f})$ for every $\lambda \in \Delta(R)$ ³. The assertion follows from the fact that an element λ of $\Delta(R)$ is a correlated equilibrium in the Selten game if and only if the global strategy $\tilde{\rho}(\lambda)$ is incentive-compatible.

- (3) Every global strategy $\tau \in \mathcal{S}(S, A)$ gives rise to an epistemic model with $\Omega = K \times S \times A$, where $\kappa : \Omega \rightarrow K, \varsigma : \Omega \rightarrow S, \alpha : \Omega \rightarrow A$ are the coordinate projections, $\mathfrak{A}_i$ is the partition generated by ς_i, α_i and $\mathbb{P}$ is given by $\mathbb{P}[k, s, a] = p[k] \cdot \sigma(s|k) \cdot \tau(a|s)$. Then

- If τ is non-communicating then the epistemic model is belief invariant.
- If τ is incentive compatible then the epistemic model is incentive compatible.
- The joint distribution of (κ, α) is $D(\sigma, \tau)$.

It follows that if $\pi = D(\sigma, \tau)$ for some incentive compatible, non-communicating global strategy then π is belief invariant Bayesian solution.

Conversely, start with an epistemic model which is belief invariant and incentive compatible. From Condition 1 of Definition 2.1 it follows that $\mathbb{P}(\kappa = k, \varsigma = s) = p[k] \cdot \sigma(s|k)$. From Condition 2 of Definition 2.1 it follows that $\mathbb{P}(\kappa = k, \varsigma = s, \alpha = a) = p[k] \cdot \sigma(s|k) \cdot \tau(a|s)$ where $\tau \in \mathcal{S}(S, A)$ is given by $\tau(a|s) = \mathbb{P}(\alpha = a | \varsigma = s)$. From Condition 3 of Definition 2.1 it follows that τ can be chosen to be non-communicating⁴. From the fact that the epistemic model is incentive compatible it follows that τ is incentive compatible. Finally, the joint distribution of κ, α is $D(\sigma, \tau)$. It follows that τ is an incentive compatible, noncommunicating global strategy such that $\pi = D(\sigma, \tau)$. □

Remark 3.3 (Strategic normal-form correlated equilibrium). The assertion of Proposition 3.2 about agent normal-form correlated equilibrium and its proof highlight the

³An anonymous referee has pointed out that the argument can be seen as a multiplayer analogue of Kuhn's Theorem: A mixture over pure strategy profiles is equivalent to a coordinated global strategy. In a single player setup, what we call global strategy is what is usually called behavioral strategy.

⁴Can be chosen' and not 'is' because $\tau(s)$ can be chosen arbitrarily if $\mathbb{P}(\varsigma = s) = 0$.

distinction with strategic normal-form correlated equilibrium. The latter solution concept can be implemented by a mediator who randomizes a pure strategy profile $\bar{f} = (f_1, \dots, f_n)$ according to some distribution λ and sends the strategy f_i to player i . If players follow the mediator recommendation then the distribution induced over plays is given by $\pi = D(\sigma, \tau)$ where $\tau = \tilde{\rho}(\lambda)$ is a coordinated global strategy. Thus, the restriction on the distribution of players actions given signals in the case of strategic normal-form correlated equilibrium is the same as in agent normal-form correlated equilibrium. However, The incentive compatibility condition of strategic normal-form correlated equilibrium is stronger than that of Definition 3.1. See Forges [12, Section 4] for a discussion of strategic and agent normal-form correlated equilibria.

4. THE PROOF OF THEOREM 2.10

Section 4.1 includes some auxiliary results that will be used in the proof. In Section 4.2 we prove the “if” direction. In Section 4.3 we prove the “only if” direction.

4.1. Preliminaries.

4.1.1. *Closure properties.* Let X, X' be finite set. The set $\mathcal{S}(X, X')$ of stochastic maps from X to X' is naturally embedded in the finite dimensional linear space $\mathbb{R}^{X \times X'}$. The following lemma asserts the closure of the sets of noncommunicating stochastic maps, coordinated stochastic maps and independent maps under convex combinations and limits.

Lemma 4.1. *Let X_i, X'_i be finite set, and let $X = \prod_i X_i$ and $X' = \prod_i X'_i$. Then the set of noncommunicating stochastic maps from X to X' and the set of coordinated stochastic maps from X to X' are polytops, and, a fortiori, compact and convex. The set of independent stochastic maps from X to X' is compact.*

The following lemma asserts the closure of sets of independent, coordinated, and noncommunicating stochastic maps under map composition. Its assertion about independent maps follows from the fact that the map composition $\circ$ and the product $\otimes$ commute. Its assertion about coordinated maps follows from the same observation and the bilinearity of map composition. Its assertion about noncommunicating maps follows from the definition.

Lemma 4.2. *Let X_i, X'_i, X''_i be finite sets, let $X = \prod_i X_i, X' = \prod_i X'_i$ and $X'' = \prod_i X''_i$, and let $\varphi \in \mathcal{S}(X, X')$ and $\psi \in \mathcal{S}(X', X'')$.*

- (1) *If φ and ψ are independent, then so is $\psi \circ \varphi$.*
- (2) *If φ and ψ are coordinated, then so is $\psi \circ \varphi$.*
- (3) *If φ and ψ are noncommunicating, then so is $\psi \circ \varphi$.*

4.1.2. *Implications of mutual garbling.* Let Z, Z' be finite sets, and let $\eta \in \mathcal{S}(Z, Z')$, $\eta' \in \mathcal{S}(Z', Z)$ and $p, p' \in \Delta(Z)$ be such that $p' = \eta(p)$ and $p = \eta'(p')$.

Consider the probability distribution $\pi \in \Delta(Z \times Z')$ that is given by $\pi[z, z'] = p[z] \cdot \eta(z'|z)$, which is the distribution generated when an $z \in Z$ is randomized according to p and then an element $z' \in Z'$ is produced via $\eta(z)$. The marginal of this distribution over Z is p and its marginal over Z' is $p' = \eta(p)$. Similarly, consider the probability distribution $\pi' \in \Delta(Z \times Z')$ that is given by $\pi'[z, z'] = p'[z']\eta'(z|z')$, which is the distribution generated when signal $z' \in Z'$ is randomized according to p' and then an element $z \in Z$ is randomized according to $\eta'(z')$. Again, the marginals of π' over Z and Z' are, respectively, $p = \eta'(p')$ and p' . In general, it need not be the case that $\pi = \pi'$, which means that the two points of view (randomizing first z and then z' , or randomizing first z' and then z) are not necessarily compatible with each other. In this section we prove that some weak version of consistency can still be recovered.

A map $\xi \in \mathcal{S}(Z, Z)$ is called *idempotent* if $\xi \circ \xi = \xi$.

Lemma 4.3. *Let Z and Z' be two finite sets, let $\eta \in \mathcal{S}(Z, Z')$, $\eta' \in \mathcal{S}(Z', Z)$ be a pair of stochastic maps, and let $p \in \Delta(Z)$ and $p' \in \Delta(Z')$ be such that $\eta(p) = p'$ and $\eta'(p') = p$. If $\xi = \eta' \circ \eta$ is idempotent, then $\eta(p) = p'$, $\eta'(p') = p$ and*

$$p[z] \cdot \eta \circ \xi(z'|z) = p'[z'] \cdot \xi \circ \eta'(z|z')$$

for every $z \in Z$ and $z' \in Z'$.

Proof. The assertion of the lemma can be written equivalently as

$$(2) \quad \sum_z p[z] \cdot (\delta_z \otimes \eta(\xi(z))) = \sum_{z'} p'[z'] \cdot (\xi(\eta'(z')) \otimes \delta_{z'})$$

We view ξ as a Markov transition over Z and use some classical properties of Markov chains. Let $Z_1, \dots, Z_k$ be the partition of Z to ergodic sets of ξ ; let p_i be their corresponding invariant distributions (such that the support of p_i is Z_i); and let $p'_i = \eta(p_i)$.

Since $\xi(p) = \eta'(\eta(p)) = \eta'(p') = p$, it follows that p is ξ -invariant and therefore,

$$(3) \quad p = \sum_i \lambda_i p_i,$$

for some $\lambda_i \geq 0$ and $\sum_i \lambda_i = 1$. From the linearity of η , it follows that:

$$(4) \quad p' = \sum_i \lambda_i p'_i.$$

Since ξ is idempotent, it is reduced to its invariant distribution on each ergodic class. Thus,

$$(5) \quad \text{if } z \in Z_i, \text{ then } \xi(z) = p_i$$

and therefore $\eta \circ \xi(z) = p'_i$ if $z \in Z_i$. From (3), it follows that the left-hand side of (2) equals $\sum_i \lambda_i p_i \otimes p'_i$.

Since the sets Z_i are the ergodic sets of ξ , it follows that if $z \in Z_i$ and $z' \in Z'$ such that $\eta(z'|z) > 0$, then the support of $\eta'(z')$ is contained in Z_i . Indeed, if $\eta'(z''|z') > 0$ for some $z'' \in Z$, then $\xi(z''|z) \geq \eta(z'|z) \cdot \eta'(z''|z') > 0$ and since Z_i is ergodic, it follows that $z'' \in Z_i$. This observation implies that the supports $Z'_1, \dots, Z'_n$ of $p'_1, \dots, p'_n$ are disjoint and that the support of $\eta'(z')$ is contained in Z_i whenever $z' \in Z'_i$, and therefore by (5) $\xi \circ \eta'(z') = p_i$ whenever $z' \in Z'_i$. By (4), the right-hand side of (2) also equals $\sum_i \lambda_i p_i \otimes p'_i$. $\square$

Recall from the theory of markov chain that for every finite set Z and every stochastic map $\xi \in \mathcal{S}(Z, Z)$ the limit

$$(6) \quad \xi^\infty = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \xi^k$$

exists and is idempotent. From Lemma 4.3 we get the following Corollary.

Corollary 4.4. *Let Z, Z' be finite sets and let $\theta : Z \rightarrow Z'$ and $\theta' : Z' \rightarrow Z$ and $p \in \Delta(Z)$ such and $p' \in \Delta(Z')$ such that $\theta(p) = p'$ and $\theta'(p') = p$. Let $\xi = \eta' \circ \eta$ and let ξ^∞ be given by (6). Let $\eta = \theta \circ \xi^\infty$ and $\eta' = \xi^\infty \circ \theta'$. Then $\eta(p) = p'$, $\eta'(p') = p$ and*

$$p[z] \cdot \eta(z'|z) = p'[z'] \cdot \eta'(z|z')$$

for every $z \in Z$ and $z' \in Z'$.

Proof. Note first that $\xi(p) = \theta'(\theta(p)) = \theta'(p') = p$ and therefore $\eta(p) = \theta(\xi^\infty(p)) = \theta(p) = p'$ and similarly that $\eta'(p') = p$. The last assertion follows from Lemma 4.3 and the facts that $\eta' \circ \eta = \xi^\infty \circ \xi \circ \xi^\infty = \xi^\infty$ is idempotent and $\eta \circ \xi^\infty = \eta$ and $\xi^\infty \circ \eta' = \eta'$. $\square$

4.2. From strategic equivalence to information garbling. This direction is an immediate consequence of Proposition 3.2. Say that an information structure σ is *larger* than an information structure σ' with respect to an equilibrium concept if, for every basic game, every equilibrium in the σ' -game is also an equilibrium in the σ -game. Note that σ and σ' are equivalent with respect to an equilibrium concept iff σ is larger than σ' and σ' is larger than σ with respect to that concept.

Assume that σ is larger than σ' with respect to Nash Equilibrium. Consider the basic game in which the prior p over K is uniform, the action set of player i is S'_i , and the payoff function is constant. Then every global strategy, in particular the identity map $\text{id} : S' \rightarrow S'$ (which is also independent) is incentive compatible. Therefore, by Proposition 3.2 $\pi = D(\sigma', \text{id}) \in \Delta(K \times S')$ is a Nash equilibrium in the game over σ' . Since σ is larger than σ' it follows by the same theorem that $\pi = D(\sigma, q)$ for some independent garbling $q \in \mathcal{S}(S, S')$. Therefore

$$\sum_k p(k) \delta_k \otimes q(\sigma(k)) = D(\sigma, q) = D(\sigma', \text{id}) = \sum_k p(k) \delta_k \otimes \sigma'(k).$$

Since $p(k) > 0$ for every k it follows that $\sigma'(k) = q(\sigma(k))$ for every k , so that $\sigma' = q \circ \sigma$. Thus, we proved that if σ is larger than σ' with respect to Nash Equilibrium then $\sigma' = q \circ \sigma$

for some independent garbling q . This implies the first assertion in Theorem 2.10. The same proof applies *mutatis mutandis* to the other assertions.

4.3. From garblings to strategic equivalence. We first prove the third assertion in Theorem 2.10. Let $q : S \rightarrow \Delta(S')$ and $q' : S' \rightarrow \Delta(S)$ be noncommunicating garblings such that $q(\sigma(k)) = \sigma'(k)$ and $q'(\sigma'(k)) = \sigma(k)$ for every $k \in K$. We must show that σ, σ' are equivalent with respect to belief invariant Bayesian solutions.

By Corollary 4.4 we can assume that

$$(7) \quad \sigma(s|k) \cdot q(s'|s) = \sigma'(s'|k) \cdot q'(s|s')$$

for every k, s, s' . Otherwise we replace q with $q \circ \xi^\infty$ and q' with $\xi^\infty \circ q'$ and the new garblings are still noncommunicating since ξ^∞ is noncommunicating by Lemma 4.2 and Lemma 4.1.

Fix a basic game and let $\pi \in \Delta(K \times A)$ be a belief invariant Bayesian equilibrium in the σ' -game. By Proposition 3.2, there exists a noncommunicating map $\tau' : S' \rightarrow A$ such that τ' is incentive-compatible in the σ' -game. Let $\tau : S \rightarrow A$ be given by

$$(8) \quad \tau = \tau' \circ q.$$

Then, τ is a noncommunicating garbling (as a composition of noncommunicating garblings), and $\tau(\sigma(k)) = \tau'(\sigma'(k))$ for every $k \in K$. In particular, it follows that

$$(9) \quad D(\sigma, \tau) = D(\sigma', \tau'),$$

i.e., the distributions over plays induced by τ in the σ -game equal those induced by τ' in the σ' -game. In light of Proposition 3.2, what remains to be shown is that τ is incentive-compatible in the σ -game. Indeed, let $v : S_i \times A_i \rightarrow A_i$ be a deviation for player 1. Under this deviation, the expected payoff of player 1 is given by:

$$(10) \quad \begin{aligned} & \sum_k p(k) \sum_{s,a} \sigma(s|k) \tau(a|s) u^1(k, v(s_i, a_i), a_{-i}) = \\ & \sum_k p(k) \sum_{s,s',a} \sigma(s|k) q(s'|s) \tau'(a|s') u^1(k, v(s_i, a_i), a_{-i}) = \\ & \sum_k p(k) \sum_{s,s',a} \sigma'(s'|k) q'(s|s') \tau'(a|s') u^1(k, v(s_i, a_i), a_{-i}) = \\ & \sum_k p(k) \sum_{s_i, s'_i, s'_{-i}, a} \sigma'(s'|k) \tau'(a|s') q'(s_i|s'_i) u^1(k, v(s_i, a_i), a_{-i}) \leq \\ & U^1(\sigma', \tau') = U^1(\sigma, \tau). \end{aligned}$$

The first equality follows from (8). The second equality follows from (7). In the third equation we use the fact that q' is noncommunicating, which implies that $\sum_{s_{-i}} q'(s_i, s_{i-1}|s'_i, s'_{i-1})$ is independent of s'_{-i} . We denote this sum by $q'(s_i|s'_i)$. The inequality follows from the

fact that τ' is incentive compatible in the σ' -game: player 1 has no incentive to randomize s_i according to $q'(s_i|s'_i)$ and then deviate to $v(s_i, a)$. The last equality follows from (9). This completes the proof of the third assertion of Theorem 2.10.

For the second assertion of Theorem 2.10, note that the same proof applies; the only difference is that all the garblings involved are coordinated. Note that in the last equality of Equation (10) we used the fact that the garbling $\xi \circ q'$ is noncommunicating. This remains valid here too, since independent and coordinated garblings are automatically noncommunicating.

For the case of independent garblings, the same proof applies with the twist that in order to achieve (7), we replace $q = q_1 \otimes \cdots \otimes q_n$ and $q' = q'_1 \otimes \cdots \otimes q'_n$ with $(q_1 \circ \xi_1^\infty) \otimes \cdots \otimes (q_n \circ \xi_n^\infty)$ and $(\xi_1^\infty \otimes q'_1) \otimes \cdots \otimes (\xi_n^\infty \otimes q'_n)$ where $\xi_i = q'_i \circ q_i$. Then the rest of the proof is the same except that all the garblings involved are independent.

5. THE PHYSICS OF GARBLINGS

The properties of the garblings introduced in this paper have a natural physical interpretation. Up to now, we have used the term ‘implementation’ in an abstract sense, without being explicit about how the garbling is implemented in the physical world. Consider, for instance the issue of implementing a garbling mechanism between two players as a bipartite physical system with inputs $s = (s_1, s_2)$ and outputs $s' = (s'_1, s'_2)$. By ‘bipartite’ we mean that the system is composed of two separate sub-systems. Each sub-system has an input and an output. Two players prepare the system together and then separate in space (say, one player goes to planet Earth and the second player to planet Mars), each holding one sub-system. The players insert the inputs s_1, s_2 to their respective sub-system and receive outputs s'_1, s'_2 . The initial physical state, before inputs are received, of the bipartite system might be random.

When we say that the sub-systems are ‘separated in space’ and when we refer to the ‘initial physical state’ of the system, we have in mind the physical meaning of the concepts ‘space’ and ‘state’. ‘Input’ and ‘output’ should also be understood as physical operations: the input is a physical manipulation which the player performs on the system while the output is the result of a physical measurement which the player performs on the system. Thus, which mechanisms can be implemented depend on the laws of physics that govern the universe. One does not have to be an expert in the laws of physics in order to understand that without any additional restrictions every garbling $q \in \mathcal{S}(S_1 \times S_2, S'_1 \times S'_2)$ can be implemented by such a physical system. This could be implemented by a pair of computers and a long cable that connects Earth and Mars, through which the computers communicate. Difficulties arise when restriction are imposed.

Suppose that the players in both locations insert their inputs *simultaneously* and that the implementation occurs *instantaneously*. Again, the terms ‘simultaneously’ and ‘instantaneously’ should be understood in their physical meaning: input should be inserted

at the same time and outputs should occur immediately, regardless of how far the players are from each other. Under this restriction the laws of physics entail that not every garbling can be physically implemented. For example, a garbling $q \in \mathcal{S}(S_1 \times S_2, S_2 \times S_1)$ in which $q(s_2, s_1 | s_1, s_2) = 1$ (that is, each player's output equals, with probability 1, to the other player's input) cannot be implemented, at least if, as modern physics asserts, the speed of information transmission is limited to the speed of light. If one player is in Earth and the other in Mars then, under any physical implementation of such a garbling, there must be a couple of minutes delay between the time player 1 inserts his input and the time player 2 reads her output.

A natural question arises then as to which garblings have an instantaneous physical implementation. The answer to this question ultimately depends on the physical theory one subscribes to. Our purpose in this section is to give the reader some intuition about the physical assumptions that correspond to the restrictions on garblings that we proposed.

We first consider classical physics, which reflects the way physicists understood the universe before the quantum revolution. At the abstract level needed for our argument, classical physics is compatible with a layman intuition about the universe. According to this view, every physical system can be described by a set Ω , the *state space*, or the *phase space* whose elements are called *states*. The state of the system *determines* the output given any input. (In Newtonian mechanics the state of the system is given by the location and momentum of all the particles in the system). The deterministic aspect of classical physics was famously illustrated by Laplace assertion that a daemon who, at a certain moment, knows the state ω will know also the outcome of all physical measurements performed on the system. For a bi-partite system, the state space can be written as $\Omega = \Omega_1 \times \Omega_2$ where Ω_i is the state space of sub-system i . The state of sub-system i and the input at sub-system i completely determine the output at this sub-system. The initial state of the bi-partite system may be chosen at random (according to some probability distribution P over $\Omega_1 \times \Omega_2$) during the preparation stage. Thus, we reach at the following definition:

Definition 5.1. A garbling $q \in \mathcal{S}(S_1 \times S_2, S'_1 \times S'_2)$ has a *classical implementation* if

$$q(s_1, s_2) = P(\{(\omega_1, \omega_2) | f_1(\omega_1, s_1) = s'_1, f_2(\omega_2, s_2) = s'_2\})$$

for some sets Ω_1, Ω_2 , some distribution $P \in \Delta(\Omega_1 \times \Omega_2)$ and some functions $f_i : \Omega_i \times S_i \rightarrow S'_i$.

A classical implementation of a garbling is sometimes called *locally deterministic*. In order to derive the outcome s'_i of a measurement at sub-system i , Laplace's Daemon relies only on the state ω_i of sub-system i and the operation s_i performed on it. It is easy to check that the garbling that have classical implementation are exactly the coordinated garblings introduced in Definition 2.5. It is clear that the garbling mentioned above

where each player’s output equals, with probability 1, to the other player’s input has no classical implementation.

It is a remarkable discovery of quantum physics that some garblings which are not coordinated can still have instantaneous implementation. This fact is usually interpreted as showing that quantum physics is incompatible with the very intuitive ‘local determinism’ idea.

We are not going to give an explicit description of the garblings permitted by quantum physics, but it turns out that all of them are noncommunicating garblings. In fact, we defined noncommunicating garblings as the ones that do not transmit information between the players. Thus, an instantaneous implementation of a noncommunicating garbling is compatible with the law, first postulated by Einstein, that information cannot travel instantaneously (or at least, this is how physicists interpret Einstein’s assertion following the quantum revolution.)

John Bell was the first to demonstrate the idea that quantum physics allows instantaneous implementation of non-coordinated garblings. In fact, it is now popular to formulate Bell’s argument in terms of payoffs the players can receive in a common interest game with incomplete information [7]. To the best of our knowledge, the first construction similar to our Definition of noncommunicating garbling was in a paper of Tsirelson [23]. In recent quantum physics literature non-communicating garblings are called *no-signaling boxes* and coordinating garblings are sometimes called *local*. The reader is referred to Barrett et al. [3] and the citations therein for more on the physics literature. Also, Shmaya [22] discusses the physics of stochastic maps in the context of Blackwell’s theory.

6. DISCUSSION

This paper originated in an attempt to generalize Blackwell’s Theorem about comparison of experiments in a single agent decision problem. The single agent setup can be viewed as a special case of the framework of our paper with the number of players $n = 1$. In this case, all the classes of garblings considered in this paper are the same, and all the equilibrium concepts are the same. Moreover, for a given information structure and a given basic game, all equilibria induces the same payoff. One way to formulate Blackwell’s Theorem is that, when $n = 1$, an information structure $\sigma : K \rightarrow S$ gives a higher payoff than the information structure $\sigma' : K \rightarrow S'$ if and only if σ is *more informative* than σ' in the sense that $\sigma' = q \circ \sigma$ for some garbling $q \in \mathcal{S}(S, S')$.

In our companion paper (Lehrer et al. [17]) we demonstrated a similar theorem to Blackwell’s for the case of multi-player common interest setup, that is when all players have the same payoff. In this case the players can get a higher equilibrium payoff under information structure σ than under σ' if and only if σ is more informative than σ' , where the relation ‘more informative’ is defined using garblings. In the case of possible conflicting interests between the players, more information on the one hand provides

more possibilities to the players, but on the other hand reduces equilibria because of the additional incentive compatibility constraints. We proved in this paper a more modest theorem: That two information structures provides the same information for the players if and only if they generate the same equilibria. As we have seen, there are several ways to define what it means for information structures to provide the same information, which correspond to different restrictions on the garblings.

REFERENCES

- [1] R. J. Aumann, *Subjectivity and correlation in randomized strategies*, J. Math. Econom. **1** (1974), no. 1, 67–96.
- [2] R. J. Aumann, *Correlated equilibrium as an expression of Bayesian rationality*, Econometrica **55** (1987), no. 1, 1–18.
- [3] J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts, *Nonlocal correlations as an information-theoretic resource*, Phys. Rev. A **71** (2005), no. 2, 022101.
- [4] D. Blackwell, *Equivalent comparisons of experiments*, Ann. Math. Statistics **24** (1953), 265–272.
- [5] D. Bergemann and S. Morris, *Correlated equilibrium in games with incomplete information* (2010), preprint.
- [6] D. Bergemann and S. Morris, *Robust Predictions in Games with Incomplete Information* (2013), preprint.
- [7] R. Cleve, P. Hoyer, B. Toner, and J. Watrous, *Consequences and Limits of Nonlocal Strategies*, Proc. of 19th IEEE Conference on Computational Complexity (2004).
- [8] K. D. Cotter, *Correlated equilibrium in games with type-dependent strategies*, J. Econom. Theory **54** (1991), no. 1, 48–68.
- [9] E. Dekel, E. Fudenberg, and S. Morris, *Interim rationalizability*, Theoretical Economics **4** (2007), no. 2, 15–40.
- [10] J. C. Ely and M. Peşki, *Hierarchies of belief and interim rationalizability*, Theoretical Economics **1** (2006), 19–65.
- [11] F. Forges, *Correlated equilibria in repeated games with lack of information on one side: a model with verifiable types*, Internat. J. Game Theory **15** (1986), no. 2, 65–82.
- [12] F. Forges, *Five legitimate definitions of correlated equilibrium in games with incomplete information*, Theory and Decision **35** (1993), no. 3, 277–310.
- [13] F. Forges, *Correlated equilibrium in games with incomplete information revisited*, Theory and Decision **61** (2006), no. 4, 329–344.
- [14] O. Gossner, *Comparison of information structures*, Games Econom. Behav. **30** (2000), no. 1, 44–63.
- [15] O. Gossner and J.-F. Mertens, *The Value of Information in Zero-Sum Games*, 2001.
- [16] J. Hirshleifer, *The private and social value of information and the reward to inventive activity*, The American Economic Review **61** (1971), no. 4, 561–574.
- [17] E. Lehrer, D. Rosenberg, and E. Shmaya, *Signaling and mediation in games with common interests*, Games Econom. Behav. **68** (2010), no. 2, 670–682.
- [18] J. Marschak and K. Miyasawa, *Economic comparability of information systems*, International Economic Review **9** (1968), no. 2, 137–174.
- [19] I. Milchtaich, *Implementability of correlated and communication equilibrium outcomes in incomplete information games*, Department of Economics, Bar Illan University, 2009.
- [20] M. Peski, *Comparison of information structures in zero-sum games*, Games and Economic Behavior **62** (2008), no. 2.

- [21] L. Samuelson and J. Zhang, *Correlated equilibria and mediated equilibria in games with incomplete information*, Penn State University, 1989.
- [22] E. Shmaya, *Comparison of information structures and completely positive maps*, J. Phys. A **38** (2005), no. 44, 9717–9727.
- [23] B. S. Tsirelson, *Some results and problems on quantum Bell-type inequalities*, Hadronic J. Suppl. **8** (1993), no. 4, 329–345.

TEL AVIV UNIVERSITY, TEL AVIV 69978, ISRAEL
E-mail address: `lehrer@post.tau.ac.il`

HEC PARIS, FRANCE.
E-mail address: `rosenberg@hec.fr`

KELLOGG SCHOOL OF MANAGEMENT, NORTHWESTERN UNIVERSITY
E-mail address: `e-shmaya@kellogg.northwestern.edu`