Correlated Equilibrium in Games with Incomplete Information

Dirk Bergemann and Stephen Morris

Econometric Society Summer Meeting
June 2012

Robust Predictions Agenda

- game theoretic predictions are very sensitive to "information structure" a.k.a. "higher order beliefs"
- suppose the modeller is uncertain about the information structure of players, what predictions can he make?
- more specifically, what if the modeller knows that the players know something (e.g., their valuations / utility functions) but perhaps they know more (e.g., they have signals about others' valuations/utility functions)
- call these robust predictions....

More Formal Statement of Basic Question

- fix a game of incomplete information, divide it into:

1. "basic game" $=$ action sets, utility functions (depending on actions and states), common prior distribution over states
2. "information structure" (aka, higher order beliefs, type space) represents players (common prior) information about states

- for the fixed game of incomplete information, "robust predictions" are the set of distributions over actions and states that could arise in Bayes Nash equilibrium in that game of incomplete information *or one in which players observed additional information*.

- Robust predictions =
- set of outcomes consistent with equilibrium given any additional information the players may observe

- Robust predictions =

- set of outcomes consistent with equilibrium given any additional information the players may observe
- set of outcomes that could arise if a mediator who knew the payoff state could privately make action recommendations

- Robust predictions $=$

- set of outcomes consistent with equilibrium given any additional information the players may observe
- set of outcomes that could arise if a mediator who knew the payoff state could privately make action recommendations
- set of incomplete information correlated equilibrium outcomes

- Robust predictions =
- set of outcomes consistent with equilibrium given any additional information the players may observe
- set of outcomes that could arise if a mediator who knew the payoff state could privately make action recommendations
- set of incomplete information correlated equilibrium outcomes
- we shall call this very permissive version of incomplete information correlated equilibrium "Bayes correlated equilibrium"

- in related work, we study the set of Bayes correlated equilibria / robust predictions in particular economic / game theoretic settings:
- "Robust Predictions in Games with Incomplete Information" (linear best response games with continuum of agents),
- "Extremal Information Structures in First Price Auctions" (joint with Ben Brooks)
- today we examine two (related) foundational issues motivated by this agenda:

1. More prior information should lead to narrower predictions: suggests ordering on (many player) information structures
2. Relation to (large) literature on incomplete information correlated equilibrium

Outline of Talk: Single Player Case

- Bayes correlated equilibrium with single player:
  what predictions can we make in a one player game ("decision problem") if we have just a lower bound on the player's information structure ("experiment")?
- we suggest a partial order on experiments:
  one experiment is more *incentive constrained* than another if it gives rise to smaller set of possible BCE outcomes across all decision problems

Blackwell's Theorem

- Blackwell (1951) introduced two classic orderings on experiments:

1. an experiment S is *more permissive* than S' if it allows the player to achieve more outcomes in every decision problem
   - Blackwell required that the set of implied ex ante payoffs be larger under S than S', and called the relation "more informed than"
   - economists often focus on the largest possible payoff and call the relation "more valuable than"
2. an experiment S is *sufficient* for experiment S' if S' can be obtained from S by a stochastic transformation

Theorem

The following are equivalent:

1. *Experiment S is more permissive than experiment S' (feasibility ordering);*
2. *Experiment S is sufficient for S' (statistical ordering);*
3. *Experiment S is more incentive constrained than S' (incentive ordering);*

- a many person equivalent of Blackwell's Theorem plus:

Theorem

The following are equivalent:

1. *Information structure S is more permissive than S' (feasibility ordering)*
2. *Information structure S is individually sufficientfor S' (statistical ordering)*
3. *Information structure S is more incentive constrained than S' (incentive ordering)*

- we use the distinction between incentive compatibility and feasibility as a lens to understand the literature on correlated equilibrium with incomplete information

1. Forges (1993, 2006): many notions of incomplete information correlated equilibrium
2. Lehrer, Rosenberg and Shmaya (2010, 2012): many multi-player versions of Blackwell's Theorem
3. Liu (2005, 2012): one more (important for us) version of incomplete information correlated equilibrium and a characterization that is useful for our ordering

- single decision maker
- finite set of *payoff states* Θ and prior

$$\psi \in \Delta(\Theta)$$

- a *decision problem* $G = (A, u)$, where A is a finite set of *actions* and

$$u : A \times \Theta \to \mathbb{R}$$

is the agent's (vNM) utility

- an *experiment* $S = (T, \pi)$, where T is a finite set of *types* (i.e., signals) and likelihood function

$$\pi : \Theta \to \Delta(T)$$

- a game of incomplete information is (G, S)

- behavior in (G, S) is described by a joint distribution of actions, signals, and states:

$$\nu \in \Delta(A \times T \times \Theta)$$

Definition

A distribution $\nu \in \Delta(A \times T \times \Theta)$ is *consistent* for (G, S) if

$$\sum_{a \in A} \nu(a, t, \theta) = \psi(\theta) \pi(t|\theta)$$

for each $t \in T$ and $\theta \in \Theta$.

Definition

A distribution $\nu \in \Delta(A \times T \times \Theta)$ is *attainable* in (G, S) if there exists $f : T \to \Delta(A)$ such that

$$\nu(a, t, \theta) = \left(\sum_{a' \in A} \nu(a', t, \theta) \right) f(a|t)$$

for each $a \in A$, $t \in T$ and $\theta \in \Theta$.

Definition

A distribution $\nu \in \Delta(A \times T \times \Theta)$ is *obedient* for (G, S) if

$$\sum_{\theta \in \Theta} \nu(a, t, \theta)\, u(a, \theta) \geq \sum_{\theta \in \Theta} \nu(a, t, \theta)\, u(a', \theta)$$

for each $a \in A$, $t \in T$ and $a' \in A$.

Definition

A distribution $\nu \in \Delta(A \times T \times \Theta)$ is *a Bayes Nash equilibrium (single agent optimal)* if it is consistent, attainable and obedient.

- say that the value of ν is

$$V \triangleq \sum_{a \in A, t \in T, \theta \in \Theta} \nu(a, t, \theta)\, u(a, \theta)$$

- if ν is optimal, its value V is the highest utility in $(G, S) : V(G, S)$

Definition

An *outcome* is a distribution $\mu \in \Delta(A \times \Theta)$ over actions and states.

- distribution $\nu \in \Delta(A \times T \times \Theta)$ gives outcome $\mu \in \Delta(A \times \Theta)$ if

$$\sum_{t \in T} \nu(a, t, \theta) = \mu(a, \theta), \quad \forall a \in A, \ \forall \theta \in \Theta.$$

- outcome $\mu \in \Delta(A \times \Theta)$ is feasible for (G, S) if there exists a consistent and attainable $\nu \in \Delta(A \times T \times \Theta)$ giving μ
- write $F(G, S)$ for the set of feasible outcomes of (G, S)

- consider separate experiments,

$$S^1 = \left(T^1, \pi^1\right), \qquad S^2 = \left(T^2, \pi^2\right)$$

- join the experiments S^1 and S^2 into $S^* = \left(T^*, \pi^*\right)$:

$$T^* = T^1 \times T^2, \quad \pi^* : \Theta \to \Delta\left(T^1 \times T^2\right)$$

Definition

S^* is a *combined experiment* of S^1 and S^2 if:

1. $T^* = T^1 \times T^2, \ \ \pi^* : \Theta \to \Delta\left(T^1 \times T^2\right)$
2. marginal of S^1 is preserved:

$$\sum_{t^2 \in T^2} \pi^*\left(\left(t^1, t^2\right) | \theta\right) = \pi^1\left(t^1 | \theta\right), \quad \forall t^1, \ \forall \theta.$$

3. marginal of S^2 is preserved:

$$\sum_{t^1 \in T^1} \pi^*\left(\left(t^1, t^2\right) | \theta\right) = \pi^2\left(t^2 | \theta\right), \quad \forall t^2, \ \forall \theta.$$

Combining Experiments and Augmenting Information

- there are multiple combined experiments S^* for any pair of experiments, since only the marginals have to match
- If S^* is combination of S and another experiment S', we say that S^* is an *information augmentation* of S.

(One Person) Robust Predictions Question

- fix a game of incomplete information (G, S)
- which outcomes can arise under optimal decision making in (G, S^*) where S^* is any information augmentation of S?
- S can be the null information structure:

$$S = \{T = \{t_0\}, \quad \pi(\theta \,|\, t_0) = \psi(\theta)\}$$

(One Person) Bayes Correlated Equilibrium

Definition

Distribution $\nu \in \Delta(A \times T \times \Theta)$ is a *Bayes correlated equilibrium* (BCE) of (G, S) if it is obedient and consistent.

- ... not necessarily attainable

Definition

An outcome μ is a BCE outcome of (G, S) if there is a BCE $\nu \in \Delta(A \times T \times \Theta)$ of (G, S) that gives μ.

(One Person) Bayes Correlated Equilibrium

Theorem

An outcome μ is a BCE outcome of (G, S) if and only if there is an information augmentation S^* of S such that μ is a Bayes Nash equilibrium outcome for (G, S^*)

Idea of Proof:

- $(\Leftarrow)$ S^* has "more" obedience constraints than S
- $(\Rightarrow)$ let ν be BCE of (G, S) supporting μ and consider augmented experiment S^* with $T^* = T \times A$ and $\pi^*(t, a|\theta) = \nu(t, a|\theta)$.

Example: Bank Run

- A bank is solvent or insolvent:

$$\Theta = \{\theta_I, \theta_S\}$$

- Each event is equally likely:

$$\psi(\cdot) = \left(\frac{1}{2}, \frac{1}{2}\right)$$

- running gives payoff 0
- not running gives payoff -1 if insolvent, y if solvent:

$$0 < y < 1$$

- $G = (A, u)$ with $A = \{a_r, a_n\}$ and u given by

	θ_S	θ_I
a_r	0	0^*
a_n	y^*	-1

Bank Run: Common Prior Only

- suppose we have the prior information only - the null information structure:

$$S_0 = (T_0, \pi), \quad T_0 = \{t_0\}$$

- parameterized consistent outcomes:

$\mu(a,\theta)$	θ_S	θ_I
a_r	$\frac{1}{2} \cdot \rho_S$	$\frac{1}{2} \cdot \rho_I$
a_n	$\frac{1}{2} \cdot (1 - \rho_S)$	$\frac{1}{2} \cdot (1 - \rho_I)$

- ρ_S : (conditional) probability of running if solvent
- ρ_I : (conditional) probability of running if insolvent

- agent may not necessarily know state θ but makes choices according to $\mu(\cdot)$
- if "advised" to run, run has to be a best response:

$$\begin{aligned} 0 &\geq \rho_S y - \rho_I \Leftrightarrow \\ \rho_I &\geq \rho_S y \end{aligned}$$

- if "advised" not to run, not run has to be a best response

$$\begin{aligned} (1-\rho_S)\, y - (1-\rho_I) &\geq 0 \Leftrightarrow \\ \rho_I &\geq (1-y) + \rho_S y \end{aligned}$$

- here, not to run provides binding constraint:

$$\rho_I \geq (1-y) + \rho_S y$$

- never to run, $\rho_I = 0, \rho_S = 0$, cannot be a BCE

- set of BCE described by (ρ_I, ρ_S)

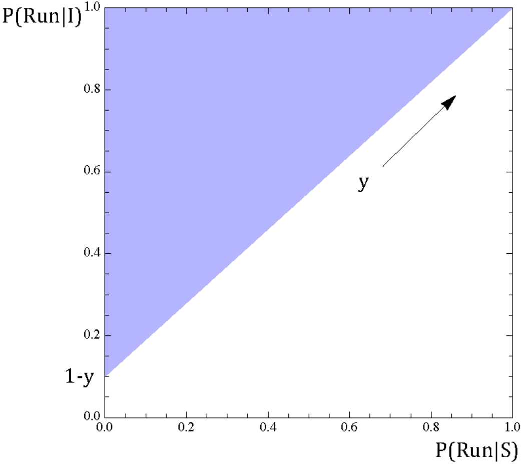

- never to run, $\rho_I = 0, \rho_S = 0$, is not be a BCE

- BCE minimizing the probability of runs has:

$$\rho_I = 1 - y, \quad \rho_S = 0$$

- Noisy stress test $T = \{t_I, t_S\}$ implements BNE:

$$\begin{array}{ccc} \pi(t \mid \theta) & \theta_I & \theta_S \\ t_I & 1-y & 0 \\ t_S & y & 1 \end{array}$$

- the bank is said to be healthy if it is solvent (always) **and** if it is insolvent (sometimes)
- solvent and unsolvent banks are bundled

Bank Run: Positive Information

- suppose player observes conditionally independent private binary signal of the state with accuracy:

$$q > \frac{1}{2}$$

- $S = (T, \pi)$ where $T = \{t_S, t_I\}$:

π	θ_S	θ_I
t_S	q	$1 - q$
t_I	$1 - q$	q

- strictly more information than null information $q = \frac{1}{2}$

Bank Run: Additional Obedience Constraints

- possibly distinct running probabilities for different types:

$$\left(\rho_I^I, \rho_S^I\right), \quad \left(\rho_I^S, \rho_S^S\right)$$

- conditional obedience constraints, say for t_S :

$$\begin{aligned} a_r &: \quad 0 \geq q\rho_S^S y - (1-q)\,\rho_I^S \\ a_n &: \quad q\left(1-\rho_S^S\right) y - (1-q)\left(1-\rho_I^S\right) \geq 0 \end{aligned}$$

or

$$\begin{aligned} a_r &: \quad \rho_I^S \geq \frac{q}{1-q} y \rho_S^S \\ a_n &: \quad \rho_I^S \geq 1 - \frac{q}{1-q} y + \frac{q}{1-q} \rho_S^S y \end{aligned}$$

Bank Run: Equilibrium Set

- set of BCE described by (ρ_I, ρ_S)

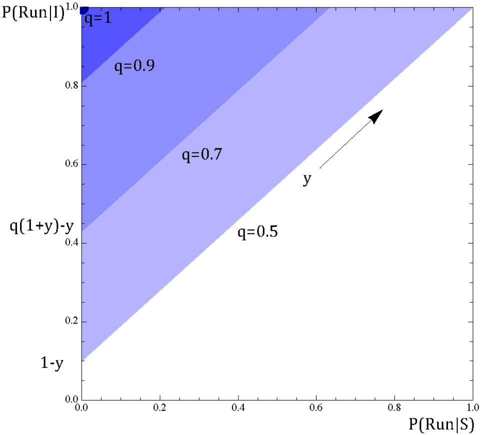

- $\rho_I = 1, \rho_S = 0$, is complete information BCE

Incentive Compatibility Ordering

- Write $BCE(G, S)$ for the set of BCE outcomes of (G, S)

Definition
Experiment S is more *incentive constrained* than experiment S' if, for all decision problems G,

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

- Note that "more incentive constrained" corresponds, intuitively, to having more information

- recall: outcome $\mu \in \Delta(A \times \Theta)$ is feasible for (G, S) if there exists an attainable and consistent $\nu \in \Delta(A \times T \times \Theta)$ giving μ
- recall: $F(G, S)$ is the set of feasible outcomes of (G, S)

Definition

Experiment S is more *permissive* than experiment S' if, for all decision problems G,

$$F(G, S) \supseteq F(G, S')$$

- $\Leftrightarrow$ the set of feasible outcomes is larger in S than S' is Blackwell's definition of "more informative"
- $\Leftrightarrow$ the largest value attained by some feasible outcome in S is larger than in S' (economists' definition of "more valuable")

Back to the Example

- suppose we have the prior information only - the null information structure:

$$S_0 = (T_0, \pi), \quad T_0 = \{t_0\}$$

- feasible outcomes $\mu(a, \theta)$ can be described by (ρ_I, ρ_S):

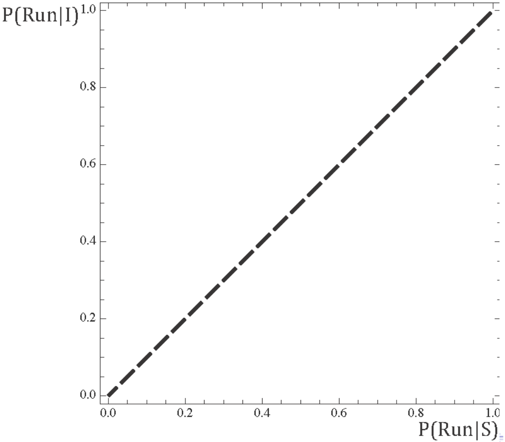

- suppose player observes conditionally independent private binary signal of the state with accuracy $q \geq \frac{1}{2}$:
- feasible outcomes $\mu(a, \theta)$ can be described by (ρ_I, ρ_S):

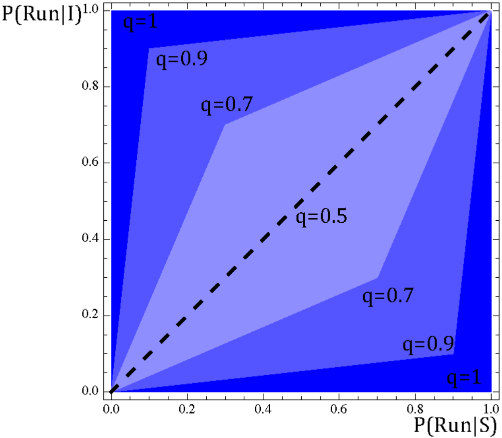

- Experiment S is sufficient for experiment S' if there exists a combination S^* of S and S' such that

$$\Pr\left(\theta|t,t'\right)=\frac{\psi\left(\theta\right)\pi^*\left(t,t'|\theta\right)}{\sum\limits_{\theta'\in\Theta}\psi\left(\theta'\right)\pi^*\left(t,t'|\theta'\right)}$$

is independent of t' :

$$\Pr\left(\theta|t,t'\right)=\Pr\left(\theta|t\right)$$

- statistic S is *sufficient* for (S,S') : to compute the posterior distribution of θ, we only need to know the value of the statistic S rather than the value of (S,S')

Theorem

The following are equivalent:

1. *Experiment S is more permissive than experiment S' (feasibility ordering)*
2. *Experiment S is sufficient for S' (statistical ordering)*
3. *Experiment S is more incentive constrained than S' (incentive ordering)*

Proof of Blackwell's Theorem Plus

- Equivalence of (1) "more permissive" and (2) "sufficient for" is due to Blackwell (in finite case, separating hyperplane argument)
- (2) "sufficient for" $\Rightarrow$ (3) "more incentive constrained":

1. take the stochastic transformation ϕ that maps S into S'
2. take any BCE $\nu \in \Delta(A \times T \times \Theta)$ of (G, S) and use ϕ to construct $\nu' \in \Delta(A \times T' \times \Theta)$
3. show that ν' is a BCE of (G, S')

Proof of Blackwell's Theorem Plus

- (3) "more incentive constrained" $\Rightarrow$ (2) "more permissive" by contrapositive
- suppose S is not more permissive than S'
- so $F(G,S) \not\supseteq F(G,S')$ for some G
- so there exists G' and $\nu' \in \Delta(A \times T' \times \Theta)$ which is feasible for (G',S') and gives outcome $\mu \in \Delta(A \times \Theta)$, with μ not feasible for (G,S)
- can choose G' so that the value V of ν' in (G',S') is V and the value every feasible μ of (G',S) is less than V
- now every there all BCE of (G',S') will have value at least V and some BCE of (G',S) will have value strictly less than V
- so $BCE(G',S) \subsetneq BCE(G',S')$

- players $i = 1, ..., I$
- (payoff) states Θ
- actions $(A_i)_{i=1}^{I}$
- utility functions $(u_i)_{i=1}^{I}$, each $u_i : A \times \Theta \to \mathbb{R}$
- state distribution $\psi \in \Delta(\Theta)$
- $G = \left((A_i, u_i)_{i=1}^{I}, \psi \right)$
- "decision problem" in the one player case

- signals (types) $(T_i)_{i=1}^{I}$
- signal distribution $\pi : \Theta \to \Delta(T_1 \times T_2 \times ... \times T_I)$
- $S = \left((T_i)_{i=1}^{I}, \pi\right)$
- "experiment" in the one player case

- fix (G, S).
- which outcomes can arise in Bayes Nash Equilibrium in (G, S^*) where S^* is an information augmentation of S?

Definition

Distribution $\nu \in \Delta(A \times T \times \Theta)$ is a *Bayes correlated equilibrium* (BCE) of (G, S) if it is obedient and consistent.

- Outcome $\mu \in \Delta(A \times \Theta)$ is BCE outcome *of* (G, S) if a BCE $\nu \in \Delta(A \times T \times \Theta)$ gives outcome μ.

Theorem

An outcome μ is a BCE outcome of (G, S) if and only if there is an information augmentation S^* of S such that μ is a BNE outcome for (G, S^*)

Proof Idea:

- $(\Leftarrow)$ feasibility of S^* has "more" obedience constraints than S
- $(\Rightarrow)$ let ν be BCE of (G, S) supporting μ and consider augmented information structure S^* with $T^* = T \times A$ and

$$\pi^* (t, a|\theta) = \nu (t, a|\theta) .$$

Definition

Information structure S is individually sufficient for information structure S' if there exists a combination S^* of S and S' such that for all $\psi(\theta)$ for each i,

$$\Pr\left(\theta \,\middle|\, t_i, t_i'\right) = \frac{\sum_{t_{-i}\in T_{-i}} \sum_{t'_{-i}\in T'_{-i}} \psi(\theta)\, \pi^*\left(t, t'|\theta\right)}{\sum_{\theta'\in\Theta} \sum_{t_{-i}\in T_{-i}} \sum_{t'_{-i}\in T'_{-i}} \psi\left(\theta'\right) \pi^*\left(t, t'|\theta'\right)}$$

is independent of t_i':

$$\Pr\left(\theta \,\middle|\, t_i, t_i'\right) = \Pr\left(\theta \,\middle|\, t_i\right).$$

Blackwell's Theorem Plus for Many Players

Theorem

The following are equivalent:

1. *Information structure S is more permissive than S' (feasibility ordering)*
2. *Information structure S is individually sufficient for S' (statistical ordering)*
3. *Information structure S is more incentive constrained than S' (incentive ordering)*

Proof of Blackwell's Theorem Plus

- Equivalence of (1) and (2) can be proved from results of LRS discussed below
- (2) "sufficient for" $\Rightarrow$ (3) "more incentive constrained" works as in the single player case

1. take the stochastic transformation ϕ that maps S into S'
2. take any BCE $\nu \in \Delta(A \times T \times \Theta)$ of (G, S) and use ϕ to construct $\nu' \in \Delta(A \times T' \times \Theta)$
3. show that ν' is a BCE of (G, S')

- need a new argument to show (3) $\Rightarrow$ (2)

New Argument: Game of Belief Elicitation

- Suppose that S is more incentive constrained than S'
- Consider a finite game $G_{\varepsilon,S}$ (like DFM 2007) where players report their beliefs and higher order beliefs, with an ε-grid of possible reports but including exact HOB types in S
- Consider BCE ν^* of $(G_{\varepsilon,S}, S)$ (for ALL $\varepsilon > 0$) where beliefs and higher order beliefs are truthfully reported
- For each $\varepsilon > 0$, there must exist a BCE ν_ε of $(G_{\varepsilon,S}, S')$ giving same outcome
- In a sequence of such BCE (as $\varepsilon \to 0$), players' types in S' cannot be giving them information about others' actions (i.e., types in S) and states.
- So S is sufficient for S'

Alternative Characterization of Sufficiency

- Say two information structures are higher order belief equivalent if and only if they give rise to the same probability distribution over Mertens-Zamir hierarchies of higher order beliefs
- Information structure S is sufficient for information structure S' if and only if S is higher order belief to a combination of S and S'
- Liu (2011) suggests an interpretation of the equivalence: in his language, information structure S is sufficient for information structure S' if there exists a combination of S and S' such that the implied stochastic mapping from $\Theta \times T$ to $\Delta(T')$ is a "correlating device"

Feasibility Conditions and Incentive Constraints

- Beyond consistency, Bayes correlated equilibrium imposed only an incentive compatibility condition (obedience)
- interest in Bayes correlated equilibrium due to link with robustness predictions.
- it should not be understood as a solution concept for a fixed incomplete information game, which is the focus of the existing literature.
- Much of the literature can be understood by thinking about additional feasibility conditions to reflect the original information structure.

Other Definitions: Stronger Feasibility Constraints

- Two crucial and seemingly very modest feasibility restrictions are:

1. join feasibility: equilibrium play cannot depend on things no one knows
2. belief invariance: information structure cannot change, so players cannot learn about the state and others' types form their action recommendations

- Belief invariant Bayesian solution (BCE + join feasibility + belief invariance) has been much discussed
- Looking at correlated equilibria of agent normal form imposes an implicit feasibility condition ("uninformed mediator") which is strictly stronger than the combination of join feasibility and belief invariance, called "agent normal form correlated equilibrium".

Other Definitions: Stronger Incentive Constraints

1. Mediator can make recommendations contingent on players' types only if they have an incentive to truthfully report them. Called "communication equilibrium".
2. Looking at correlated equilibrium of strategic form rules out type dependent recommendations, called "strategic form correlated equilibrium"

Generalizing Blackwell's Theorem

- we saw - in both the one and the many player case - that "more information" helps by relaxing feasibility constraints and hurts by imposing incentive constraints
- in the one player case, if we focus on the set of attainable payoffs (like Blackwell) or the best payoffs (as in the economics literature), incentive constraints cannot hurt
- In the many player case they can...
- Bayes correlated equilibrium imposes no feasibility constraints, just incentive constraints...
- it leads to clean generalization of Blackwell's theorem relating "incentive constrained" ordering and sufficiency

Conclusion

- a permissive notion of correlated equilibrium in games of incomplete information: Bayes correlated equilibrium
- BCE renders robust predicition operational, embodies concern for robustness to strategic information
- leads to a multi-agent generalization of Blackwell's single agent information ordering