Mechanism Design and Incomplete Information

Arrow Lecture

Stephen Morris

July 5, 2017

Mechanism Design

Mechanism design is the "reverse engineering" part of economic theory. Normally, economists study existing economic institutions and try to predict or explain what outcomes the institutions generate. But in mechanism design we reverse direction: we start by identifying the outcomes we want and then ask what institutions could be designed to achieve those outcomes.

A Tripartite Distinction of Informational Assumptions (in Economic Theory or Economic Design)

1. Perfect Information
   - Everything is common knowledge among "players"
2. Complete but Imperfect Information
   - Even if there is not perfect information (e.g., there is uncertainty and asymmetric information), there is common knowledge about the structure of the environment
3. Incomplete Information
   - No assumptions about anything

Informational Assumptions in Games

von Neumann and Morgenstern "Theory of Games and Economic Behavior" 1944

1. Perfect Information Games
   - There is common knowledge of the structure of a game being played: players, the order in which they move, previous moves, payoffs, etc...
   - LEADING EXAMPLE: Chess
2. Complete but Imperfect Information
   - There is common knowledge of the structure of the game being played: players, rules of the game, feasible strategies, payoffs, etc....; but may not know past or current actions of other players or exogenous uncertainty
   - LEADING EXAMPLE: Poker
3. Incomplete Information
   - There is not common knowledge of the structure of the game being played
   - Leading EXAMPLE: All economic environments of interest?

A Pessimistic Assessment

von Neumann and Morgenstern "Theory of Games and Economic Behavior" 1944

...we cannot avoid the assumption that all subjects under consideration are completely informed about the physical characteristics of the situation in which they operate

- Aumann (1987) wrote "The common knowledge assumption underlies all of game theory and much of economic theory. Whatever be the model under discussion ... the model itself must be assumed common knowledge; otherwise the model is insufficiently specified, and the analysis incoherent."

John Harsanyi 1967/68

- incomplete Information is not a problem
- we can incorporate any incomplete information without loss of generality!

John Harsanyi Part 1: Type Spaces

- there is a set of states Θ that we care about
- two players, Ann and Bob (generalize straightforwardly to many players)
- each player has a space of possible "types": T_A, T_B
- write $\pi_A(t_B, \theta | t_A)$ for the probability that type t_A of Ann assigns to both Bob being type t_B and the state being θ; so we have

$$\pi_A : T_A \to \Delta(T_B \times \Theta)$$

and analogously

$$\pi_B : T_B \to \Delta(T_A \times \Theta)$$

- The state space Θ can embed a lot of stuff...
  - in game theory, it can encompass the rules of the game and payoffs
  - in mechanism design, it can encompass players' preferences over outcomes, beliefs about others' preferences, payoff relevant states and so on

John Harsanyi Part 2: Universal Type Spaces

- Ann is characterized by...
  1. her belief about the state
  2. her belief about the state and the Bob's belief about the state
  3. her belief about the state and [Bob's belief about the state and Ann's belief about the state]
  4. and so on....
- So Ann is characterized by this infinite sequence of such higher order beliefs, or universal types
- "universal type space" T^* satisfies $T^* \approx \Delta\left(T^* \times \Theta\right)$
- We can assume that this structure is common knowledge
- Incomplete information is not a problem after all!

Filling in the Details...

Harsanyi is actually rather confusing on part 2 and some high-powered people formalized this over the next 25 years:

1. Aumann (1976) had to define "common knowledge" (needed to formalize the statements above) from discussions with **Arrow** and Hahn
2. Mertens and Zamir (1987) and Brandenburger and Dekel (1993) formalized the universal type space (dealing with both conceptual and mathematical issues glossed over in my discussion above)
3. Aumann (1989, 1999) syntactic formulation of "without loss of generality common knowledge" claim

The Misunderstanding of John Harsanyi

- the good news:
  - by working with the universal type space, we can dispense with common knowledge assumptions
- the bad news:
  - the economics profession (in general and in mechanism design) went straight back to make unrealistic complete information assumptions, by working with "small" and otherwise simple type spaces (e.g., independent types)
- this lecture:
  - implications of the "misunderstanding" for mechanism design

The Misunderstanding of John Harsanyi and Mechanism Design

Wilson (1987):

> *Game theory has a great advantage in explicitly analyzing the consequences of trading rules that presumably are really common knowledge; it is deficient to the extent that it assumes other features to be common knowledge, such as one agent's probability assessment about another's preferences or information.*
>
> *I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality."*

The Misunderstanding of John Harsanyi and Mechanism Design

- Mechanism design
  - we would really like to assume that there is complete information about the *game/mechanism*
  - it is particularly desirable to relax common knowledge assumptions about the environment, because optimal mechanisms are otherwise too finely tuned
- Contrast this with economic theory / game theory
  - really important to relax common knowledge of the mechanism (John Sutton and IO)
  - common knowledge of the environment is (at least a bit) less of a problem
- One response to misunderstanding: do not address "incomplete information", focus on simple mechanisms, computational constraints, worst case analysis, etc...
- Another response:
  - take relaxing common knowledge assumptions seriously and allowing real incomplete information in mechanism design

An Important Restriction on Type Space: Player Specific Parameters

- suppose that Ann's preferences are summarized by a parameter $\theta_A \in \Theta_A$ (known to Ann), and similarly for Bob ("private values")
- natural to consider slightly different type spaces:
  - Ann has a set of types T_A, where a type is characterized by a payoff parameter $\widehat{\theta}_A(t_A) \in \Theta_A$ and a belief $\widehat{\pi}_A(t_A) \in \Delta(T_B)$
  - similarly for Bob

Universal Type Space with Player Specific Parameters

- Ann's universal type is her payoff parameter and
  1. her belief about Bob's payoff parameter
  2. her belief about Bob's belief and his payoff parameter
  3. and so on....
- Ann's universal type space is $T_A^* \approx \Theta_A \times \Delta\left(T_B^*\right)$
- Is a subset of our first universal type space

Type Space Restrictions = Implicit Common Knowledge Assumptions

- Common to assume:
  1. naive type space (identify types with payoff parameters)
  2. common prior (beliefs could have been derived from common prior and Bayes updating)
  3. and either *independence* or beliefs determine payoff parameters (*BDP*: Neeman 2004) implied by generic beliefs on naive type space
- Sometimes implicitly or explicitly trying to implement on all types spaces in some class (e.g., all naive common prior independent type spaces)
- Implementing on the universal type space is the same (modulo technicalities) as implementing on all types spaces

Funny Result 1: Full Surplus Extraction

- Consider the private good allocation problem with private values and transfers. Easy to implement the efficient allocation. But two key results about revenue:
  - with independent naive common prior type space, buyers earn information rent
  - with BDP naive common prior type space, efficient allocation and full surplus extraction
    - players can be given a strictly positive incentive to truthfully announce their types via bets at no expected cost
- One response:
  - BDP does (or does not) hold generically on the universal type space
- Nuanced response:
  - There is not full surplus extraction on the universal type space
  - Take a position on which types in the universal type space are relevant

Funny Result 2: Prior Extraction

- Consider a public goods problem with private values and budget balanced transfers. Two key public good results:
  - Not possible to implement efficient choice in dominant strategies
  - Possible to implement efficient choice in (Bayes) Nash equilibrium
    - with independent types, AGV (see also **Arrow**)
- But what if the prior is not known? Two responses:
  - back to dominant strategies and negative results
  - prior extraction: ask players to report their common prior and shoot them if they report something different
- Alternative nuanced response: relax union of common prior naive type spaces assumption to universal type space. Nuanced conclusion:
  - Implemention of the efficient outcome in Bayes Nash equilibrium on universal type space may or may not be equivalent to dominant strategies implementation

Relaxing Private Values Assumption

- Maintained common knowledge assumption in discussion so far: private values
- Let's relax this assumption
- Suppose that values are interdependent
- Ann's value of an object is $v_A = \theta_A + \gamma\theta_B$ for some $0 < \gamma < 1$
- Analogously, Bob's value is $v_B = \theta_B + \gamma\theta_A$

Three Interpretations

1. θ_A is Ann's consumption value but it is possible that Ann will have to re-sell to Bob, extracting proportion γ of Bob's value
2. Ann and Bob each have a signal that confounds a common value and private value component (cannot be distinguished)

Implicit Common Knowledge Assumptions and Interdependent Values

- In example, we have single good interdependent values example, we had

$$v_A = \theta_A + \gamma\theta_B \text{ and } v_B = \theta_B + \gamma\theta_A$$

- By linear algebra, we have

$$\theta_A = \frac{1}{1-\gamma^2}\left(v_A - \gamma v_B\right) \text{ and } \theta_B = \frac{1}{1-\gamma^2}\left(v_B - \gamma v_A\right)$$

- So if we considered the player specific payoff parameter universal type space for (θ_A, θ_B), we were implicitly assuming that there was common knowledge that Ann knows $v_A - \gamma v_B$ and Bob knows $v_B - \gamma v_A$
- Whether this makes sense depends on the interpretation

Canonical Preference Higher-Order Preference Types

Should actually distinguish "higher order preference types", e.g.,

1. first order valuation: Ann's unconditional value of an object,
2. second order belief and valuation:
   - Ann's belief about Bob's first order valuation
   - Ann's valuation conditional on Bob's first order valuations
3. third order belief and valuation:
   - Ann's belief about Bob's second order type
   - Ann's valuation conditional on Bob's second order type
4. and so on

Thus we are looking higher order preferences over acts

Universal Higher-Order Preference Type Space

- Can represent all higher-order preference types by universal space $T^* = \Delta\left(T^* \times \{0, 1\}\right)$
- In this representation, Ann's probability of state 1 is her unconditional valuation of the object
- Higher order preference types correspond exactly to what would be learnt about players
- selling on the higher-order preference type space is complicated

Conclusion

- Incomplete information has not been fully incorporated into mechanism design
- Results are driven by implicit common knowledge whose role is sometimes not well understood
- But relaxing all common knowledge assumptions may be possible but unhelpful
- Focus on which are reasonable common knowledge assumptions and make them explicit

New Questions

- Summer School included some discussions of existing results
- New questions:
  - Public Goods
  - Common Knowledge of Independent Priors