Strategically Valuable Information[*]

Josh Cherry[†] Lones Smith[‡]

Department of Economics
University of Michigan
and
Department of Economics
University of Wisconsin

December 7, 2014

Abstract

In a 1953 classic paper, Blackwell showed that experiment A is sufficient for B if and only if a decision maker can attain a larger set of payoffs with A than with B in any decision problem. In this paper, we ask the Blackwell equivalence question in a strategic setting. In other words, does there exist a partial order on information held by players in a game that reflects "more" or "better" information, which coincides precisely with the ability to induce more equilibrium payoff vectors in all Bayesian games? If so, we say that it is "strategically more valuable". In this paper, we define a meaningful sense in which information structures can be compared by how "strategically informative" they are. Combining the two notions, we answer our original question in the affirmative: There exists an intuitive definition and characterization of the partial order *more strategically informative*, and it is equivalent to the partial order *more strategically valuable*. The conditions we provide are easily checked, are useful in an array of economic settings, and have straightforward geometric interpretations.

Our main theorem applies to a wide variety of economic environments of interest endowed with commonly used information structures. For example, sunspots are a frequently used tool in general equilibrium theory. Our results provide a natural partial ordering on sunspot equilibria, regardless of the environment in which they operate. The centerpiece application is to repeated games with private monitoring where the more strategically informative order ranks monitoring structures. Consequently, we can show when a change in monitoring structure will weakly expand the set of sequential equilibria. This mirrors a classic result of **?** for repeated games with imperfect public monitoring.

[*]We are grateful for helpful comments from seminar participants at the University of Michigan.

[†]email: jscherry@umich.edu. Josh is thankful for the financial support of NSF IGERT grants 0114368 and 0654014.

[‡]email: lones@ssc.wisc.edu. Lones is grateful for the ongoing financial support of the NSF that has supported this work.

1 Introduction

In a 1953 classic paper, Blackwell compares two partial orderings on experiments, or informative Bayesian signals. The first ranking is by statistical sufficiency — experiment A is sufficient for B if B is a statistical garbling of A, or equivalently, B can be attained by adding noise to A. The second ranking is the economic value of the signals for decision problems with state-dependent payoffs. Experiment A is more valuable than B if the decision maker can attain a larger set of payoffs with A than with B in any decision problem. Blackwell showed that experiment A is sufficient for B if and only if it is more valuable.

In this paper, we ask the Blackwell equivalence question in a strategic setting. In other words, does there exist a partial order on information held by players in a game that reflects "more" or "better" information, which coincides precisely with the ability to induce more equilibrium payoff vectors in all Bayesian games? If so, we say that it is "strategically more valuable". In this paper, we define a meaningful sense in which information structures can be compared by how "strategically informative" they are. Combining the two notions, we answer our original question in the affirmative: There exists an intuitive definition and characterization of the partial order *more strategically informative*, and it is equivalent to the partial order *more strategically valuable*. The conditions we provide are easily checked, are useful in an array of economic settings, and have straightforward geometric interpretations.

Our main theorem applies to a wide variety of economic environments of interest endowed with commonly used information structures. For example, sunspots are a frequently used tool in general equilibrium theory. Our results provide a natural partial ordering on sunspot equilibria, regardless of the environment in which they operate. The centerpiece application is to repeated games with private monitoring. In a repeated game with private monitoring, each period players simultaneously choose actions, after which each player privately observes a signal informative of the action profile most recently played. Neither the actions nor signal realizations are ever observed by any other player. The more strategically informative order ranks monitoring structures — the probability distribution on private signals — in much the same way. Consequently, we can show when a change in monitoring structure will weakly expand the set of sequential equilibria. This mirrors a classic result of **?** for repeated games with imperfect public monitoring.

The most closely related paper to this one is that by **?**. He studies the same question, and succeeds in obtaining a different (albeit equivalent) characterization. We feel that the current approach is superior in several dimensions. First, our condition is easier to verify — it corresponds geometrically to well understood statistical concepts. Second, while the proofs in **?** are indirect and complex, we provide a straightforward, illuminating proof that leverages the separation argument at the core of Blackwell's Theorem.

2 Model

2.1 The Informational Setting

We begin by considering a standard multi-player Bayesian environment, as in **?**. The incomplete information is about the state of the world ω an element of the state space Ω, which we assume is a Borel set. The set $N = \{1, \ldots, n\}$ indexes players. For each player i there exists a (possibly infinite) *partition* $\mathcal{P}_i$ of Ω, a mutually exclusive and exhaustive family of subsets of Ω. At a state $\omega \in \Omega$ the set $p_i(\omega) \in \mathcal{P}_i$ is the set of states indistinguishable by player i, which we will call a *cell*. Similarly, let $\mathcal{P} = \{\mathcal{P}_i\}_{i=1}^n$ be the *joint partition* of all players.

An *information structure* $\mathcal{I}$ is a triple $(\Omega, \mu, \{\mathcal{P}_i\}_{i=1}^n)$ of a state space Ω, a measure μ over Ω, and a joint partition $\mathcal{P} = \{\mathcal{P}_i\}_{i=1}^n$. Assuming that $\Omega = [0,1]$ and μ is the Lebesgue measure entails no loss of generality, so in much of what follows we do so. It then suffices to describe an information structure by the joint partition $\mathcal{P}$.

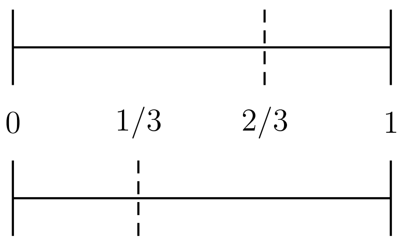

Figure 1: Information Structure $\mathcal{P}$

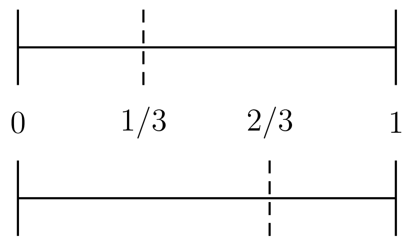

Figure 2: Information Structure $\mathcal{Q}$

Consider the following two player information structures: $\mathcal{P}_1 = \{[0,1/3),[1/3,1]\}$ $\mathcal{P}_2 = \{[0,2/3),[2/3,1]\}$ and $\mathcal{Q}_1 = \{[0,2/3),[2/3,1]\}$ $\mathcal{Q}_2 = \{[0,1/3),[1/3,1]\}$.

The information structure $\mathcal{Q}$ is obtained from $\mathcal{P}$ by 'relabeling', in this case reversing the direction of the interval $[0,1]$. Since the state is not payoff relevant, it is natural to require that these information structures have equivalent strategic effects in any environment. In fact, by the same logic any (measure preserving) permutation of the state space should be similarly inconsequential and for the purposes of this paper be treated as equivalent. Let $[\mathcal{P}]$ be the set of information structures equivalent to $\mathcal{P}$: $\mathcal{Q} \in [\mathcal{P}]$ if there exists a transformation t in the permutation group $T(\Omega)$ such that $\mathcal{P} = t(\mathcal{Q})$.[1] We can then define the equivalence relation $\sim$: $\mathcal{P} \sim \mathcal{Q}$ if and only if $\mathcal{Q} \in [\mathcal{P}]$. This relation is reflexive, symmetric and transitive. In what follows, we will work on the quotient space of information structures modulo $\sim$, defined by the canonical projection from the space of all information structures to the quotient space. Therefore, without loss of generality, we will not distinguish between elements of equivalence class.

Let $\sigma(\mathcal{P}_i)$ be the σ-algebra generated by $\mathcal{P}_i$, and $\sigma(\mathcal{P})$ the σ-algebra generated by $\mathcal{P}$ on the product space Ω^n. Information structure $\mathcal{P}$ *refines* $\mathcal{Q}$ if $\sigma(\mathcal{Q}) \subseteq \sigma(\mathcal{P})$. Let $\sigma(\mathcal{P}_i)$ be the σ-algebra generated by $\mathcal{P}_i$, and $\sigma(\mathcal{P})$ the σ-algebra generated by $\mathcal{P}$ on the product space Ω^n. Information structure $\mathcal{P}$ *refines* $\mathcal{Q}$ if $\sigma(\mathcal{Q}) \subseteq \sigma(\mathcal{P})$. Two sub-$\sigma$-algebras $\mathcal{F}, \mathcal{G}$ are *conditionally independent given σ-algebra* $\mathcal{H}$ — written $(\mathcal{F} \perp \mathcal{G})|\mathcal{H}$ — if any two events $F \in \mathcal{F}$ and $G \in \mathcal{G}$ are conditionally independent given $\mathcal{H}$, namely:

$$\mu(F \cap G|\mathcal{H}) = \mu(F|\mathcal{H})\mu(G|\mathcal{H})$$

Information structure $\mathcal{P}$ is *more strategically informative* than $\mathcal{Q}$ — written $\mathcal{P} \gg \mathcal{Q}$ — if $\mathcal{P}$ refines $\mathcal{Q}$ and for every player i $(\mathcal{P}_i \perp \mathcal{Q}_{-i})|\mathcal{Q}_i$. The first part of the condition requires

[1] A *permutation group* (of transformations) T is a set of measure preserving bijective automorphisms on a set X such that: (a) If $t_1 \in T$ and $t_2 \in T$, then $t_1 \circ t_2 \in T$; (b) For all transformations $t \in T$ there exists $t^{-1} \in T$ such that for all elements $x \in X$we have $(t^{-1} \circ t)(x) = x$, i.e. there exists an inverse; and (c) There exists $e \in T$ such that $e(x) = x$, i.e. there exists an identity.

that each player have a finer partition. The second part requires that the information added by $\mathcal{P}$ (relative to $\mathcal{Q}$) be safely ignored: if players $-i$ ignore the added information, nothing is lost by player i by doing so as well. Conditional independence means that player i's information $\mathcal{Q}_i$ is sufficient for $\mathcal{Q}_{-i}$. Therefore, it does not help player i make inferences about the other players.

Example 1.

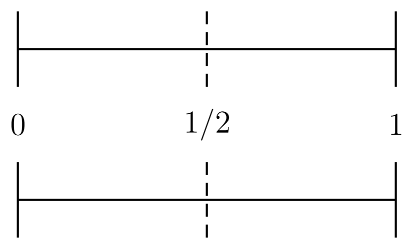

Figure 3: Sunpot with Two Octcomes

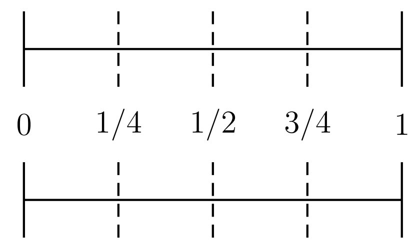

Figure 4: Sunspot with Four Outcomes

Consider the following two information structures, each for two players:

$$\mathcal{P}_1 = \mathcal{P}_2 = \{[0, 1/2), [1/2, 1]\}$$

and

$$\mathcal{Q}_1 = \mathcal{Q}_2 = \{[0, 1/4), [1/4, 1/2), [1/2, 3/4), [3/4, 1]\}$$

In this example, $\mathcal{P}$ is a sunspot with two outcomes and $\mathcal{Q}$ is a sunspot with four outcomes. $\mathcal{P}$ is strategically more informative than $\mathcal{Q}$. $\square$

Example 2.

Now, consider the following three information structures, each for two players:

$$\mathcal{P}_1 = \mathcal{P}_2 = \{[0, 1/2), [1/2, 1]\}$$

$$\mathcal{Q}_1 = \{[0, 1/3), [1/3, 1]\} \quad \mathcal{Q}_2 = \{[0, 2/3), [2/3, 1]\}$$

$$\mathcal{R}_1 = \{[0, 1/6), [1/6, 1/2), [1/2, 2/3), [2/3, 1]\} \quad \mathcal{R}_2 = \{[0, 1/3), [1/3, 1/2), [1/2, 5/6), [5/6, 1]\}$$

In this example, $\mathcal{P}$ is again a sunspot with two outcomes. The information structure $\mathcal{Q}$ is the signal used in Aumann's original paper on correlated equilibrium **?**. Neither $\mathcal{P} \gg \mathcal{Q}$ nor $\mathcal{Q} \gg \mathcal{P}$. However, $\mathcal{R} \gg \mathcal{P}$ and $\mathcal{R} \gg \mathcal{Q}$.

Example 3 (Bivariate Gaussian Signals)**.**
Suppose there are two players who each observe a signal generated by a bivariate normal distribution with variance normalized to 1 and covariance ρ. With a slight abuse of notation, let $\mathcal{P}_\rho$ refer to the information structure induced by bivariate Gaussian signals with covariance ρ. Any correlation is more strategically informative than independent signals,

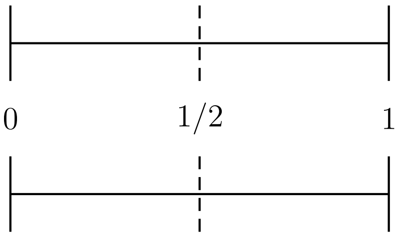

Figure 5: Information Structure $\mathcal{P}$

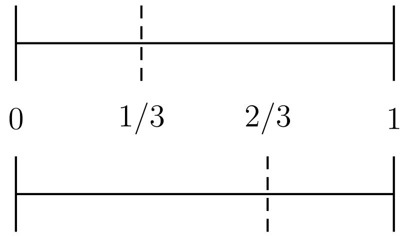

Figure 6: Information Structure $\mathcal{Q}$

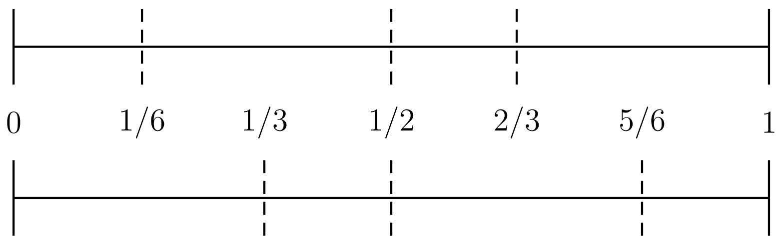

Figure 7: Information Structure $\mathcal{R}$ is constructed by embedding $\mathcal{Q}$ on both sides of $\mathcal{P}$

that is if $\rho \neq 0$ we have $\mathcal{Q}_\rho \gg \mathcal{Q}_{\mathbf{0}}$, where $\mathbf{0}$ denotes independence. It should also be clear that by a simple transformation $\mathcal{Q}_\rho$ and $\mathcal{Q}_{-\rho}$ are equally strategically informative. However, for distinct non-zero covariances p, q inducing information structures $\mathcal{P}$ and $\mathcal{Q}$, respectively, neither $\mathcal{P} \gg \mathcal{Q}$ nor $\mathcal{Q} \gg \mathcal{P}$.

Example 4 (Multivariate Gaussian Signals)**.**
When the information structure is composed of normally distributed signals, even when there are n players, the more strategically informative ordering has a particularly convenient representation. Suppose there is a vector valued random variable $X \sim N(0, \Sigma_X)$, where Σ_X is an $n \times n$ positive semi-definite (covariance) matrix. Then the state space $\Omega = \mathbb{R}^n$. Each of n players is informed of the projection of X onto Ω_i, i.e. the i-th component X_i of X. Suppose the random variable $Y \sim N(0, \Sigma_Y)$. Let Σ_{XY} be the covariance matrix for the random variable $[X, Y]$. Let M_i be an $(n+1) \times (n+1)$ matrix obtained by eliminating rows and columns 1 to $i-1$ and $i+1$ to n. M_i is the covariance matrix for player i's information under X and all players information under Y. By a known characterization of conditional independence for multivariate Gaussian distributions, we have the following result. *The random variable X is more strategically informative than Y if and only if M_i is singular for every player i.*

2.2 The Strategic Environment

We now turn to the strategic aspects of these environments. An n-player game G in normal form is a set of players $N = \{1, 2, \ldots, n\}$, a set A_i of pure actions available to each player i, and a function g from the set of action profiles $A = \prod_i (A_i)$ into $\mathbb{R}^n$. Let Γ_n be the set of all n-player games in normal form. A game $G \in \Gamma_n$ *extended by* $\mathcal{P}$ allows each player to choose his action conditional on the cell of his partition in which he finds himself. Therefore, a strategy is now a function f_i from $\mathcal{P}_i$ to A_i, where $f_i(p_i)$ is the action played by player i in cell p_i. The expected payoff function $E[g(f(p)|\mathcal{P})]$ is evaluated in the obvious way by using Bayes rule. A strategy profile f is a *Bayes-Nash equilibrium* if

$$E[g_i(f(p)|\mathcal{P})] \geq E[g_i(f_i'(p_i), f_{-i}(p_{-i})|\mathcal{P})]$$

for all players i and all strategies f_i'. If f is a Bayes-Nash equilibrium, then $E[g(f(p)|\mathcal{P})]$ is a Bayes-Nash equilibrium payoff. Let $\Pi(G|\mathcal{P})$ be all Bayes-Nash equilibrium payoffs in the game G extended by information structure $\mathcal{P}$.

An information structure $\mathcal{P}$ is *more strategically valuable* than information structure $\mathcal{Q}$ — written $\mathcal{P} \sqsupseteq \mathcal{Q}$ — if $\Pi(G|\mathcal{P}) \supseteq \Pi(G|\mathcal{Q})$ for all games $G \in \Gamma_n$.

3 The Equivalence Result

We are now in a position to state the main result of this paper, which we prove in the appendix.

Theorem 1. *Information structure $\mathcal{P}$ is more strategically valuable than information structure $\mathcal{Q}$ if and only if $\mathcal{P}$ is more strategically informative than $\mathcal{Q}$.*

$$\mathcal{P} \sqsupseteq \mathcal{Q} \Leftrightarrow \mathcal{P} \gg \mathcal{Q}$$

We now return to our previous example to demonstrate the implications of Theorem 1. As we showed, neither $\mathcal{P}$ nor $\mathcal{Q}$ is more strategically informative than the other. Theorem 1 implies that they are also not ranked by the more strategically valuable relation. To see this, we will demonstrate a game in which information structure $\mathcal{P}$ is induces an equilibrium payoff not attained with $\mathcal{Q}$, and vice versa. In the coordination game depicted in Figure 8, the signal provided by $\mathcal{P}$ generates the payoff $(3/2, 3/2)$. This payoff requires perfect coordination and so cannot be obtained with $\mathcal{Q}$. Similarly, the information structure $\mathcal{Q}$ is necessary to obtain the payoff $(10/3, 10/3)$ in the game of chicken, depicted in Figure 9.

	L	R
U	(2,2)	(0,0)
D	(0,0)	(1,1)

Figure 8: Pure Coordination Game

	L	R
U	(4,4)	(1,5)
D	(5,1)	(0,0)

Figure 9: Game of Chicken

4 Economic Applications

4.1 Application to Repeated Games

? shows that in a repeated game with imperfect public monitoring, the equilibrium payoff set expands in the accuracy of the public signal. Our result extends this line of thought to repeated games with private monitoring. In particular, we show that the set of sequential equilibrium payoffs grows when the monitoring structure becomes more strategically informative.

A repeated game is played in periods $1, 2, \ldots$ Each period, every player $i \in N = \{1, 2, \ldots, n\}$ chooses an action a_i from a finite action set A_i. After play any period, each player receives a private message m_i from a finite set M_i. A *monitoring structure* ψ is a collection of $|A|$ probability distributions $\{\psi(\cdot|a) \in \Delta(M) \,|\, a \in A\}$ on the message profile set $M = \prod_i M_i$. Let the set of all monitoring structures be Ψ. After an action profile a is realized, a message profile $m = (m_1, \ldots, m_n)$ is drawn with chance $\psi(m|a)$, and each player i is then privately informed of his component message m_i.

In each period, a player observes his realized action $a_i \in A_i$ and private message m_i. Let the null history h_i^1 be player i's history before play begins. A *private history* h_i^t is the complete record of player i's past actions $(a_i^1, \ldots, a_i^{t-1})$ and past private messages $(m_i^1, \ldots, m_i^{t-1})$, including the null history. Let H_i^t be the set of all possible private histories h_i^t for player i, and $H_i = \bigcup_{t=1}^{\infty} H_i^t$ the set of all such histories of any length. A (behavior) strategy s_i is a sequence of functions $\{s_i^t\}_{t=1}^{\infty}$, where $s_i^t : H_i^t \to \Delta(A_i)$ for every period $t = 1, 2, 3, \ldots$ In other words, it maps every private into a mixed action. Let $\mathcal{S}$ be the space of all such strategy profiles $s = (s_1, \ldots, s_n)$. Given the strategy profile $s \in \mathcal{S}$, Bayes' rule and the Law of Total Probability naturally imply beliefs and behavior at all future information sets.

Each length t private history, together with a strategy profile, implies an ex-ante distribution on the product space $A^t \times M^t$. Each player, being informed only of his own actions and signals, entertains a natural partition of $A^t \times M^t$. For each period t, this information structure is defined endogenously by the distribution of mixed actions and monitoring signals. Let $\Phi^t(\psi, s)$ be the (ex-ante expected) information structure in period t under monitoring structure ψ and strategy profile s.

A monitoring structure ψ^1 is *more strategically informative* than ψ^2 if

$$\left(\psi_i^2(\cdot|a) \perp \psi_{-i}^1(\cdot|a)\right) \,\big|\, \psi_i^1(\cdot|a)$$

for every player i and every action profile a. In a static context, this precisely coincides with the previous definition.

Lemma 1. *If a monitoring structure ψ^1 is more strategically informative than ψ^2, then $\Phi^t(\psi^1, s)$ is more strategically informative than $\Phi^t(\psi^2, s)$ for any strategy profile s and any period t.*

Let $G_\psi(\delta)$ denote the infinitely repeated game of private monitoring with monitoring structure ψ, played in periods $t = 1, 2, 3, \ldots$. Payoffs are discounted as usual by the factor $0 < \delta < 1$. Let $v_i : \mathcal{S} \to \mathbb{R}$ be the discounted average payoff for player i in the repeated game $G_\psi(\delta)$. While more precisely presented in the Appendix, here we write that player i's discounted average payoff starting in period t from the strategy profile s is $v_i^t(s|h_i^t)$. Then a strategy profile s is a *sequential equilibrium* of $G_\psi(\delta)$ if and only if no player can ever profitably deviate, i.e. $v_i(s|h_i^t) \geq v_i(\tilde{s}_i, s_{-i}|h_i^t)$ for every private history h_i^t and strategy $\tilde{s}_i : H_i \to \Delta(A_i)$ of every player i. Since playing a Nash equilibrium of G after every history is a sequential equilibrium, existence is guaranteed. Let V_ψ be the set of sequential equilibrium payoff vectors of the mediated game $G_\psi(\delta)$

Suppose monitoring structure ψ^1 is more strategically valuable than monitoring structure ψ^2. The main theorem implies that the sequential equilibrium repeated game payoff set with ψ^1 contains that of ψ^2. In the appendix, we prove:

Theorem 2. *If a monitoring structure ψ^1 is more strategically informative than ψ^2, then $V_{\psi^1} \supseteq V_{\psi^2}$*

4.2 Sunspots in General Equilibrium

This result has important implications for general equilibrium theory. Since **?** showed that Walrasian settings can be interpreted as games, our result applies to markets as well. Public signals that are not payoff relevant is often used in these environments to add convexity to outcomes. These signals, popularly known as sunspots, make goods divisible and help with equilibrium selection. Consequently, sunspots play a central role in many general equilibrium models.

We reinterpret competitive markets as games and model sunspots precisely with information structures. By doing so, the main result of this paper allows us to say with certainty when "better" public information forces the set of sunspot equilibria to grow. For example, as in Example XX, adding nested outcomes necessarily makes the set of equilibria (weakly) expand. However, suppose all agents observe the realization of a Gaussian random variable before acting. In this case the variance of the random variable is of no consequence; all sunspots with any positive variance are equally strategically valuable since observations can be transformed by any non-zero scalar.

5 Conclusion

At first glance, the main result should extend to games of incomplete information by introducing the player "Nature", as in the tradition of Harsanyi. However, doing so requires the same conditional independence conditions between each player and Nature as between any two players. For intuition, consider the following simple example. There are two players $\{1, 2\}$, two states of the world $\{l, h\}$ and two pure actions for each player $\{L, H\}$. The game is zero-sum, with player 1 as the maximizer. He earns a payoff of 1 when his action matches the state, and zero otherwise. More precise information about the state increases his payoff, which necessarily lowers his opponent's payoff.

Appendix

Proof that more strategically informative is sufficient for more strategically valuable:

Assume that $\mathcal{P}$ is more strategically informative than $\mathcal{P}$. Let $G \in \Gamma_n$ be an arbitrary game and $v \in \Pi(G, \mathcal{Q})$ an arbitrary Bayes-Nash equilibrium payoff attained by the strategy profile f. Since $\mathcal{P}$ is a refinement of $\mathcal{Q}$, the strategy profile f is measurable with respect to $\mathcal{Q}$. Furthermore, since $\mathcal{P}_i$ is conditionally independent of $\mathcal{Q}_{-i}$, f is still a Bayes-Nash equilibrium.

Proof of more strategically informative is necessary for more strategically valuable:

Suppose that $\mathcal{P}$ is more strategically informative than $\mathcal{Q}$. Since $\Pi(G, \mathcal{Q}) \subseteq \Pi(G, \mathcal{P})$ for every game $G \in \Gamma_n$, the inclusion in particular holds for all *decision games*: a game G^{D_i} in which players $-i$ earn a payoff of zero for every action profile, and the remaining player i — the decision maker — has non-null payoffs. Equilibrium strategies are by definition known to the decision maker, and so this game is equivalent to a standard choice under uncertainty problem, where the players $-i$ (with null payoffs) take the role of Nature. Let $\Pi_i(G^{D_i}, \mathcal{P})$ be set of payoffs to admissible decision rules in G^{D_i}. By Blackwell's Theorem, the hypothesis $\Pi_i(G^{D_i}, \mathcal{Q}) \subseteq \Pi_i(G^{D_i}, \mathcal{P})$ for all decision games is equivalent to $\mathcal{P}_i$ being sufficient for $\mathcal{Q}_{-i}$. Then by the Fisher factorization theorem, $\mathcal{P}_i$ is sufficient for $\mathcal{Q}_{-i}$ if and only if $\mathcal{P}_i$ it can be factored into two terms, one of which is conditionally independent of $\mathcal{Q}$.