Bayesian mechanism design revisited

Peter Hammond Conference, University of Warwick

C. d'Aspremont and J. Crémer

CORE UCLouvain and TSE Toulouse

December 2-4, 2022

Prologue

- Peter Hammond retiring? An administrative decision, but we very much hope to continue benefiting from his inspiring work and advice.

Prologue

- Peter Hammond retiring? An administrative decision, but we very much hope to continue benefiting from his inspiring work and advice.
- First time I heard about Peter: He was corresponding with Louis Gevers when we were working on Social Choice with Interpersonal Comparisons (1974?). Peter was then in Australia. Naively, I even thought he would stay there for ever...

Prologue

- Peter Hammond retiring? An administrative decision, but we very much hope to continue benefiting from his inspiring work and advice.
- First time I heard about Peter: He was corresponding with Louis Gevers when we were working on Social Choice with Interpersonal Comparisons (1974?). Peter was then in Australia. Naively, I even thought he would stay there for ever...
- But my intuition was wrong: Peter has had an impressive international career, with permanent positions in several prestigious universities (Essex, Stanford, Warwick), and visiting positions in many places around the world (CORE-Louvain-la-Neuve, EUI-Florence, IAS-Jerusalem, IFS-Kiel, IER-Hitotsubashi and many others). Happily, our paths crossed many times at some of these places.

Prologue

- Peter Hammond retiring? An administrative decision, but we very much hope to continue benefiting from his inspiring work and advice.
- First time I heard about Peter: He was corresponding with Louis Gevers when we were working on Social Choice with Interpersonal Comparisons (1974?). Peter was then in Australia. Naively, I even thought he would stay there for ever...
- But my intuition was wrong: Peter has had an impressive international career, with permanent positions in several prestigious universities (Essex, Stanford, Warwick), and visiting positions in many places around the world (CORE-Louvain-la-Neuve, EUI-Florence, IAS-Jerusalem, IFS-Kiel, IER-Hitotsubashi and many others). Happily, our paths crossed many times at some of these places.
- Peter has received Honorary doctorates from the University of Kiel and the University of Oslo.

Prologue

- Peter Hammond retiring? An administrative decision, but we very much hope to continue benefiting from his inspiring work and advice.
- First time I heard about Peter: He was corresponding with Louis Gevers when we were working on Social Choice with Interpersonal Comparisons (1974?). Peter was then in Australia. Naively, I even thought he would stay there for ever...
- But my intuition was wrong: Peter has had an impressive international career, with permanent positions in several prestigious universities (Essex, Stanford, Warwick), and visiting positions in many places around the world (CORE-Louvain-la-Neuve, EUI-Florence, IAS-Jerusalem, IFS-Kiel, IER-Hitotsubashi and many others). Happily, our paths crossed many times at some of these places.
- Peter has received Honorary doctorates from the University of Kiel and the University of Oslo.
- Most impressive of all is the variety of topics to which he has fundamentally contributed.

Prologue: Peter Hammond's list of topics

1. Growth and Exhaustible Resources
2. Rational Individual Choice and Consequentialism
3. Game Theory and Consequentialism
4. Consequentialist Social Choice and Utilitarian Ethical Theory
5. Social Choice: General
6. Social Choice with Interpersonal Comparisons
7. Social Choice with Individual and Group Rights
8. Distributional Objectives in Welfare Economics
9. General Equilibrium Theory and Market Efficiency
10. Gains from Trade and Migration
11. Widespread Externalities and the f-Core
12. Equilibrium in Incomplete Markets
13. **Private Information and Incentive Constraints**
14. Cost-Benefit Analysis, Policy Reform, and Welfare Measurement
15. Welfare, Information and Uncertainty
16. Miscellaneous in Welfare Economics and Ethics
17. Continuum of Random Variable

Prologue: Topic 13: Private information and Incentive Constraints: some articles

- “Symposium on Incentive Compatibility: Introduction,” Review of Economic Studies 46 (1979), 181–184.
    - “The Implementation of Social Choice Rules:Some General Results on Incentive Compatibility,” 185–216 (with P. Dasgupta, E. Maskin)
    - “Straightforward Individual Incentive Compatibility in Large Economies,” 263-282.
- “On Imperfect Information and Optimal Pollution Control,” Review of Economic Studies 47 (1980), 857–860.(with P. Dasgupta, E. Maskin)
- “Markets as Constraints: Multilateral Incentive Compatibility in Continuum Economies,” Review of Economic Studies 54 (1987), 399–412
- “Incentives and Allocation Mechanisms,” in R. van der Ploeg (ed.) Advanced Lectures in Quantitative Economics (Academic Press, 1990)
- “**A Revelation Principle for (Boundedly) Bayesian Rationalizable Strategies**,” in R.P. Gilles and P.H.M. Ruys (eds.), Imperfections and Behavior in Economic Organizations (Kluwer, 1994), 39–70.

Prologue: The robustness of the Revelation Principle

- Peter's paper extends the Harsanyi (1967/68) argument (and Mertens-Zamir1985 construction) about players' types.
  - A player's type should not only contains a description of his beliefs and his payoffs. It should also contain a description of his strategic behaviour (to focus on one equilibrium).
  - The paper shows how to construct an equivalent direct mechanism (with truth-telling being a best-response). This construction depends upon *external prior beliefs* over the set of agents' Bayesian rationalizable strategies in the game form
  - In the bounded rationality case: "players are assumed to use models no more complicated than those in which they can solve the appropriate expected utility maximization problem.").
- Two special "simple" cases are also developed.
  - Dominant strategy incentive constraints: the outcome (and the direct mechanism) are then independent of the beliefs and behaviour types.
  - Reduced Bayesian model concentrating on payoff-types: the revelation principle remains valid (but external beliefs still play a role). Our work is within this reduced model.

Outline

- Introduction
- The framework
- Bayesian implementation
- The No-freeness case
- Introducing Freeness
- Concluding Remarks

Introduction

Introduction: Mechanism design

There are multiple applications and the mechanism designer can be socially- or self-interested.

- An important application is the Pareto-efficent funding of the production of a public good (or the reduction of a public bad, e.g. pollution) where the mechanism designer objective is to maximize the total welfare of the consumers. Their "willingness-to-pay" for the public good are private information.

Introduction: Mechanism design

There are multiple applications and the mechanism designer can be socially- or self-interested.

- An important application is the Pareto-efficent funding of the production of a public good (or the reduction of a public bad, e.g. pollution) where the mechanism designer objective is to maximize the total welfare of the consumers. Their "willingness-to-pay" for the public good are private information.
- Another important application is the selling of an object through an auction. The mechanism designer may be the seller who wants to maximize his benefit. The "values" of the auctioned object for the buyers are private information.

Introduction: two approaches

In mechanisms design, two approaches have been used: The Ex-post approach and the Bayesian approach

- (i) *Ex-post incentive compatibility*: Truthtelling is a Nash equilibrium whatever the true types of all agents (called an expost equilibrium). Ex-post incentive compatibility is primary due to Hurwicz (1972). With private values it is equivalent to Dominant-strategy incentive compatibility (Groves1973, Green-Laffont1976).

Introduction: two approaches

In mechanisms design, two approaches have been used: The Ex-post approach and the Bayesian approach

- (i) *Ex-post incentive compatibility*: Truthtelling is a Nash equilibrium whatever the true types of all agents (called an expost equilibrium). Ex-post incentive compatibility is primary due to Hurwicz (1972). With private values it is equivalent to Dominant-strategy incentive compatibility (Groves1973, Green-Laffont1976).
- When side-payments are allowed, Vickrey-Clarke-Groves (Vickrey 1961, Clarke 1971, and Groves 1973) mechanisms ensure both Dominant-strategy IC and efficiency of the decision rule (total welfare is maximised, but the budget may not be balanced)

Introduction: two approaches

In mechanisms design, two approaches have been used: The Ex-post approach and the Bayesian approach

- (i) *Ex-post incentive compatibility*: Truthtelling is a Nash equilibrium whatever the true types of all agents (called an expost equilibrium). Ex-post incentive compatibility is primary due to Hurwicz (1972). With private values it is equivalent to Dominant-strategy incentive compatibility (Groves1973, Green-Laffont1976).
- When side-payments are allowed, Vickrey-Clarke-Groves (Vickrey 1961, Clarke 1971, and Groves 1973) mechanisms ensure both Dominant-strategy IC and efficiency of the decision rule (total welfare is maximised, but the budget may not be balanced)
- (ii) *Bayesian incentive compatibility* (BIC) Truthtelling is a Bayesian equilibrium (Harsanyi 1967/68): In the Bayesian approach the incentive compatibility requirement is weaker. It was introduced in Arrow (1979) and AGV79 to solve the balanced budget problem (with imposed particpation). Assumptions on the beliefs have to be imposed.

Introduction: Bayesian approach

- Recently in the field of **algorithmic mechanism design** relaxing ex post IC to BIC has also proved useful. It helps to circumvent computational intractability (using approximation algorithms) (Hartline-Lucier, 2018)

Introduction: Bayesian approach

- Recently in the field of **algorithmic mechanism design** relaxing ex post IC to BIC has also proved useful. It helps to circumvent computational intractability (using approximation algorithms) (Hartline-Lucier, 2018)
- This field is mostly motivated by Internet applications where procedures are highly frequently repeated and distributions can be estimated: *e.g.* online auctions (eBay), advertising auctions (Google, Yahoo!, MSN).

Introduction: Bayesian approach

- Recently in the field of **algorithmic mechanism design** relaxing ex post IC to BIC has also proved useful. It helps to circumvent computational intractability (using approximation algorithms) (Hartline-Lucier, 2018)
- This field is mostly motivated by Internet applications where procedures are highly frequently repeated and distributions can be estimated: *e.g.* online auctions (eBay), advertising auctions (Google, Yahoo!, MSN).
- As stressed by Albert-Conitzer-Lopomo-Stone (2018) using **automated mechanism design** techniques: in settings where valuations are correlated, stronger results are possible. "If a mechanism designer intends to maximally exploit a correlated valuations setting, she must use information about the distribution".

Introduction: Two opposite kinds of assumptions on the beliefs

Both in the general collective decision context and in auction design, results on BIC have been derived using two opposite kinds of assumptions on the distribution of agent types.

- **In the first kind**, it was assumed that the distribution satisfies an *independence condition* implying what we call "*full-freeness*" and meaning that each agent's beliefs are common knowledge.

Introduction: Two opposite kinds of assumptions on the beliefs

Both in the general collective decision context and in auction design, results on BIC have been derived using two opposite kinds of assumptions on the distribution of agent types.

- **In the first kind**, it was assumed that the distribution satisfies an *independence condition* implying what we call "*full-freeness*" and meaning that each agent's beliefs are common knowledge.
- Under independence finding a Pareto optimal (efficient and budget-balanced) BIC mechanism that satisfies voluntary participation (Interim Individual Rationality) seemed an impossibility after Myerson & Satterthwaite (1983) result (for two agents: one buyer, one seller).

Introduction: Two opposite kinds of assumptions on the beliefs

Both in the general collective decision context and in auction design, results on BIC have been derived using two opposite kinds of assumptions on the distribution of agent types.

- **In the first kind**, it was assumed that the distribution satisfies an *independence condition* implying what we call "*full-freeness*" and meaning that each agent's beliefs are common knowledge.
- Under independence finding a Pareto optimal (efficient and budget-balanced) BIC mechanism that satisfies voluntary participation (Interim Individual Rationality) seemed an impossibility after Myerson & Satterthwaite (1983) result (for two agents: one buyer, one seller).
- However Makowski and Mezzetti (1994), have derived an interesting (but restrictive) necessary and sufficient condition to get existence.

Introduction: no-freeness

The second kind of conditions that have been imposed on the distribution of types is at the other extreme.

- It implies a "*no-freeness*" assumption: for any agent, different types mean different beliefs *i.e.* “Beliefs Determine Preference” in Heifetz and Neeman (2006) or "Beliefs announcement" in Johnson-Pratt-Zeckhauser (1990)

Introduction: no-freeness

The second kind of conditions that have been imposed on the distribution of types is at the other extreme.

- It implies a "*no-freeness*" assumption: for any agent, different types mean different beliefs *i.e.* “Beliefs Determine Preference” in Heifetz and Neeman (2006) or "Beliefs announcement" in Johnson-Pratt-Zeckhauser (1990)
- In auction theory, a well-known assumption of correlation of buyers' valuations is the *Crémer-McLean condition*. It implies no-freeness and allows the seller to extract the whole surplus.

Introduction: no-freeness

The second kind of conditions that have been imposed on the distribution of types is at the other extreme.

- It implies a "*no-freeness*" assumption: for any agent, different types mean different beliefs *i.e.* “Beliefs Determine Preference” in Heifetz and Neeman (2006) or "Beliefs announcement" in Johnson-Pratt-Zeckhauser (1990)
- In auction theory, a well-known assumption of correlation of buyers' valuations is the *Crémer-McLean condition*. It implies no-freeness and allows the seller to extract the whole surplus.
- In the collective decision context, conditions implying no-freeness were proposed by Matsushima (2007) and Kosenok and Severinov (2008) to implement under IIR any public decision mechanism (with more than two agents). In Kosenok and Severinov the conditions are necessary and sufficient and full surplus-sharing is obtained.

Introduction: independence classes

- The main purpose of this paper is to provide results on intermediate cases, in-between full-freeness and no-freeness, thus generalising the preceding results

Introduction: independence classes

- The main purpose of this paper is to provide results on intermediate cases, in-between full-freeness and no-freeness, thus generalising the preceding results
- These in-between results are based on the notion of "*free belief equivalence classes*" associated to an agent beliefs (on the other agents types conditional on its own type). It is a partition of his types such that, within each equivalence class, these conditional beliefs are identical.

- The main purpose of this paper is to provide results on intermediate cases, in-between full-freeness and no-freeness, thus generalising the preceding results
- These in-between results are based on the notion of "*free belief equivalence classes*" associated to an agent beliefs (on the other agents types conditional on its own type). It is a partition of his types such that, within each equivalence class, these conditional beliefs are identical.
- In the full-freeness case, each agent's partition of types has a single element (the set of all his types) and, in the no-freeness case, each element of an agent's partition contains a single type.

Introduction: independence classes

- The main purpose of this paper is to provide results on intermediate cases, in-between full-freeness and no-freeness, thus generalising the preceding results
- These in-between results are based on the notion of "*free belief equivalence classes*" associated to an agent beliefs (on the other agents types conditional on its own type). It is a partition of his types such that, within each equivalence class, these conditional beliefs are identical.
- In the full-freeness case, each agent's partition of types has a single element (the set of all his types) and, in the no-freeness case, each element of an agent's partition contains a single type.
- We will also introduce a weak notion of efficiency - "Independence Class efficiency" - which requires the outcome to be efficient within each equivalence class for the given decision rule.

Introduction: automated mechanism design

- Albert, Connitzer and Lopomo (2015), using an *automated mechanism design* approach to auction design (with a single buyer responding to an external signal), have a similar structure.

Introduction: automated mechanism design

- Albert, Connitzer and Lopomo (2015), using an *automated mechanism design* approach to auction design (with a single buyer responding to an external signal), have a similar structure.
- They introduce a partition of the buyer types so that within each element of the partition the conditional distribution over the external signal remains identical.

Introduction: automated mechanism design

- Albert, Connitzer and Lopomo (2015), using an *automated mechanism design* approach to auction design (with a single buyer responding to an external signal), have a similar structure.
- They introduce a partition of the buyer types so that within each element of the partition the conditional distribution over the external signal remains identical.
- As the size of the equivalent classes increases the analogue to the Crémer-McLean condition is weakened and the optimal revenue of the seller decreases.

The Framework

The basic framework

Consider a collective decision mechanism to which participate $n \geq 3$ agents belonging to the set $\mathcal{N} = \{1, 2, \ldots, n\}$ with a compact set $\mathcal{X}$ of outcomes.

- The private information of agent $i \in \mathcal{N}$ is represented by his *type* α_i which belongs to a finite set $\mathcal{A}_i$.

The basic framework

Consider a collective decision mechanism to which participate $n \geq 3$ agents belonging to the set $\mathcal{N} = \{1, 2, \ldots, n\}$ with a compact set $\mathcal{X}$ of outcomes.

- The private information of agent $i \in \mathcal{N}$ is represented by his *type* α_i which belongs to a finite set $\mathcal{A}_i$.
- The *utility* of agents is assumed to be fully transferable: the payoff of agent i is

$$u_i(x; \alpha_i) + t_i.$$

with t_i a *monetary transfer* (which could be negative).

The basic framework

Consider a collective decision mechanism to which participate $n \geq 3$ agents belonging to the set $\mathcal{N} = \{1, 2, \dots, n\}$ with a compact set $\mathcal{X}$ of outcomes.

- The private information of agent $i \in \mathcal{N}$ is represented by his *type* α_i which belongs to a finite set $\mathcal{A}_i$.
- The *utility* of agents is assumed to be fully transferable: the payoff of agent i is
$$u_i(x; \alpha_i) + t_i.$$
with t_i a *monetary transfer* (which could be negative).
- The type of agent i affects his utility, but may also affect his *beliefs* about the types of the other agents: These conditional beliefs are represented by a probability distribution $p_i(\alpha_{-i} \mid \alpha_i)$ over $\mathcal{A}_{-i} = \prod_{j \in \mathcal{N}_{-i}} \mathcal{A}_j$.

The basic framework

Consider a collective decision mechanism to which participate $n \geq 3$ agents belonging to the set $\mathcal{N} = \{1, 2, \ldots, n\}$ with a compact set $\mathcal{X}$ of outcomes.

- The private information of agent $i \in \mathcal{N}$ is represented by his *type* α_i which belongs to a finite set $\mathcal{A}_i$.
- The *utility* of agents is assumed to be fully transferable: the payoff of agent i is

$$u_i(x; \alpha_i) + t_i.$$

  with t_i a *monetary transfer* (which could be negative).
- The type of agent i affects his utility, but may also affect his *beliefs* about the types of the other agents: These conditional beliefs are represented by a probability distribution $p_i(\alpha_{-i} \mid \alpha_i)$ over $\mathcal{A}_{-i} = \prod_{j \in \mathcal{N}_{-i}} \mathcal{A}_j$.
- We assume *consistency*: $\exists p\,(.)$ such that $p\,(\alpha) = p\,(\alpha_{-i} \mid \alpha_i)\,p\,(\alpha_i)$, for all $\alpha \in \mathcal{A}$, and all $i \in \mathcal{N}$. An *information structure* is denoted $(\mathcal{A}, p)$, where p is the common prior.

Budget-balanced mechanisms with voluntary participation

- We consider *direct revelation mechanisms* (s, t) (each agent supposed to reveal its type), with $s : \mathcal{A} \longrightarrow X$ a public decision rule and $t : \mathcal{A} \longrightarrow \mathbb{R}^n$ a transfer rule.

Budget-balanced mechanisms with voluntary participation

- We consider *direct revelation mechanisms* (s, t) (each agent supposed to reveal its type), with $s : \mathcal{A} \rightarrow X$ a public decision rule and $t : \mathcal{A} \rightarrow \mathbb{R}^n$ a transfer rule.
- We impose that the transfer rule t be *budget-balanced (BB)*,

$$\sum_{i \in N} t_i(\alpha) = 0, \text{ for all } \alpha \in \mathcal{A},$$

Budget-balanced mechanisms with voluntary participation

- We consider *direct revelation mechanisms* (s, t) (each agent supposed to reveal its type), with $s : \mathcal{A} \to X$ a public decision rule and $t : \mathcal{A} \to \mathbb{R}^n$ a transfer rule.
- We impose that the transfer rule t be *budget-balanced (BB)*,

$$\sum_{i \in N} t_i(\alpha) = 0, \text{ for all } \alpha \in \mathcal{A},$$

- Suppose there is a vector of reservation utility levels $\{\bar{U}_i(\alpha_i) : \alpha_i \in \mathcal{A}_i, i \in N\}$. The mechanism (s, t) is *interim individually rational (IIR)* if, for all $\alpha_i \in \mathcal{A}_i$ and all $i \in \mathcal{N}$,

$$\sum_{\alpha_{-i} \in \mathcal{A}_{-i}} \left[u_i\left(s\left(\alpha_i, \alpha_{-i}\right); \alpha_i, \alpha_{-i}\right) + t_i\left(\alpha_i, \alpha_{-i}\right)\right)\right] p\left(\alpha_{-i} \mid \alpha_i\right) \geq \bar{U}_i(\alpha_i).$$

Efficient collective mechanisms

- If the public decision rule s satisfies

$$\sum_{i \in N} u_i(s(\alpha); \alpha) \geq \sum_{i \in N} u_i(x; \alpha), \text{ for all } x \in X \text{ and all } \alpha \in \mathcal{A},$$

it is said to be *(ex post) efficient (EF)*. (*Pareto optimality* $=$ *EF*$+$*BB*.)

Efficient collective mechanisms

- If the public decision rule s satisfies

$$\sum_{i \in N} u_i(s(\alpha); \alpha) \geq \sum_{i \in N} u_i(x; \alpha), \text{ for all } x \in X \text{ and all } \alpha \in \mathcal{A},$$

  it is said to be *(ex post) efficient (EF)*. (*Pareto optimality* $=$ *EF*$+$*BB*.)
- If a decision rule s is *efficient* then, for all $\alpha \in \mathcal{A}$, and all $x \in \{x \in X \mid x = s(\alpha) \text{ for some } \alpha \in \mathcal{A}\}$

$$\sum_{i \in \mathcal{N}} u_i(s(\alpha); \alpha_i) \geq \sum_{i \in \mathcal{N}} u_i(x; \alpha_i).$$

Equivalence classes of free beliefs.

The set of types $\mathcal{A}_i$ of each agent $i \in N$ can be partitioned into *equivalence classes of free beliefs*.

- An equivalence class is a subset Q_i of $\mathcal{A}_i$ such that $p(. \mid \alpha_i) = p(. \mid \widetilde{\alpha}_i)$ whenever $\alpha_i, \widetilde{\alpha}_i \in Q_i$. We denote by $Q_i(\alpha_i)$, the element of the partition to which α_i belongs and let $Q(\alpha) = (Q_1(\alpha_1), ..., Q_n(\alpha_n))$

Equivalence classes of free beliefs.

The set of types $\mathcal{A}_i$ of each agent $i \in N$ can be partitioned into *equivalence classes of free beliefs*.

- An equivalence class is a subset Q_i of $\mathcal{A}_i$ such that $p(. \mid \alpha_i) = p(. \mid \widetilde{\alpha}_i)$ whenever $\alpha_i, \widetilde{\alpha}_i \in Q_i$. We denote by $Q_i(\alpha_i)$, the element of the partition to which α_i belongs and let $Q(\alpha) = (Q_1(\alpha_1), ..., Q_n(\alpha_n))$
- The partition of $\mathcal{A}_i$ is denoted $\mathcal{Q}^i(\mathcal{A}_i)$. For each i and each Q_i, one can define a probability distribution $P(Q_{-i} \mid Q_i)$ on $\times_{j \neq i}^n \mathcal{Q}^j(\mathcal{A}_j)$ from the distribution $p(\alpha_{-i} \mid \alpha_i)$ on $\mathcal{A}_{-i}$. Let $P(Q_{-i} \mid Q_i) = \sum_{\alpha_{-i} \in Q_{-i}} p(\alpha_{-i} \mid Q_i)$, with $p(\alpha_{-i} \mid Q_i) \equiv p(\alpha_{-i} \mid \alpha_i)$, for any $\alpha_i \in Q_i$. When every Q_i is a singleton, P coincides with p.

Equivalence classes of free beliefs.

The set of types $\mathcal{A}_i$ of each agent $i \in N$ can be partitioned into *equivalence classes of free beliefs*.

- An equivalence class is a subset Q_i of $\mathcal{A}_i$ such that $p(. \mid \alpha_i) = p(. \mid \widetilde{\alpha}_i)$ whenever $\alpha_i, \widetilde{\alpha}_i \in Q_i$. We denote by $Q_i(\alpha_i)$, the element of the partition to which α_i belongs and let $Q(\alpha) = (Q_1(\alpha_1), ..., Q_n(\alpha_n))$
- The partition of $\mathcal{A}_i$ is denoted $\mathcal{Q}^i(\mathcal{A}_i)$. For each i and each Q_i, one can define a probability distribution $P(Q_{-i} \mid Q_i)$ on $\times_{j \neq i}^n \mathcal{Q}^j(\mathcal{A}_j)$ from the distribution $p(\alpha_{-i} \mid \alpha_i)$ on $\mathcal{A}_{-i}$. Let $P(Q_{-i} \mid Q_i) = \sum_{\alpha_{-i} \in Q_{-i}} p(\alpha_{-i} \mid Q_i)$, with $p(\alpha_{-i} \mid Q_i) \equiv p(\alpha_{-i} \mid \alpha_i)$, for any $\alpha_i \in Q_i$. When every Q_i is a singleton, P coincides with p.
- Assuming consistency,

$$
\begin{aligned}
P(Q) &= \sum_{\alpha_i \in Q_i} \sum_{\alpha_{-i} \in Q_{-i}} p(\alpha_{-i} \mid \alpha_i) \, p(\alpha_i) \\
&= P(Q_{-i} \mid Q_i) \, P(Q_i) .
\end{aligned}
$$

Mechanisms: Independence Class Efficient

Introducing the notion of belief equivalence classes for every agent allows for a weaker notion of efficiency: "Efficiency within equivalence classes".

- A decision rule s is *Independence Class efficient* (IC-efficient) if, for all $\alpha \in \mathcal{A}$,

$$\sum_i u_i(s(\alpha; \alpha_i) \geq \sum_i u_i(x; \alpha_i) \text{ for all } x \in X\left(Q(\alpha)\right),$$

  where

$$X(Q\left(\alpha\right)) = \{x \mid x = s(\alpha') \text{ for some } \alpha' \in Q\left(\alpha\right)\}.$$

Mechanisms: Independence Class Efficient

Introducing the notion of belief equivalence classes for every agent allows for a weaker notion of efficiency: "Efficiency within equivalence classes".

- A decision rule s is *Independence Class efficient* (IC-efficient) if, for all $\alpha \in \mathcal{A}$,

$$\sum_i u_i(s(\alpha; \alpha_i) \geq \sum_i u_i(x; \alpha_i) \text{ for all } x \in X\left(Q(\alpha)\right),$$

  where

$$X(Q\left(\alpha\right)) = \{x \mid x = s(\alpha') \text{ for some } \alpha' \in Q\left(\alpha\right)\}.$$

- Every efficient decision rule is IC-efficient.

Bayesian Implementation

Bayesian Incentive Compatibility

- A mechanism (s, t) is *Bayesian Incentive Compatible (BIC)* if, for all $\alpha_i' \in \mathcal{A}_i$, $\alpha_i \in \mathcal{A}_i$ and $i \in N$

$$\sum_{\alpha_{-i} \in \mathcal{A}_{-i}} \left[u_i\left(s\left(\alpha_{-i}, \alpha_i\right); \alpha_i\right) + t_i\left(\alpha_{-i}, \alpha_i\right)\right)\right] p\left(\alpha_{-i} \mid \alpha_i\right)$$

$$\geq \sum_{\alpha_{-i} \in \mathcal{A}_{-i}} \left[u_i\left(s\left(\alpha_{-i}, \alpha_i'\right); \alpha_{-i}\right) + t_i\left(\alpha_{-i}, \alpha_i'\right)\right)\right] p\left(\alpha_{-i} \mid \alpha_i\right).$$

i.e., the "truthful" strategy vector $(\widetilde{\alpha}_i\left(\alpha_i\right) \equiv \alpha_i, \forall i)$ is a Bayesian equilibrium.

(i) BIC within independence classes
(ii)Ensuring IIR
(iii) BIC across independence classes

Bayesian Incentive Compatibility

- A mechanism (s, t) is *Bayesian Incentive Compatible (BIC)* if, for all $\alpha_i' \in \mathcal{A}_i$, $\alpha_i \in \mathcal{A}_i$ and $i \in N$

$$\sum_{\alpha_{-i} \in \mathcal{A}_{-i}} \left[u_i\left(s\left(\alpha_{-i}, \alpha_i\right); \alpha_i\right) + t_i\left(\alpha_{-i}, \alpha_i\right)\right] p\left(\alpha_{-i} \mid \alpha_i\right)$$
$$\geq \sum_{\alpha_{-i} \in \mathcal{A}_{-i}} \left[u_i\left(s\left(\alpha_{-i}, \alpha_i'\right); \alpha_{-i}\right) + t_i\left(\alpha_{-i}, \alpha_i'\right)\right] p\left(\alpha_{-i} \mid \alpha_i\right).$$

  i.e., the "truthful" strategy vector $(\widetilde{\alpha}_i\left(\alpha_i\right) \equiv \alpha_i, \forall i)$ is a Bayesian equilibrium.

- For s to be "implemented", we have three requirements:

(i) BIC within independence classes
(ii)Ensuring IIR
(iii) BIC across independence classes

(i) BIC within independence classes

An easy (and constructive) way to get BIC **within independent classes**:
Define Arrow-AGV mechanism (s, θ) with s IC-efficient:
Letting, for all i and all α_i,

$$g_i(\alpha_i) = \sum_{\alpha_{-i}} \sum_{j \neq i} u_j\left(s\left(\alpha_{-i}, \alpha_i\right), \alpha_j\right) p_i(\alpha_{-i} \mid \alpha_i).$$

The monetary transfer to agent i is defined as (to ensure budget-balance)

$$\theta_i\left(\alpha_{-i}, \alpha_i\right) = \left[g_i\left(\alpha_i\right) - \frac{1}{n-1} \sum_{j \neq i} g_j\left(\alpha_j\right)\right]$$

Then, by IC-efficiency, BIC holds for every agent i, within each independence class Q_i.

(ii) A lemma to implement interim individual rationality.

Lemma (1)

Given an information structure $\{\mathcal{A}, p\}$ *with* p *consistent, and numbers* $\{D_i(\alpha_i)\}$*, there exists a balanced transfer function* t *which satisfies*

$$\sum_{\alpha_{-i} \in \mathcal{A}_{-i}} t_i(\alpha_{-i}, \alpha_i) p(\alpha_{-i} \mid \alpha_i) \geq -D_i(\alpha_i), \text{ for all } i \in \mathcal{N} \text{ and all } \alpha_i \in \mathcal{A}_i$$

if and only if

$$\sum_{i \in \mathcal{N}} \sum_{\alpha_i \in \mathcal{A}_i} D_i(\alpha_i)\ p(\alpha_i) \geq 0$$

The non-negative expected surplus condition

The following lemma was first proved in Matsushima (2007, proposition 1)

Lemma (2)

For any decision rule s, there exists a balanced transfer rule t satisfying IIR if and only if the **non-negative expected surplus condition** *holds, i.e.*

$$S \equiv \sum_{i \in N} \sum_{\alpha} u_i(s(\alpha); \alpha_i, \alpha_{-i}) p(\alpha) - \sum_{i \in N} \sum_{\alpha_i} \bar{U}_i(\alpha_i) p(\alpha_i) \geq 0.$$

For the proof, apply lemma (1) letting

$$D_i(\alpha_i) = \sum_{\alpha_{-i} \in \mathcal{A}_{-i}} u_i(s(\alpha_i, \alpha_{-i}); \alpha_i, \alpha_{-i}) p(\alpha_{-i} \mid \alpha_i) - \bar{U}_i(\alpha_i).$$

(iii) BIC across independent classes

The following condition is to ensure BIC **across equivalence classes (while preserving IIR)**. It adapt previous conditions used in the "imposed participation" case: Compatibility condition (or *Condition C*) in AGV79 and *Condition B* in AGV82.

- *Condition* $\overline{B}$: Assume $n \geq 3$. An information structure $\{\mathcal{Q}, P\}$ satisfies *Condition* $\overline{B}$ if and only if there exists a balanced transfer $\tau^{\overline{B}}$ defined on $\mathcal{Q}$ such that, for any Q_i in $\mathcal{Q}^i$

$$\sum_{Q_{-i}\in\mathcal{Q}^{-i}} t_i^{\bar{B}}(Q_{-i}, Q_i)P(Q_{-i} \mid Q_i) > \sum_{Q_{-i}\in\mathcal{Q}^{-i}} t_i^{\bar{B}}(Q_{-i}, \widetilde{Q}_i)P(Q_{-i} \mid Q_i), \text{ for} \tag{BBIC}$$

and

$$\sum_{Q_{-i}} t_i^{\bar{B}}(Q_{-i}, Q_i)P(Q_{-i} \mid Q_i) = 0 \text{ for all } Q_i \in \mathcal{Q}^i \tag{BIIR}$$

(iii) BIC across independent classes

The following condition is to ensure BIC **across equivalence classes (while preserving IIR)**. It adapt previous conditions used in the "imposed participation" case: Compatibility condition (or *Condition C*) in AGV79 and *Condition B* in AGV82.

- *Condition* $\overline{B}$: Assume $n \geq 3$. An information structure $\{\mathcal{Q}, P\}$ satisfies *Condition* $\overline{B}$ if and only if there exists a balanced transfer $\tau^{\overline{B}}$ defined on $\mathcal{Q}$ such that, for any Q_i in $\mathcal{Q}^i$

$$\sum_{Q_{-i} \in \mathcal{Q}^{-i}} t_i^{\bar{B}}(Q_{-i}, Q_i) P(Q_{-i} \mid Q_i) > \sum_{Q_{-i} \in \mathcal{Q}^{-i}} t_i^{\bar{B}}(Q_{-i}, \widetilde{Q}_i) P(Q_{-i} \mid Q_i), \text{ for} \tag{BBIC}$$

  and

$$\sum_{Q_{-i}} t_i^{\bar{B}}(Q_{-i}, Q_i) P(Q_{-i} \mid Q_i) = 0 \text{ for all } Q_i \in \mathcal{Q}^i \tag{BIIR}$$

- *Under no-freeness*, the information structure $\{\mathcal{Q}, P\}$ can be identified to $\{\mathcal{A}, p\}$ since $Q_i(\alpha_i) = \{\alpha_i\}$ so that we simply say $\{\mathcal{A}, p\}$ satisfies *Condition* $\overline{B}$.

This condition can be ensured constructively

Index the agents modulo n. Addition and subtraction on the indices of agents are such that $n+1 \equiv 1$ and $1-1 \equiv n$.

- For every $i \in \mathcal{N}$, and for $j = i+1$ and $j = i-1$, assume that, for all i and all Q_i, $\widetilde{Q}_i$ in $\mathcal{Q}^i$,

$$\begin{aligned} P\left(Q_{-i-(i+1)} \mid Q_i\right) &\equiv \sum_{Q_{i+1}} P\left(Q_{-i} \mid Q_i\right) > 0, \\ \text{and } P\left(Q_{-i-(i-1)} \mid Q_i\right) &\equiv \sum_{Q_{i-1}} P\left(\alpha_{-i} \mid \alpha_i\right) > 0, \end{aligned}$$

  for all $Q_{-i-(i+1)}$ and all $Q_{-i-(i-1)}$.
- We can define the following *transfer scoring rule* which satifies (generically) the two requirements of *Condition* $\overline{B}$:

$$\begin{aligned} \bar{t}_i(Q) = {} & \log P\left(Q_{-i-(i-1)} \mid Q_i\right) - \log P\left(Q_{-i-(i-1)} \mid Q_{i-1}\right) \\ & + \log P\left(Q_{-i-(i+1)} \mid Q_i\right) - \log P\left(Q_{-i-(i+1)} \mid Q_{i+1}\right) \end{aligned}$$

The no-freeness case

No-freeness and full surplus-sharing

The *expected surplus is denoted*
$S = \sum_{i \in N} \sum_{\alpha_i} \left[\sum_{\alpha_{-i}} u_i(s(\alpha); \alpha_i) p\left(\alpha_{-i} \mid \alpha_i\right) - \overline{U}_i\left(\alpha_i\right) \right] p\left(\alpha_i\right)$.
Under no-freeness, Condition $\overline{B}$ is equivalent to Identifiability+Cremer-McLean conditions in Kosenok & Severinov 2008 (lemma A.2 p.144). We have the following result:

Theorem (1)

Assume no-freeness. For any decision rule s, any utility functions $u_i\left(x; \alpha\right)$, any reservation utilities $\overline{U_i}\left(\alpha_i\right)$ such that $S \geq 0$, and for all utility levels $V_i(\alpha_i) \geq \overline{U_i}(\alpha_i)$, there exists a BIC, IIR, balanced mechanism (s, t) which implements s and satisfies

$$\sum_{\alpha_{-i}} \left[u_i\left(s\left(\alpha_{-i}, \alpha_i\right); \alpha_i\right) + t_i\left(\alpha_{-i}, \alpha_i\right)\right] p(\alpha_{-i} \mid \alpha_i) = V_i\left(\alpha_i\right),$$

if and only if $\{\mathcal{A}, p\}$ satisfies condition $\overline{B}$.

Full-surplus-sharing: proof

Under no-freeness, $Q_i(\alpha_i) = \alpha_i$, for all α_i and all i.
Since $S \geq 0$, there is (by lemma 2) a balanced transfer rule τ such that *IIR* holds

$$\sum_{\alpha_{-i}} \left[u_i\left(s\left(\alpha_i, \alpha_{-i}\right); \alpha_i\right) + \tau_i\left(\alpha_i, \alpha_{-i}\right)\right] p\left(\alpha_{-i} \mid \alpha_i\right) \geq \overline{U}_i\left(\alpha_i\right), \text{ for all } \alpha_i, \text{ all } i.$$

Let $t = \tau + K\tau^{\overline{B}}$: t is balanced and, with K large enough, BIC is satisfied.
Since $\sum_{\theta_{-i}} K\tau_i^{\overline{B}}(\alpha_{-i}, \alpha_i)p_i(\alpha_{-i} \mid \alpha_i) = 0$, participation still holds.
This proves suffncency. We omit the proof of necessity (much longer).

Introducing freeness

Introducing freeness: Reinforcing the surplus condition.

- As soon as some free types are introduced (say a pair $(\alpha_i, \widetilde{\alpha}_i)$ s.t. $p(\alpha_{-i} \mid \alpha_i) = p(\alpha_{-i} \mid \widetilde{\alpha}_i)$), **we need to reinforce the nonnegative surplus condition**, and the more so with increasing freeness.

Introducing freeness: Reinforcing the surplus condition.

- As soon as some free types are introduced (say a pair $(\alpha_i, \widetilde{\alpha}_i)$ s.t. $p(\alpha_{-i} \mid \alpha_i) = p(\alpha_{-i} \mid \widetilde{\alpha}_i)$), **we need to reinforce the nonnegative surplus condition**, and the more so with increasing freeness.
- With no-freeness or constant payoffs within equivalence classes it coincides with the nonnegative expected surplus condition.

Introducing freeness: Reinforcing the surplus condition.

- As soon as some free types are introduced (say a pair $(\alpha_i, \widetilde{\alpha}_i)$ s.t. $p(\alpha_{-i} \mid \alpha_i) = p(\alpha_{-i} \mid \widetilde{\alpha}_i)$), **we need to reinforce the nonnegative surplus condition**, and the more so with increasing freeness.
- With no-freeness or constant payoffs within equivalence classes it coincides with the nonnegative expected surplus condition.
- At the limit, for the independent case ,we get the strong condition introduced by Makowski-Mezzetti (1994).

Reinforcing the surplus condition.

To *reinforce the nonnegative expected surplus condition*, let us first define

$$S_i(\alpha_i) = \sum_{\alpha_{-i}} \left\{ \sum_j [u_j(s(\alpha_i, \alpha_{-i}); \alpha_j) - \overline{U}_j(\alpha_j)] \right\} p(\alpha_{-i} \mid \alpha_i),$$

which is the expected surplus conditional on the type of agent i being α_i. Notice that the expected value of $S_i(\alpha_i)$ is independent of i as

$$\sum_{\alpha_i} p(\alpha_i) S_i(\alpha_i) = \sum_{\alpha} \left\{ \sum_j [u_j(s\,(\alpha; \alpha_j)) - \overline{U}_j(\alpha_j) \right\} p(\alpha)] = S,$$

where S is the expected social surplus.

Reinforcing the surplus condition.

We also define

$$S_i^{\min}(Q_i) = \min_{\alpha_i \in Q_i} S_i(\alpha_i).$$

The quantity

$$\sum_{\alpha_i \in Q_i} S_i(\alpha_i) p(\alpha_i \mid Q_i) - S_i^{\min}(Q_i) = \sum_{\alpha_i \in Q_i} \left[S_i(\alpha_i) - S_i^{\min}(Q_i)\right] p(\alpha_i \mid Q_i)$$

is a measure of the variability of S_i within the independence class Q_i.

- The reinforced nonnegative expected surplus condition is:

$$S \geq \sum_i \sum_{\alpha_i} \left[S(\alpha_i) - S_i^{\min}(Q_i(\alpha_i))\right] p(\alpha_i),$$

Implementing IC-efficient collective decision rules with IIR.

Theorem (2)

For any IC-efficient collective decision rule s, any utility functions $u_i\left(x;\alpha\right)$*, any reservation utilities* $\overline{U_i}\left(\alpha_i\right)$ *such that the reinforced nonnegative expected surplus condition holds, there exists a BIC and IIR balanced transfer rule which implements s if and only if* $\left\{\mathcal{Q}, P\right\}$ *satisfies Condition* $\overline{B}$*.*

(i) To prove necessity, assume implementation of any IC-efficient rule. In particular it implements all collective decision rules which are constant on equivalence classes. For such rules, the condition reduces to $S \geq 0$. Hence, the 'only if' part of the proof of Theorem 1 applies.

Implementing IC-efficient collective decision rules: proving sufficiency

(ii) Sufficiency can be proved in three steps.
First step: Construct an Arrow-AGV mechanism (s, θ) with s IC-efficient and taking into account reservation utility levels: for all i and all α_i, we let

$$g_i(\alpha_i) = \sum_{\alpha_{-i}} \sum_{j \neq i} u_j\left(s\left(\alpha_{-i}, \alpha_i\right), \alpha_j\right) p_i(\alpha_{-i} \mid \alpha_i) - \overline{U}_j(\alpha_j).$$

The monetary transfer to agent i is defined as before (to ensure budget-balance)

$$\theta_i\left(\alpha_{-i}, \alpha_i\right) = \left[g_i\left(\alpha_i\right) - \frac{1}{n-1} \sum_{j \neq i} g_j\left(\alpha_j\right)\right]$$

Then, BIC holds for every agent i, within each equivalence class Q_i.

Implementing IC-efficient collective decision rules (second step))

Second step (to ensure *IIR*): We need to find a balanced transfer rule $\overline{\tau}$ which is constant within equivalence class (to preserve BIC) and satisfies, for any $\alpha_i \in \mathcal{A}_i$ and $i = 1, ..., n$:

$$\sum_{\alpha_{-i}} \overline{\tau}_i(\alpha_i, \alpha_{-i}) \, p_i(\alpha_{-i} \mid \alpha_i) \geq -D_i(\alpha_i) \text{ with}$$

$$D_i(\alpha_i) = \min_{\alpha_i' \in Q_i(\alpha_i)} \sum_{\alpha_{-i}} \left[u_i(s(\alpha_i', \alpha_{-i}); \alpha_i') + \theta_i(\alpha_i', \alpha_{-i}) - \overline{U}_i(\alpha_i') \right] p(\alpha_{-i} \mid \alpha$$

Applying Lemma 1, it can be shown that this is true if and only if the "reinforced nonnegative expected surplus condition" holds, *i.e.* this condition is equivalent to

$$\sum_{i \in \mathcal{N}} \sum_{\alpha_i \in \mathcal{A}_i} D_i(\alpha_i) \; p(\alpha_i) \geq 0$$

Implementing IC-efficient collective decision rules (third step)

Third step: Unless full-freeness holds, we still have to ensure BIC across equivalence classes.
Since $\{\mathcal{Q}, P\}$ satisfies *Condition* $\overline{B}$, there exists a balanced transfer rule $\tau^{\overline{B}}$ defined on $\mathcal{Q}$ satisfying BIC strictly.
Choosing $K > 0$ large enough, we get BIC across equivalence classes (while preserving IIR) and the mechanism

$$\left(s, \theta + \overline{\tau} + K\tau^{\overline{B}}\right)$$

is both budget balanced and BIC.

Concluding remarks

- Bayesian mechanism design enlarges the possibilities for applications (e.g. Internet applications). It requires conditions on the conditional beliefs of the agents. We have introduced conditions between the the two traditional extremes, full-freeness (independence) and no-freeness (BDP)

- Bayesian mechanism design enlarges the possibilities for applications (e.g. Internet applications). It requires conditions on the conditional beliefs of the agents. We have introduced conditions between the the two traditional extremes, full-freeness (independence) and no-freeness (BDP)
- These in-between conditions are based on the notion of 'free belief equivalence classes' and are necessary and sufficient for implementing IC-efficient collective decision rules

- Bayesian mechanism design enlarges the possibilities for applications (e.g. Internet applications). It requires conditions on the conditional beliefs of the agents. We have introduced conditions between the the two traditional extremes, full-freeness (independence) and no-freeness (BDP)
- These in-between conditions are based on the notion of 'free belief equivalence classes' and are necessary and sufficient for implementing IC-efficient collective decision rules
- For voluntary participation (IIR) one needs more than two agents and some restrictions on the payoffs (non-negative expected surplus condition or more, the strongest being Makowski-Mezetti surplus condition when independence holds).

Robust mechanisms?

- Albert-Conitzer-Lopomo-Stone (2018) introduces what they call $\varepsilon-$**robust mechanisms** that extend robust mechanisms (Bergemann and Morris 2005) allowing for a small probability of violation of incentive compatibility and individual rationality when there is uncertainty over the distribution of bidders.

Robust mechanisms?

- Albert-Conitzer-Lopomo-Stone (2018) introduces what they call $\varepsilon-$**robust mechanisms** that extend robust mechanisms (Bergemann and Morris 2005) allowing for a small probability of violation of incentive compatibility and individual rationality when there is uncertainty over the distribution of bidders.
- Assuming the true distribution satisfies the Crémer-McLean condition, they are able to compute these mechanisms efficiently and learn nearly optimal mechanisms (with sufficient correlation) using a polynomial number of samples from the true distribution.

Robust mechanisms?

- Albert-Conitzer-Lopomo-Stone (2018) introduces what they call $\varepsilon-$**robust mechanisms** that extend robust mechanisms (Bergemann and Morris 2005) allowing for a small probability of violation of incentive compatibility and individual rationality when there is uncertainty over the distribution of bidders.
- Assuming the true distribution satisfies the Crémer-McLean condition, they are able to compute these mechanisms efficiently and learn nearly optimal mechanisms (with sufficient correlation) using a polynomial number of samples from the true distribution.
- Hopefully, our results "suggest that there are likely to be computationally feasible robust mechanisms that approximately achieve budget balance and social efficiency" (assuming the true distribution satisfies *Condition* $\overline{B}$?)