Strategic Distinguishability and Robust Virtual Implementation

Dirk Bergemann and Stephen Morris
Brown University

June 2008

Interdependent Preferences

- preferences are frequently assumed to be interdependent for informational or psychological reasons
- what are the observable implications of interdependent preferences
- "revealed preference" is well-developed theory to understand individual choice with independent preferences
- an approach to "strategic revealed preference" is suggested to understand interdependent preferences

Strategic Distinguishability

- each agent's preference depends on the "payoff types" of all agents
- two types of an agent are "strategically indistinguishable" if in every game there exists some *common* action which each type might rationally choose given some beliefs and higher-order beliefs
- two types of an agent are "strategically distinguishable" if there exists a game such that those types must rationally choose different messages whatever their beliefs and higher-order beliefs
- we characterize strategic distinguishability for general environments:
  - basic idea: types are strategically distinguishable if there is not too much interdependence of preferences

Strategic Revealed Preference

- strategically distinguishable types reveal information through choice
- information revelation in mechanism design:
  in a specific game do different types act differently in specific equilibrium?
  - specific game: direct mechanism of given social choice function
  - specific equilibrium: truthtelling
- in contrast, here we ask does there *exist* a game such that strategically distinguishable types always act differently

Maximally Revealing Mechanism

- construction of a canonical game to identify strategically distinguishable types
  - for all beliefs and higher order beliefs
  - maximally revealing mechanism

Robust Virtual Implementation

- social choice function maps payoff type profiles to outcomes
- "robust implementation": there exists a mechanism such that every equilibrium delivers the socially desired outcome whatever players' beliefs and higher order beliefs about others' types
- "robust virtual implementation": there exists a mechanism such that every equilibrium delivers the socially desired outcome with probability at least $1 - \varepsilon$ whatever players' beliefs and higher order beliefs about others' types

Robust Virtual Implementation

- necessary conditions:
    1. ex post incentive compatibility
    2. robust measurability: strategically indistiguishable always receive same allocation
- sufficiency: extending an argument of Abreu-Matsushima 1992
    - virtual (instead of exact) implementation: specific social choice function is chosen with probability $1-\varepsilon$ (rather than 1)
    - insert maximally revealing mechanism with probability ε

Outline

1. Introduction
2. Auction Example
3. Environment and Solution Concepts
4. Strategic Distinguishability: A Characterization Result
5. Robust Virtual Implementation

Auction Example

- I agents with quasilinear utility
- agent i has type $\theta_i \in \Theta_i = [0, 1]$
- agent i's valuation of a single object is

$$v_i(\theta) = \theta_i + \gamma \sum_{j \neq i} \theta_j$$

- $\gamma \in \mathbb{R}$ measures the intensity of the interdependence
- $\gamma = 0$: private values, no interdependence

Interdependence and Strategic Distinguishability

- with $v_i(\theta) = \theta_i + \gamma \sum_{j \neq i} \theta_j$ suppose:
  1. $\gamma \geq \frac{1}{I-1}$
  2. every low θ_i valuation agent was convinced that other agents were high θ_j agents, and vica versa
  3. in particular, each payoff θ_i is convinced that his opponents are types
  $$\theta_j = \frac{1}{2} + \frac{1}{\gamma(I-1)}\left(\frac{1}{2} - \theta_i\right)$$
- then common knowledge that everyone's valuation of the object is $\frac{1}{2}(1 + \gamma(I-1))$
- thus all types strategically indistinguishable if $\gamma \geq \frac{1}{I-1}$
- we will later establish that all types are strategically distinguishable in this example if $\gamma < \frac{1}{I-1}$

Second Price Auction

- private values $\gamma = 0$ so $v_i = \theta_i$
- second price sealed bid auction
  - object goes to highest bidder
  - winner pays second highest bid
- truth-telling is a dominant strategy, but there are inefficient equilibria

Approximate Second Price Auction

- with probability $1 - \varepsilon$,
  - allocate object to highest bidder
  - winner pays second highest bid
- for each i, with probability $\frac{\varepsilon b_i}{I}$
  - i gets object
  - pays $\frac{1}{2} b_i$
- truth-telling is a strictly dominant strategy so we can guarantee Robust Virtual Implementation

Modified Second Price Auction

- $\gamma > 0$, $v_i = \theta_i + \gamma \sum_{j \neq i} \theta_j$
- generalized second price sealed bid auction
  - object goes to highest bidder
  - winner pays $\max_{j \neq i} b_j + \gamma \sum_{j \neq i} b_j$
- if $\gamma \leq 1$, truth-telling is a "ex post" equilibrium but there are inefficient ex post equilibria ("ex post incentive compatibility")

Modified Second Price Auction

- with probability $1-\varepsilon$,
  - allocate object to highest bidder i
  - winner pays $\max_{j\neq i} b_j + \gamma \sum_{j\neq i} b_j$
- for each i with probability $\frac{\varepsilon b_i}{I}$,
  - i gets object
  - pays $\frac{1}{2} b_i + \gamma \sum_{j\neq i} b_j$

truth telling is a strict ex post equilibrium

in Auction Example

- if $\gamma \geq \frac{1}{I-1}$, inefficient multiple equilibria in the ε modified second price auction AND ALL OTHER mechanisms
- if $\gamma < \frac{1}{I-1}$, robust virtual implementation can be achieved using the ε modified second price auction

Robust Virtual Implementation Results in General Environments

Necessary and Sufficient Conditions:

1. Ex Post Incentive Compatibility
   - in example, $\gamma \leq 1$
2. "Robust Measurability" or Not Too Much Interdependence
   - in example, $\gamma < \frac{1}{I-1}$

Section 3: ENVIRONMENT AND SOLUTION CONCEPTS

Environment

- I agents
- lottery outcome space $Y = \Delta(X)$, X finite
- finite "payoff" types Θ_i
- vNM utilities: $u_i : Y \times \Theta \to \mathbb{R}$

Mechanism

A mechanism $\mathcal{M}$ is a collection $\left((M_i)_{i=1}^{I}, g\right)$

- each M_i is a finite message set
- outcome function $g : M \to Y$

Rationalizable Messages

- initialize at $S_i^{\mathcal{M},0}(\theta_i) = M_i$, inductive step:
- $S_i^{\mathcal{M},k+1}(\theta_i) =$

$$\left\{ m_i \;\middle|\; \begin{array}{l} \exists\ \mu_i \in \Delta\left(\Theta_{-i} \times M_{-i}\right) \text{ s.t.:} \\ (1)\ \mu_i\left(\theta_{-i}, m_{-i}\right) > 0 \Rightarrow m_{-i} \in S_{-i}^{\mathcal{M},k}\left(\theta_{-i}\right) \\ (2)\ m_i \in \underset{m_i'}{\arg\max} \sum\limits_{\theta_{-i}, m_{-i}} \mu_i\left(\theta_{-i}, m_{-i}\right) u_i\left(g\left(m_i', m_{-i}\right), \theta\right) \end{array} \right\}$$

- limit set

$$S_i^{\mathcal{M}}(\theta_i) = \bigcap_{k \geq 0} S_i^{\mathcal{M},k}(\theta_i).$$

- $S_i^{\mathcal{M}}(\theta_i)$ are *rationalizable actions* of type θ_i in mechanism $\mathcal{M}$

Epistemic Foundations: Framework

- Type Space $\mathcal{T} = \left(T_i, \widehat{\pi}_i, \widehat{\theta}_i\right)_{i=1}^{I}$
  1. T_i countable types of agent i
  2. $\widehat{\pi}_i : T_i \to \Delta\left(T_{-i}\right)$ (belief type component)
  3. $\widehat{\theta}_i : T_i \to \Theta_i$ (payoff type component)
- incomplete information game $\left(\mathcal{T}, \mathcal{M}\right)$
  - i's strategy: $\sigma_i : T_i \to \Delta\left(M_i\right)$
  - strategy profile σ is an equilibrium if $\sigma_i\left(m_i | t_i\right) > 0$ implies m_i is in

$$\arg\max_{m_i'} \sum_{t_{-i}, m_{-i}} \widehat{\pi}_i\left[t_i\right]\left(t_{-i}\right)\left(\prod_{j \neq i} \sigma_j\left(m_j | t_j\right)\right) u_i\left(g\left(m_i', m_{-i}\right), \widehat{\theta}\left(t\right)\right)$$

Epistemic Foundations: Result

PROPOSITION. $m_i \in S_i^{\mathcal{M}}(\theta_i)$ if and only if there exist

1. a type space $\mathcal{T}$,
2. an equilibrium σ of $(\mathcal{T}, \mathcal{M})$ and
3. a type $t_i \in T_i$, such that
   3.1 $\sigma_i(m_i|t_i) > 0$ and
   3.2 $\widehat{\theta}_i(t_i) = \theta_i$.

Brandenburger and Dekel (1987), Battigalli (1996), Bergemann and Morris (2001), Battigalli and Siniscalchi (2003).

Section 4:
STRATEGIC DISTINGUISHABILITY

Strategic Distinguishability

DEFINITION. Types θ_i and θ_i' are strategically indistinguishable if

$$S^{\mathcal{M}}(\theta_i) \cap S^{\mathcal{M}}(\theta_i') \neq \varnothing$$

for every $\mathcal{M}$.

DEFINITION. Types θ_i and θ_i' are strategically equivalent if

$$S^{\mathcal{M}}(\theta_i) = S^{\mathcal{M}}(\theta_i')$$

for every $\mathcal{M}$.

Preference Relation

DEFINITION. R_{θ_i,λ_i} is a preference relation of agent i with payoff type θ_i and conjecture $\lambda_i \in \Delta(\Theta_{-i})$ about types of others:

$$yR_{\theta_i,\lambda_i}y' \Leftrightarrow \sum_{\theta_{-i}\in\Theta_{-i}} \lambda_i(\theta_{-i})\, u_i(y,\theta) \geq \sum_{\theta_{-i}\in\Theta_{-i}} \lambda_i(\theta_{-i})\, u_i(y',\theta)$$

- write $\Psi_i \subseteq \Theta_j$ for subset and $\Psi_{-i} = \{\Psi_j\}_{j\neq i}$ for profile of subsets
- possible preference profiles if i assigns probability 1 to his opponents' types to be $\theta_{-i} \in \Psi_{-i}$:

$$\mathcal{R}_i(\theta_i, \Psi_{-i}) = \{R \in \mathcal{R} \,|\, R = R_{\theta_i,\lambda_i} \text{ for some } \lambda_i \in \Delta(\Psi_{-i})\}$$

Defining Separability

- with:
  $\mathcal{R}_i\left(\theta_i, \Psi_{-i}\right) = \left\{R \in \mathcal{R} \,\middle|\, R = R_{\theta_i, \lambda_i} \text{ for some } \lambda_i \in \Delta\left(\Psi_{-i}\right)\right\}$

DEFINITION. Type set profile Ψ_{-i} separates Ψ_i if

$$\bigcap_{\theta_i \in \Psi_i} \mathcal{R}_i\left(\theta_i, \Psi_{-i}\right) = \varnothing.$$

- Ψ_{-i} separates Ψ_i if whatever realized preference of i, we can rule out at least one possible type of i.

Iterative Definition of Separability

- iterateively delete type sets of i that are separated by some type set profile Ψ_{-i}

$$\begin{aligned}\Xi_i^0 &= 2^{\Theta_i} \\ \Xi_i^{k+1} &= \left\{ \Psi_i \in 2^{\Theta_i} \middle| \Psi_{-i} \text{ doesn't separate } \Psi_i \text{ for some } \Psi_{-i} \in \Xi_{-i}^k \right.\end{aligned}$$

and limit type set profile is

$$\Xi_i^* = \bigcap_{k \geq 0} \Xi_i^k$$

Pairwise Inseparable

DEFINITION. Types θ_i and θ_i' are pairwise inseparable if

$$\left\{\theta_i, \theta_i'\right\} \in \Xi_i^*,$$

and we write $\theta_i \sim \theta_i'$.

- note $\sim$ is reflexive, symmetric, but not necessarily transitive

Back to the Auction Example

- I bidders
- bidder i has type $\theta_i \in \Theta_i = [0, 1]$
- bidder i's valuation is $v_i(\theta, m_i) = \theta_i + \gamma \sum_{j \neq i} \theta_j - m_i$
- set of possible preferences = set of possible valuations

$$V_i(\theta_i, \Psi_{-i}) = \left[\theta_i + \gamma \sum_{j \neq i} \min \Psi_j \quad , \quad \theta_i + \gamma \sum_{j \neq i} \max \Psi_j \right]$$

Separability in the Auction Example I

- now Ψ_{-i} separates Ψ_i if and only if

$$\bigcap_{\theta_i \in \Psi_i} V_i\left(\theta_i, \Psi_{-i}\right) = \varnothing$$

- suppose $\theta_i, \theta_i' \in \Psi_i$ and $\theta_i < \theta_i'$;
- there exist $\lambda_i, \lambda_i' \in \Delta\left(\Psi_{-i}\right)$ such that $R_{\theta_i,\lambda_i} = R_{\theta_i',\lambda_i'}$ iff

$$\theta_i + \gamma \sum_{j \neq i} \max \Psi_j \;\geq\; \theta_i' + \gamma \sum_{j \neq i} \min \Psi_j$$

Separability in the Auction Example II

- Ψ_i is separable given Ψ_{-i} if and only if

$$\max \Psi_i - \min \Psi_i > \gamma \left(\sum_{j \neq i} \max \Psi_j - \min \Psi_j \right)$$

- thus

$$\Xi_i^1 = \{ \Psi_i \, | \max \Psi_i - \min \Psi_i \leq [\gamma (I - 1)] \}$$

and iteratively:

$$\Xi_i^k = \left\{ \Psi_i \, \middle| \max \Psi_i - \min \Psi_i \leq [\gamma (I - 1)]^k \right\}$$

Pairwise Inseparability in the Auction Example

- If $\gamma \geq \frac{1}{I-1}$, all θ_i, θ_i' are pairwise inseparable
- If $\gamma < \frac{1}{I-1}$, $\theta_i \neq \theta_i' \Rightarrow \theta_i$ and θ_i' are pairwise separable
- pairwise separability requires "not too much interdependence"

Fixed Point Characterization

Consider a collection of sets $\Xi = (\Xi_i)_{i=1}^{I}$, each $\Xi_i \subseteq 2^{\Theta_i}$.

DEFINITION. A collection Ξ is mutually inseparable if, for each i and $\Psi_i \in \Xi_i$, there exists $\Psi_{-i} \in \Xi_{-i}$ such that Ψ_{-i} does not separate Ψ_i.

LEMMA. Types θ_i and θ_i' are pairwise inseparable if and only if there exists mutually inseparable Ξ such that $\{\theta_i, \theta_i'\} \subseteq \Psi_i$ for some $\Psi_i \in \Xi_i$.

Strategic Distinguishability

DEFINITION. Types θ_i and θ_i' are strategically indistinguishable if

$$S^{\mathcal{M}}(\theta_i) \cap S^{\mathcal{M}}(\theta_i') \neq \varnothing$$

for every $\mathcal{M}$.

THEOREM 1. Types θ_i and θ_i' are strategically indistinguishable if and only if they are pairwise inseparable.

Sufficiency of Pairwise Separability I

PROPOSITION 1: If θ_i and θ_i' are indistinguishable, then

$$S_i^{\mathcal{M}}(\theta_i) \cap S_i^{\mathcal{M}}(\theta_i') \neq \varnothing$$

in any mechanism $\mathcal{M}$.
Suppose Ξ is mutually inseparable
Fix any finite mechanism.

Sufficiency of Pairwise Separability II

By induction on k, for each k, i and $\Psi_i \in \Xi_i$, there exists a common action $m_i^k(\Psi_i)$ such that $m_i^k(\Psi_i) \in S_i^k(\theta_i)$ for each $\theta_i \in \Psi_i$

1. True by definition for $k = 0$.
2. Suppose true for $k - 1$
   - fix any i and $\Psi_i \in \Xi_i$
   - since Ξ is mutually inseparable, $\exists \Psi_{-i} \in \Xi_{-i}$, R_i and, for each $\theta_i \in \Psi_i$, $\lambda_i^{\theta_i} \in \Delta(\Psi_{-i})$ such that $R_{\theta_i, \lambda_i^{\theta_i}} = R_i$
   - $m_i^k(\Psi_i)$ be any element of the argmax under R_i of $g\left(m_i, m_{-i}^{k-1}(\Psi_{-i})\right)$
   - by construction, $m_i^k(\Psi_i) \in S_i^{\mathcal{M},k}(\theta_i)$ for all $\theta_i \in \Psi_i$.

Necessity of Pairwise Separability

PROPOSITION 2: There exists a mechanism $\mathcal{M}^*$ such that if $\theta_i \nsim \theta_i'$, then

$$S_i^{\mathcal{M}^*}(\theta_i) \cap S_i^{\mathcal{M}^*}(\theta_i') = \varnothing.$$

PROOF: By construction of "maximally revealing mechanism".

Construction of Maximally Revealing Mechanism I

uniform lottery $\bar{y} : \bar{y}(x) \triangleq 1/|X|$

KEY LEMMA:

Type set profile Ψ_{-i} separates Ψ_i iff there exists $\widetilde{y} : \Psi_i \to Y$ such that

$$\sum_{\theta_i \in \Psi_i} \left(\widetilde{y}(\theta_i) - \overline{y}\right) = 0$$

and, for each $\theta_i \in \Psi_i$ and $\lambda_i \in \Delta(\Psi_{-i})$,

$$\widetilde{y}(\theta_i) P_{\theta_i,\lambda_i} \overline{y}.$$

Construction of Maximally Revealing Mechanism II

LEMMA (Morris 1994, Samet 1998): Let $V_1, ..., V_L$ be closed, convex, subsets of the N-dimensional simplex Δ^N. These sets have an empty intersection if and only if there exist $z_1, ..., z_L \in \mathbb{R}^N$ such that

$$\sum_{l=1}^{L} z_l = 0$$

and

$$v \cdot z_l > 0$$

for each $l = 1, ..., L$ and $v \in V_l$.

Key lemma follows from this duality lemma, letting $\Theta_i = \{1, ..., L\}$ and V_l be the set of possible utility weights of type $\theta_i = l$ with any $\lambda_i \in \Delta(\Psi_{-i})$.

Construction of Maximally Revealing Mechanism III

- let $B^Y(\theta_{i,}\lambda_i)$ be the agents most preferred lotteries in the set Y given type θ_i and belief λ_i:

$$B_i^Y(\theta_i,\lambda_i) = \left\{ y \in Y \,\middle|\, yR_{\theta_i,\lambda_i}y' \text{ for all } y' \in Y \right\}$$

TEST SET LEMMA. There exists a finite set $Y^* \subseteq Y$ such that

1. for each i, θ_i and $\lambda_i \in \Delta(\Theta_{-i})$, $B_i^{Y^*}(\theta_i,\lambda_i) \neq Y^*$
2. for each i, Ψ_i and Ψ_{-i}, if Ψ_{-i} separates Ψ_i, then for each $\theta_i \in \Psi_i$ and $\lambda_i \in \Delta(\Psi_{-i})$, there exists $\theta_i' \in \Psi_i$ such that

$$B_i^{Y^*}(\theta_i,\lambda_i) \cap B_i^{Y^*}(\theta_i',\Psi_{-i}) = \varnothing.$$

Mechanism in Words

- each player makes K simultaneous announcements:
  1. an element of test set Y^*
  2. a profile of first round announcements of other players he thinks possible, plus an element of Y^*
  3. a profile of second round announcements of other players he thinks possible, plus an element of Y^*
  4. .....
- all chosen outcomes selected with positive probability, with much higher weight on "earlier" announcements

Mechanism in Formulae

mechanism $\mathcal{M}^{K,\varepsilon} = \left(\left(M_i^K \right)_{i=1}^{I}, g^{K,\varepsilon} \right)$ parameterized by

1. $\varepsilon > 0$
2. integer K

- i's message set is M_i^K where
  - $M_i^0 = \left\{ \overline{m}_i^0 \right\}$
  - $M_i^{k+1} = M_i^k \times M_{-i}^k \times Y^*$
- typical element $m_i^k = \left\{ \overline{m}_i^0, r_i^1, y_i^1, ..., r_i^k, y_i^k \right\}$
- allocation rule:

$$g^{K,\varepsilon}(m) = \overline{y} + \frac{1-\varepsilon^K}{1-\varepsilon} \frac{1}{I} \sum_{k=1}^{K} \varepsilon^{k-1} \sum_{i=1}^{I} \mathbb{I}\left(r_i^k, m_{-i}^{k-1} \right) \left(y_i^k - \overline{y} \right)$$

where

$$\mathbb{I}\left(r_i^k, m_{-i}^{k-1} \right) = \begin{cases} 1, \text{ if } r_i^k = m_{-i}^{k-1} \\ 0, \text{ otherwise} \end{cases}$$

Conclusion of Proof of Proposition 2

1. Let

$$\overline{\Theta}_i^k\left(m_i^k\right)=\overline{\Theta}_i^k\left(\left(m_i^{k-1},r_i^k,y_i^k\right)\right)=\left\{\theta_i\;\middle|\;\begin{array}{l}\theta_i\in\overline{\Theta}_i^{k-1}\left(m_i^{k-1}\right)\\ \overline{\Theta}_{-i}^{k-1}\left(r_i^k\right)\neq\varnothing\\ y_i^k\in B_i\left(\theta_i,\overline{\Theta}_{-i}^{k-1}\left(r_i^k\right)\right)\end{array}\right.$$

2. There exists $\bar{\varepsilon}>0$ such that

$$\left\{\theta_i\in\Theta_i\;\middle|\;m_i^k\in S_i^{\mathcal{M}^{k,\varepsilon}}\left(\theta_i\right)\right\}\subseteq\overline{\Theta}_i^k\left(m_i^k\right)$$

for all $\varepsilon\leq\bar{\varepsilon}$ and $m_i^k\in M_i^k$.

3. There exists $\bar{\varepsilon}>0$ and K such that

$$\left\{\theta_i\in\Theta_i\;\middle|\;m_i^K\in S_i^{\mathcal{M}^{K,\varepsilon}}\left(\theta_i\right)\right\}\in\Xi_i^*$$

for all $\varepsilon\leq\bar{\varepsilon}$ and $m_i^K\in M_i^K$.

Section 5: ROBUST VIRTUAL IMPLEMENTATION

Definitions Reminder

- "implementation": requires ALL equilibria deliver the right outcome, a.k.a. full implementation
- "robust": same mechanism works independent of agents' beliefs and higher order beliefs about the environment
- "virtual": enough if correct outcome arises with probability $1-\varepsilon$

DEFINITION: A social choice function $f : \Theta \to Y$.
Write $\|y - y'\|$ for the Euclidean distance between a pair of lotteries y and y', i.e.,

$$\|y - y'\| = \sqrt{\sum_{x \in X} \left(y\left(x\right) - y'\left(x\right)\right)^2}.$$

DEFINITION: Social choice function f is robustly ε-implementable if there exists a mechanism $\mathcal{M}$ such that

$$m \in S^{\mathcal{M}}\left(\theta\right) \Rightarrow \|g\left(m\right) - f\left(\theta\right)\| \leq \varepsilon.$$

DEFINITION: Social choice function f is robustly virtually implementable if, for every $\varepsilon > 0$, f is robustly ε-implementable.

Result

DEFINITION: Social choice function f satisfies ex post incentive compatibility if, for all i, θ_i, θ_{-i} and θ_i':

$$u_i\left(f\left(\theta_i, \theta_{-i}\right), \left(\theta_i, \theta_{-i}\right)\right) \geq u_i\left(f\left(\theta_i', \theta_{-i}\right), \left(\theta_i, \theta_{-i}\right)\right).$$

DEFINITION: Social choice function f satisfies robust measurability if $\theta_i \sim \theta_i' \Rightarrow f\left(\theta_i, \theta_{-i}\right) = f\left(\theta_i', \theta_{-i}\right), \forall \theta_{-i}$

THEOREM 2. Social choice function f is robustly virtually implementable if and only if f satisfies ex post incentive compatibility and robust measurability.

Abreu-Matsushima (1992) Incomplete Information

- Standard "Bayesian" incomplete information setting, i.e., common knowledge of common prior on type space
- Necessary conditions for virtual implementation
  - Bayesian incentive compatibility
  - Abreu-Matsushima measurability: types are iteratively distinguishable
  - reduces to "value distinction" in private values case

Adding Robustness

- with robustness, full implementation equivalent to belief free version of iterated deletion of strictly dominated strategies
- generalizing Abreu-Matsushima, necessary conditions become:

1. Ex post incentive compatibility (instead of Bayesian IC)
   - Bergemann-Morris "Robust Mechanism Design"
2. robust measurability as belief free version of AM measurability

Intermediate Notions of Robustness

Artemov-Kunimoto-Serrano (2008) consider type space with

- given finite payoff types $\theta_i \in \Theta_i$
- given finite first-order beliefs $q_i(\theta_i | \theta_{-i})$

and general type space T_i is assumed to be consistent with payoff types and first-order beliefs

- in the presence of a type diversity condition, incentive compatibility and AM measurability is necessary and sufficient for robust virtual implementation
- some tension between rich type space and type diversity

Exact Implementation I

following Maskin methods, necessary and sufficient conditions for exact robust implementation - using ANY mechanism: (Bergemann-Morris "Robust Implementation in General Mechanisms" (2008))

1. ex post incentive compatibility
2. "robust monotonicity": not too much interdependence

Exact Implementation II

in large class of economically interesting "monotonic aggregator" environments:
(Bergemann-Morris "Robust Implementation in Direct Mechanisms" (2007))

1. robust monotonicity = robust measurability
2. natural generalization of $\gamma < \frac{1}{I-1}$ condition
3. if robust virtual implementation is possible, it arises in modified direct mechanism

Conclusion

- strategic distinguishability:
  information revelation through choice in some game
- strategic distinguishability = not too much interdependence
- information revelation in maximally revealing mechanism
- virtual implementation via maximally revealing mechanism
- robust virtual implementation leads to sharp possibility but also impossibility results