Strategy-proofness in the Large

Eric Budish
U. Chicago

"Beyond Strategyproofness" panel at INFORMS / EC 18

Strategy-proofness in Market Design

- Strategyproofness (SP) – reporting your preferences truthfully is a dominant strategy – is perhaps the predominant notion of incentives in market design
  - Frequently imposed as a theoretical design requirement, across a wide variety of auction, assignment and matching problems
  - Explicit role in recent real-world reforms in school choice, kidney exchange, two-sided matching (Roth, 2008)
- Many reasons why SP is so heavily emphasized relative to Bayesian or Nash implementation:
  1. Wilson doctrine (Bergemann Morris, 2005)
  2. Strategically simple for participants (Fudenberg Tirole, 1991)
  3. SP as fairness: unsophisticated players are not disadvantaged (Friedman 1991, Pathak Sonmez 2008)
  4. SP mechanisms generate information about preferences that may be useful to policy makers (Roth, 2008)

The Limits of SP in Market Design

However, in numerous market design contexts, impossibility theorems indicate that SP severely limits what is possible

- General equilibrium / Walrasian mechanism: Hurwicz's (1972) impossibility theorem
- Stable matching: Roth's (1982) impossibility theorem
- Multi-unit assignment: Papai's (2001) and Ehlers-Klaus's (2003) dictatorship theorems
- School choice: Abdulkadiroğlu, Pathak and Roth's (2009) impossibility theorem
- Quasi-linear setting: Green-Laffont's (1977) VCG theorem, in light of Ausubel-Milgrom (2006)
- Many, many others

Takeaway: SP may be attractive, but it is expensive!

This Paper

This Paper

Goal: propose a new criterion of approximate strategyproofness and show that it is a useful second-best

This Paper

Goal: propose a new criterion of approximate strategyproofness and show that it is a useful second-best

Strategyproof in the Large (SP-L): for any agent, any full-support iid probability distribution of the other agents' reports, and any $\epsilon > 0$, in a large enough market the agent maximizes his expected payoff to within ϵ by reporting his preferences truthfully.

This Paper

Goal: propose a new criterion of approximate strategyproofness and show that it is a useful second-best

Strategyproof in the Large (SP-L): for any agent, any full-support iid probability distribution of the other agents' reports, and any $\epsilon > 0$, in a large enough market the agent maximizes his expected payoff to within ϵ by reporting his preferences truthfully.

- Heuristically: SP-L requires that an agent who regards a mechanism's "prices" as exogenous to her report can do no better than report truthfully
  - Could be traditional prices (e.g. auction) or price-like statistics (e.g. matching)

This Paper

Goal: propose a new criterion of approximate strategyproofness and show that it is a useful second-best

Strategyproof in the Large (SP-L): for any agent, any full-support iid probability distribution of the other agents' reports, and any $\epsilon > 0$, in a large enough market the agent maximizes his expected payoff to within ϵ by reporting his preferences truthfully.

- Heuristically: SP-L requires that an agent who regards a mechanism's "prices" as exogenous to her report can do no better than report truthfully
  - Could be traditional prices (e.g. auction) or price-like statistics (e.g. matching)
- Positioning "in between" approx SP and approx Bayes-Nash
  - Weaker than approximate SP: any full-support probability distribution of opponent reports, rather than any realization
  - Stronger than approximate Bayes Nash, which assumes common knowledge of the true probability distribution.

Problem	Manipulable in the Large	SP-L
Multi-Unit Auctions	Pay-As-Bid	Uniform Price
Single-Unit Assignment	Boston Mechanism	Probabilistic Serial HZ Pseudomarket
Multi-Unit Assignment	Bidding Points Auction HBS Draft	CEEI Generalized HZ
Matching	Priority Match	Deferred Acceptance
Other		Double Auctions Walrasian Mechanism

Observations

- Organizes Milton Friedman on auctions, Al Roth on matching
- Extant theory argument for Approx IC in large markets –> SP-L
- Manipulable in the Large –> Empirical Evidence of Problems in Practice
- We would *not* get this classification with ϵ-SP: too demanding

Obtaining the Classification

- To show that a mechanism is not SP-L: suffices to produce an example of a profitable manipulation in the large-market limit
  - Relatively straightforward for each of the mechanisms in the table (App. B)

Obtaining the Classification

- To show that a mechanism is not SP-L: suffices to produce an example of a profitable manipulation in the large-market limit
  - Relatively straightforward for each of the mechanisms in the table (App. B)
- To show that a mechanism is SP-L, we provide two sufficient conditions
- Condition 1: Envy freeness
  - A direct mechanism $\{(\Phi^n)_{\mathbb{N}}, T\}$ is **envy-free** (**EF**) if, for all i, j, n, t:

$$u_{t_i}[\Phi_i^n(t)] \geq u_{t_i}[\Phi_j^n(t)].$$

  - Proposition: EF$\rightarrow$SP-L
- This condition covers most of the mechanisms in the table (including Uniform-Price Auctions)

EF→SP-L: Idea of Proof

- Decompose the gain to type t_i from misreporting as t_j as
  1. Gain from receiving t_j's bundle, holding fixed the realized empirical distribution of types
  2. Gain from affecting the distribution of the realized empirical distribution of types
- Envy-Freeness directly implies that (1) is non-positive (so long as the realized empirical has full support, which has probability going to one)
- A probabilistic argument establishes that (2) becomes negligible in large markets
  - Relies on full-support and iid: else, there could be a realized empirical where agent i single-handedly affects the probability by a non-vanishing amount (e.g. the probability that zero players report t_j)
  - Relies on ex-interim perspective of SP-L: for instance, uniform-price auctions are envy free, but it is always possible to construct a realizations of others' reports where t_i is pivotal and prefers to report as t_j

Obtaining the Classification

- Condition 2: Envy freeness "but for tie breaking"
  - A direct mechanism $\{(\Phi^n)_{\mathbb{N}}, T\}$ is **envy-free but for tie breaking** (**EF-TB**) if for each n there exists a function $x^n : (T \times [0,1])^N \to \Delta(X_0^n)$, symmetric over its coordinates, such that

$$\Phi^n(t) = \int_{l \in [0,1]^n} x^n(t,l)dl$$

    and, for all i, j, n, t, and l, if $l_i \geq l_j$ then

$$u_{t_i}[x_i^n(t,l)] \geq u_{t_i}[x_j^n(t,l)].$$

- Proposition: EF-TB$\to$SP-L
- Covers the rest of the mechanisms in the table
  - Approximate CEEI and Deferred Acceptance are EF-TB but are not EF (per an example in Bogomolnaia and Moulin, 2001)
- Proof more involved, see Appendix A for details

Relationship to Theory Literature on Large Markets

Problem	Manipulable in the Large	SP-L
Multi-Unit Auctions	Pay-As-Bid	**Uniform Price**
Single-Unit Assignment	Boston Mechanism	**Probabilistic Serial** **HZ Pseudomarket**
Multi-Unit Assignment	Bidding Points Auction HBS Draft	**CEEI** **Generalized HZ**
Matching	Priority Match	**Deferred Acceptance**
Other		**Double Auctions** **Walrasian Mechanism**

Relationship to Theory Literature on Large Markets

- We obtain results for several mechanisms whose large-market incentives properties are well understood:
  - Uniform-Price Auctions (Swinkels, 2001)
  - Probabilistic Serial (Kojima Manea, 2010)
  - Deferred Acceptance (Immorlica Mahdian 2005; Kojima Pathak, 2009)
  - Double Auctions (Rustichini Satterthwaite Williams 1994; Cripps Swinkels 2006)
  - Walrasian Mechanism (Roberts and Postlewaite 1976; Jackson and Manelli 1997)
- As well as for some mechanisms whose large-market properties are less well understood:
  - Hylland-Zeckhauser Pseudomarket (1979) and its generalization (Budish, Che, Kojima and Milgrom, 2012)
  - Approximate CEEI (Budish, 2011)

Relationship to Theory Literature on Large Markets

- Moreover, we obtain these results using a single notion of approximate incentive compatibility, SP-L
- Previous literature has used different notions, tailored for each mechanism
  - Roberts and Postlewaite: truthful reporting is ex-post approximately optimal for all opponent reports where eqm prices vary continuously with reports
  - RSW: exact Bayes-Nash equilibria
  - Swinkels: both exact and approximate Bayes-Nash equilibria
  - Kojima and Pathak: approximate Nash equilibria, with complete information on one side of the market and incomplete on the other side. Also approximate Bayes-Nash
  - Kojima and Manea: exact SP, in a large enough finite market
- Tradeoffs
  - SP-L weaker than many of the previous notions (or non-comparable)
  - We require finite type, action, outcome spaces

Relationship to Empirical Literature on Manipulability

Problem	Manipulable in the Large	SP-L
Multi-Unit Auctions	**Pay-As-Bid**	Uniform Price
Single-Unit Assignment	**Boston Mechanism**	Probabilistic Serial HZ Pseudomarket
Multi-Unit Assignment	**Bidding Points Auction** **HBS Draft**	CEEI Generalized HZ
Matching	**Priority Match**	Deferred Acceptance
Other		Double Auctions Walrasian Mechanism

Relationship to Empirical Literature on Manipulability

- For each of the mechanisms in the Manipulable in the Large column of the table, there is empirical evidence that participants strategically misreport their preferences in practice.
- Also evidence that some participants fail to play best responses, and that this undermines efficiency, fairness, or other design objectives
  - Pay-as-bid auctions: Friedman (1991), Jegadeesh (1993), Brenner et al. (2009)
  - Boston mechanism: Abdulkadiroğlu et al (2006, 2009)
  - Bidding points auction: Krishna and Ünver (2008), Budish (2011)
  - HBS draft mechanism: Budish and Cantillon (2012)
  - Priority match: Roth (1990, 1991, 2002)

SP-L is cheap relative to Bayes-Nash

- Theorem (stated informally):
  - Consider a social choice function F defined over Ω^*, the union of all common prior, i.i.d., full support type spaces
  - If F is continuous and limit Bayes-Nash implementable, then F is approximately SP-L implementable

- Translation: SP-L is cheap relative to Bayes-Nash implementation, at least in large, anonymous markets
- Example: market design debate over the Boston mechanism
- N.B. result generalizes to accommodate some kinds of discontinuities as well (see appendix)

Summary

- We propose SP-L as a second-best alternative to SP:
  1. Many of the benefits of SP design favor SP-L design as well
     - Wilson doctrine, strategic simplicity, fairness, data
  2. Classification of non-SP mechanisms favors SP-L
     - Organizes Friedman on auctions, Roth on matching
     - Organizes theory literature on approx IC in large markets
     - Empirical evidence: It is mechanisms that not only are not SP but that are *not even SP-L* that have problems in practice
  3. Under some assumptions, SP-L is approximately costless relative to Bayes-Nash or Nash

Hope for the paper

- Hope: paper will be viewed as formal justification for focusing on SP-L when confronting a new market design problem for which there are no good SP solutions
  - Example: Budish (JPE 2011) on course allocation, given Papai-Ehlers-Klaus dictatorship theorems (implemented at Wharton in Fall 2013, now several more)
  - Fudenberg: "SP is a virus"
- Caveat: of course, large-market limit is an abstraction – frequently useful, always imperfect
  - Just as the assumption of price-taking behavior is a useful if imperfect abstraction in some other parts of economics
- In environments where this abstraction is compelling: **consider designing a mechanism that is SP-L!**