Strategyproofness in the Large

Eduardo Azevedo and Eric Budish

University of Pennsylvania Theory Seminar
Nov 2012

Strategyproofness 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)

Strategyproofness 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)

The Limits of SP in Market Design

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

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.

This Paper

Argument for SP-L as a second-best:

This Paper

Argument for SP-L as a second-best:

1. In large markets, SP-L approximates many of the formal advantages of SP over Bayes-Nash or Nash implementation
    - Wilson doctrine, strategic simplicity, fairness

This Paper

Argument for SP-L as a second-best:

1. In large markets, SP-L approximates many of the formal advantages of SP over Bayes-Nash or Nash implementation
   - Wilson doctrine, strategic simplicity, fairness
2. Classification of non-SP mechanisms supports SP-L

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

This Paper

Argument for SP-L as a second-best:

1. In large markets, SP-L approximates many of the formal advantages of SP over Bayes-Nash or Nash implementation
   - Wilson doctrine, strategic simplicity, fairness
2. Classification of non-SP mechanisms supports 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

This Paper

Argument for SP-L as a second-best:

1. In large markets, SP-L approximates many of the formal advantages of SP over Bayes-Nash or Nash implementation
   - Wilson doctrine, strategic simplicity, fairness
2. Classification of non-SP mechanisms supports 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. Main theory result: in large markets, SP-L is in a certain sense costless to satisfy relative to Bayes-Nash or Nash
   - Key conditions: finite type and outcome spaces, private values, (semi)-anonymity, (quasi)-continuity
   - Proof is by construction: given a mechanism with Bayes-Nash equilibria, construct an SP-L mechanism that implements approximately the same outcomes

This Paper

Argument for SP-L as a second-best:

1. In large markets, SP-L approximates many of the formal advantages of SP over Bayes-Nash or Nash implementation
   - Wilson doctrine, strategic simplicity, fairness
2. Classification of non-SP mechanisms supports 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. Main theory result: in large markets, SP-L is in a certain sense costless to satisfy relative to Bayes-Nash or Nash
   - Key conditions: finite type and outcome spaces, private values, (semi)-anonymity, (quasi)-continuity
   - Proof is by construction: given a mechanism with Bayes-Nash equilibria, construct an SP-L mechanism that implements approximately the same outcomes

Overall: SP-L approximates the benefits of SP, while being approximately costless to impose

Roadmap

- Introduction
- **Environment**
- Strategyproof in the Large
- Classification of non-SP Mechanisms
- Constructing SP-L Mechanisms
- Discussion and Extensions
- Conclusion

Preliminaries

Preliminaries

- **Finite outcome space**, X_0, with $X = \Delta X_0$
  - E.g. consumption bundle (+ transfer), school, match partner

Preliminaries

- **Finite outcome space**, X_0, with $X = \Delta X_0$
  - E.g. consumption bundle (+ transfer), school, match partner
- **Finite type space**, T. For each $t_i \in T$ there is a vNM utility function $u_{t_i} : X \to [0, 1]$
  - Preferences are **private values**

Preliminaries

- **Finite outcome space**, X_0, with $X = \Delta X_0$
  - E.g. consumption bundle (+ transfer), school, match partner
- **Finite type space**, T. For each $t_i \in T$ there is a vNM utility function $u_{t_i} : X \to [0, 1]$
  - Preferences are **private values**
- For each market size $n \in \mathbb{N}$, there is an arbitrary set $Y_n \subseteq (X_0)^n$ of feasible allocations in an economy with n agents
  - E.g.: capacity of each object in X_0 grows linearly with n
  - Notice: X_0 held fixed as n grows (T as well)

Preliminaries

- **Finite outcome space**, X_0, with $X = \Delta X_0$
  - E.g. consumption bundle (+ transfer), school, match partner
- **Finite type space**, T. For each $t_i \in T$ there is a vNM utility function $u_{t_i} : X \to [0, 1]$
  - Preferences are **private values**
- For each market size $n \in \mathbb{N}$, there is an arbitrary set $Y_n \subseteq (X_0)^n$ of feasible allocations in an economy with n agents
  - E.g.: capacity of each object in X_0 grows linearly with n
  - Notice: X_0 held fixed as n grows (T as well)
- A **Mechanism** consists of a **finite action space** A, and a sequence $(\Phi^n)_{n\in\mathbb{N}}$ of allocation functions

$$\Phi^n : A^n \to \Delta((X_0)^n)$$

each of which satisfies feasibility

Anonymity

- We limit attention to mechanisms that are **anonymous**
  - Each agent's outcome is a common function of her own action and the distribution of all actions
  - More formally: each function $\Phi^n(\cdot)$ is invariant to permutations

Anonymity

- We limit attention to mechanisms that are **anonymous**
  - Each agent's outcome is a common function of her own action and the distribution of all actions
  - More formally: each function $\Phi^n(\cdot)$ is invariant to permutations
- Analysis generalizes to semi-anonymous mechanisms as defined in Kalai (2004)
  - Finite set of groups
  - Each agent's outcome is a function of
    - her own action
    - what group she belongs to
    - the distribution of actions within each group

Limit Mechanisms

Limit Mechanisms

- Given anonymity, we can think about mechanisms from the perspective of a generic agent i

Limit Mechanisms

- Given anonymity, we can think about mechanisms from the perspective of a generic agent i
- Let $\phi^n(a_i, m)$ be the random allocation agent i gets under mechanism $\{(\Phi^n)_{n\in\mathbb{N}}, A\}$ when
  - There are n agents total
  - Agent i plays action a_i
  - The other $n-1$ agents play **iid** according to action distribution $m \in \Delta A$ ("ex interim")

$$\phi^n(a_i, m) = \sum_{a_{-i}} \Phi_i^n(a_i, a_{-i}) \cdot \Pr(a_{-i}|a_{-i} \sim iid(m)) \quad (1)$$

Limit Mechanisms

- Given anonymity, we can think about mechanisms from the perspective of a generic agent i
- Let $\phi^n(a_i, m)$ be the random allocation agent i gets under mechanism $\{(\Phi^n)_{n\in\mathbb{N}}, A\}$ when
  - There are n agents total
  - Agent i plays action a_i
  - The other $n-1$ agents play **iid** according to action distribution $m \in \Delta A$ ("ex interim")

$$\phi^n(a_i, m) = \sum_{a_{-i}} \Phi^n_i(a_i, a_{-i}) \cdot \Pr(a_{-i} | a_{-i} \sim iid(m)) \quad (1)$$

- The function $\phi^\infty : A \times \Delta A \to X$ is the **limit of mechanism** $\{(\Phi^n)_{\mathbb{N}}, A\}$ if, for all a_i, m:

$$\phi^\infty(a_i, m) = \lim_{n\to\infty} \phi^n(a_i, m)$$

The randomness in our definition of the limit is useful for two reasons, one technical and one substantive:

1. Well defined for any $m \in \Delta A$, not just rationals
2. Any specific empirical distribution of opponent play becomes increasingly rare as the market grows large

- Example: $A = \{Heads, Tails\}$, $m = 0.5$,
  - Expected number of *Heads* is of course $\frac{n}{2}$
  - But likelihood of exactly $\frac{n}{2}$ *Heads* goes to zero as $n \to \infty$

The randomness in our definition of the limit is useful for two reasons, one technical and one substantive:

1. Well defined for any $m \in \Delta A$, not just rationals
2. Any specific empirical distribution of opponent play becomes increasingly rare as the market grows large

- Example: $A = \{Heads, Tails\}$, $m = 0.5$,
  - Expected number of *Heads* is of course $\frac{n}{2}$
  - But likelihood of exactly $\frac{n}{2}$ *Heads* goes to zero as $n \to \infty$
- Economic interpretation: so long as being pivotal is a "knife edge" event – e.g., exactly $\frac{n}{2}$ *Heads* is the knife edge – agents will regard the probability of being pivotal as zero in the large-market limit
  - Allows us to think of m as encoding prices (or price-like statistics) which are exogenous from the perspective of each agent
  - Note role of full support

- Most (if not all?) practical market design mechanisms have limits as we have defined them
- But it is very easy to construct examples that do not.
  - E.g., if a mechanism behaves like a uniform-price auction when n is even and like a pay-as-bid auction when n is odd it will not have a limit
- From here forward we limit attention to mechanisms that have limits

Roadmap

- Introduction
- Environment
- **Strategyproof in the Large**
- Classification of non-SP Mechanisms
- Constructing SP-L Mechanisms
- Discussion and Extensions
- Conclusion

Strategyproofness in the Large

Strategyproofness in the Large

- Consider direct mechanisms, $A = T$.

Strategyproofness in the Large

- Consider direct mechanisms, $A = T$.
- Mechanism $\{(\Phi^n)_{\mathbb{N}}, T\}$ is **strategyproof** (SP) if for all t_i, t_i' in T, all n and t_{-i} in T^{n-1}

$$u_{t_i}[\Phi_i^n(t_i, t_{-i})] \geq u_{t_i}[\Phi_i^n(t_i', t_{-i})]$$

Strategyproofness in the Large

- Consider direct mechanisms, $A = T$.
- Mechanism $\{(\Phi^n)_{\mathbb{N}}, T\}$ is **strategyproof** (SP) if for all t_i, t_i' in T, all n and t_{-i} in T^{n-1}

$$u_{t_i}[\Phi_i^n(t_i, t_{-i})] \geq u_{t_i}[\Phi_i^n(t_i', t_{-i})]$$

- Mechanism $\{(\Phi^n)_{\mathbb{N}}, T\}$ is **strategyproof in the large** (SP-L) if for any t_i, t_i' in T, and any full support distribution of types $m \in \Delta T$

$$u_{t_i}[\phi^\infty(t_i, m)] \geq u_{t_i}[\phi^\infty(t_i', m)]$$

Strategyproofness in the Large

- Consider direct mechanisms, $A = T$.
- Mechanism $\{(\Phi^n)_{\mathbb{N}}, T\}$ is **strategyproof** (SP) if for all t_i, t_i' in T, all n and t_{-i} in T^{n-1}

$$u_{t_i}[\Phi_i^n(t_i, t_{-i})] \geq u_{t_i}[\Phi_i^n(t_i', t_{-i})]$$

- Mechanism $\{(\Phi^n)_{\mathbb{N}}, T\}$ is **strategyproof in the large** (SP-L) if for any t_i, t_i' in T, and any full support distribution of types $m \in \Delta T$

$$u_{t_i}[\phi^\infty(t_i, m)] \geq u_{t_i}[\phi^\infty(t_i', m)]$$

- Else, the mechanism is **manipulable in the large**.

SP-L vs. Related Concepts

- **SP-L**: for any t_i, t_i', m

$$u_{t_i}[\phi^\infty(t_i, m)] \geq u_{t_i}[\phi^\infty(t_i', m)]$$

SP-L vs. Related Concepts

- **SP-L**: for any t_i, t_i', m

$$u_{t_i}[\phi^\infty(t_i, m)] \geq u_{t_i}[\phi^\infty(t_i', m)]$$

- **SP-L (alt. statement)**: for any t_i, t_i', m, and any $\epsilon > 0$, there exists n_0 such that if $n > n_0$ we have

$$u_{t_i}[\phi^n(t_i, m)] \geq u_{t_i}[\phi^n(t_i', m)] - \epsilon.$$

SP-L vs. Related Concepts

- **SP-L**: for any t_i, t_i', m

$$u_{t_i}[\phi^{\infty}(t_i, m)] \geq u_{t_i}[\phi^{\infty}(t_i', m)]$$

- **SP-L (alt. statement)**: for any t_i, t_i', m, and any $\epsilon > 0$, there exists n_0 such that if $n > n_0$ we have

$$u_{t_i}[\phi^{n}(t_i, m)] \geq u_{t_i}[\phi^{n}(t_i', m)] - \epsilon.$$

- **Approximate Bayes-Nash**: for the true prior $\mu_0 \in \Delta T$, any t_i, t_i', and any $\epsilon > 0$, there exists n_0 such that if $n > n_0$ we have

$$u_{t_i}[\phi^{n}(t_i, \mu_0)] \geq u_{t_i}[\phi^{n}(t_i', \mu_0)] - \epsilon.$$

SP-L vs. Related Concepts

- **SP-L**: for any t_i, t_i', m

$$u_{t_i}[\phi^{\infty}(t_i, m)] \geq u_{t_i}[\phi^{\infty}(t_i', m)]$$

- **SP-L (alt. statement)**: for any t_i, t_i', m, and any $\epsilon > 0$, there exists n_0 such that if $n > n_0$ we have

$$u_{t_i}[\phi^{n}(t_i, m)] \geq u_{t_i}[\phi^{n}(t_i', m)] - \epsilon.$$

- **Approximate Bayes-Nash**: for the true prior $\mu_0 \in \Delta T$, any t_i, t_i', and any $\epsilon > 0$, there exists n_0 such that if $n > n_0$ we have

$$u_{t_i}[\phi^{n}(t_i, \mu_0)] \geq u_{t_i}[\phi^{n}(t_i', \mu_0)] - \epsilon.$$

- **Approximate SP**: for any t_i, t_i', and any $\epsilon > 0$, there exists n_0 such that if $n > n_0$, for any $t_{-i} \in T^{n-1}$, we have

$$u_{t_i}[\Phi_i^{n}(t_i, t_{-i})] \geq u_{t_i}[\Phi_i^{n}(t_i', t_{-i})] - \epsilon.$$

Formal Appeal of SP-L

SP-L: for any $t_i,\ t_i',\ m:\ u_{t_i}[\phi^\infty(t_i, m)] \geq u_{t_i}[\phi^\infty(t_i', m)]$.

1. Wilson Doctrine
   - Bergemann and Morris (2005): agents' behavior and hence mechanism outcomes should be insensitive to beliefs
   - SP mechanisms comply exactly: truthful play is exactly optimal for any beliefs
   - SP-L mechanisms comply approximately: truthful play is approximately optimal for a wide range of beliefs
2. Strategic Simplicity
   - For any full-support beliefs m, and any cost $c > 0$ of calculating an optimal response, in a large enough market it is optimal to simply report truthfully and avoid the cost c
3. Fairness to Unsophisticated Players
   - For any full-support distribution of play m, and any cost $c > 0$, in a large enough market the cost of being non-sophisticated, and just reporting truthfully, is less than c

Roadmap

- Introduction
- Environment
- Strategyproof in the Large
- **Classification of non-SP Mechanisms**
- Constructing SP-L Mechanisms
- Discussion and Extensions
- Conclusion

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

Plan

1. Go over multi-unit auctions example in detail
2. Briefly describe how we obtain the rest of the classification
3. Relate to extant theory literature on large markets
4. Relate to extant empirics literature on manipulability

Example: Multi-Unit Auctions

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

Example: Multi-Unit Auctions

- Neither uniform-price nor pay-as-bid auction is SP (Ausubel and Cramton, 2002)
- We will show that uniform-price auction is SP-L, while pay-as-bid is Manipulable in the Large
- Example illustrates several aspects of the definition of SP-L, in particular
  - Ex-interim as opposed to ex-post perspective to manipulations
  - Full support assumption

Example: Multi-Unit Auctions

Basic setup:

- There are kn units of a homogeneous good, with $k \in \mathbb{Z}_+$.
- To simplify notation, we assume that agents' preferences take the form of linear utility functions, up to a capacity limit.
  - Specifically, each agent i's type t_i consists of a per-unit value v_i and a maximum capacity q_i, with $V = \{1, \ldots, \bar{v}\}$ the set of possible values, $Q = \{1, \ldots, \bar{q}\}$ the set of possible capacity limits, and $T = V \times Q$.
  - Set of outcomes is given by $X_0 = (V \times Q) \cup \emptyset$, with an outcome consisting of a per-unit payment and quantity.
- Bids consist of a per-unit value and a max capacity, so the action set $A = T$

Uniform-Price Auction: Finite

- Given a vector of n reports t, let $D(p;t)$ denote demand at price p. The market clearing price is:

$$p^*(t) = \max\{p \in V : \frac{D(p;t)}{n} \geq k\}.$$

Uniform-Price Auction: Finite

- Given a vector of n reports t, let $D(p;t)$ denote demand at price p. The market clearing price is:

$$p^*(t) = \max\{p \in V : \frac{D(p;t)}{n} \geq k\}.$$

- The uniform-price auction allocates each bidder i her demanded quantity at $p^*(t)$, with the exception that bids with $v_i = p^*(t)$ are rationed with equal probability. Formally, $\Phi_i^n(t)$ allocates according to

Reported Value	Expected Number of Units
$v_i < p^*(t)$	0
$v_i = p^*(t)$	$\bar{r} \cdot q_i$
$v_i > p^*(t)$	q_i

and all winning bidders pay $p^*(t)$ per unit they receive. The rationing constant $\bar{r}$ is set to clear the market

Uniform-Price Auction: Large-Market Limit

Uniform-Price Auction: Large-Market Limit

- Let $\rho^*(m)$ denote the price that clears supply and *average demand* given bid distribution $m \in \Delta T$:

$$\rho^*(m) = \max\{p \in V : E[D(p; t_i) | t_i \sim m] \geq k\}.$$

Uniform-Price Auction: Large-Market Limit

- Let $\rho^*(m)$ denote the price that clears supply and *average demand* given bid distribution $m \in \Delta T$:

$$\rho^*(m) = \max\{p \in V : E[D(p; t_i)|t_i \sim m] \geq k\}.$$

- **Generic case**: expected demand at price $\rho^*(m)$ strictly exceeds supply, that is,

$$E[D(\rho^*(m); t_i)|t_i \sim m] > k.$$

- In this case, as the market grows large, the realized price will be equal to $\rho^*(m)$ with probability converging to one. Limit mechanism allocates each bidder their demand at $\rho^*(m)$, again with rationing. Formally, $\phi^\infty(t_i, m)$ gives player i

Reported Value	Expected Number of Units
$v_i < \rho^*(m)$	0
$v_i = \rho^*(m)$	$\bar{r} \cdot q_i$
$v_i > \rho^*(m)$	q_i

at a per unit price of $\rho^*(m)$.

Knife-Edge Case

- There is also a "knife-edge" case, in which expected demand at price $\rho^*(m)$ is exactly equal to supply, that is,

$$E[D(\rho^*(m); t_i)|t_i \sim m] = k$$

- In this case, price is stochastic even in the large-market limit
- Given large n, the realized per-capita demand at price $\rho^*(m)$ is
  - weakly greater than per-capita supply k with probability of about $\frac{1}{2}$
  - strictly smaller than per-capita supply k with probability of about $\frac{1}{2}$.
- Therefore, the price in our limit will be $\rho^*(m)$ with probability of $\frac{1}{2}$, and $\rho^*(m) - 1$ with probability of $\frac{1}{2}$.
- Key point: even though the price is sometimes $\rho^*(m)$ and sometimes $\rho^*(m) - 1$, the probability that bidder i is pivotal converges to zero

The Uniform-Price Auction is SP-L

- The argument that the uniform-price auction is SP-L is now straightforward
  - Choose any type t_i, and any full support distribution m
  - The description of ϕ^{∞} above implies truthful reporting is a dominant strategy in our limit
- Note: argument would not go through if
  - We required that truthful reporting is optimal for any *realization* of opponent reports
  - We relaxed full support (Swinkels 2001 example)

Pay-as-Bid Auction

Pay-as-Bid Auction

- Exactly the same allocation function as uniform-price auction
- Difference: winning bidders pay their bid, not the market clearing price

Pay-as-Bid Auction

- Exactly the same allocation function as uniform-price auction
- Difference: winning bidders pay their bid, not the market clearing price
- Clearly, bidders benefit from misreporting, even in the limit
  - Given distribution m, limit price $\rho^*(m)$, a bidder of type $t_i = (v_i, q_i)$ with $v_i > \rho^*(m) + 1$ strictly prefers to misreport as $t_i' = (\rho^*(m) + 1, q_i)$

Pay-as-Bid Auction

- Exactly the same allocation function as uniform-price auction
- Difference: winning bidders pay their bid, not the market clearing price
- Clearly, bidders benefit from misreporting, even in the limit
  - Given distribution m, limit price $\rho^*(m)$, a bidder of type $t_i = (v_i, q_i)$ with $v_i > \rho^*(m) + 1$ strictly prefers to misreport as $t_i' = (\rho^*(m) + 1, q_i)$
- Hence, the pay-as-bid auction is not SP-L.

Pay-as-Bid Auction

- Exactly the same allocation function as uniform-price auction
- Difference: winning bidders pay their bid, not the market clearing price
- Clearly, bidders benefit from misreporting, even in the limit
  - Given distribution m, limit price $\rho^*(m)$, a bidder of type $t_i = (v_i, q_i)$ with $v_i > \rho^*(m) + 1$ strictly prefers to misreport as $t_i' = (\rho^*(m) + 1, q_i)$
- Hence, the pay-as-bid auction is not SP-L.
- Notice that t_i's optimal misreport depends on m, and that an unsophisticated bidder who bids truthfully can suffer a large loss relative to optimal behavior.

Pay-as-Bid Auction

- Exactly the same allocation function as uniform-price auction
- Difference: winning bidders pay their bid, not the market clearing price
- Clearly, bidders benefit from misreporting, even in the limit
  - Given distribution m, limit price $\rho^*(m)$, a bidder of type $t_i = (v_i, q_i)$ with $v_i > \rho^*(m) + 1$ strictly prefers to misreport as $t_i' = (\rho^*(m) + 1, q_i)$
- Hence, the pay-as-bid auction is not SP-L.
- Notice that t_i's optimal misreport depends on m, and that an unsophisticated bidder who bids truthfully can suffer a large loss relative to optimal behavior.
- Hence, in contrast to the uniform-price auction, the pay-as-bid auction is neither strategically simple nor fair to unsophisticated bidders.

Friedman (1991): "you do not have to be a specialist" to figure out how to participate in the uniform price auction, because you can just indicate "the maximum amount you are willing to pay for different quantities ... if you bid a higher price [than the market clearing price], you do not lose as you do under the current [pay-as-bid] method."

Obtaining the Classification

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

Obtaining the Classification

- To show that a mechanism is not SP-L: suffices to produce an example of a profitable manipulation in the limit, as we did for pay-as-bid
  - 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 limit, as we did for pay-as-bid
  - 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

Obtaining the Classification

- To show that a mechanism is not SP-L: suffices to produce an example of a profitable manipulation in the limit, as we did for pay-as-bid
  - 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

Obtaining the Classification

- To show that a mechanism is not SP-L: suffices to produce an example of a profitable manipulation in the limit, as we did for pay-as-bid
  - 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→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

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)

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$\rightarrow$SP-L

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→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:

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)

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

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
  - Our analysis does not yield an understanding of the exact forces pushing away from truthful behavior in finite markets (e.g., first-order condition analysis of RSW)

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)

Relationship to Empirical Literature on Manipulability

- By contrast, to the best of our knowledge, there are no empirical examples where an SP-L market design is shown to be harmfully manipulated in a large market
- Evidence to date thus suggests that the relevant distinction for practice is SP-L vs. not SP-L, rather than SP vs. not SP.
  - More conservatively: SP vs. SP-L vs. not SP-L
  - Caution: no evidence one way or the other for many of the SP-L mechanisms in the table

Roadmap

- Introduction
- Environment
- Strategyproof in the Large
- Classification of non-SP Mechanisms
- **Constructing SP-L Mechanisms**
- Discussion and Extensions
- Conclusion

Constructing SP-L Mechanisms from Bayes-Nash Mechanisms

- Goal: show that, in large markets, SP-L is in a well-defined sense approximately costless to impose relative to Bayes-Nash (or Nash) incentive compatibility
- Together with the analysis above that suggests that SP-L is attractive, completes our argument that SP-L is a useful second best to SP

Plan

1. Give the construction
2. Quasi-continuity
3. State construction theorem, and sketch proof
4. Discussion

Construction: Preliminaries

- Definition: A **Limit Bayes-Nash Equilibrium** at a given prior $\mu \in \Delta T$ is a strategy $\sigma^* : T \to \Delta A$ such that for all t_i, a_i

$$u_{t_i}[\phi^\infty(\sigma^*(t_i), \sigma^*(\mu))] \geq u_{t_i}[\phi^\infty(a_i, \sigma^*(\mu))]$$

  - where $\sigma^*(\mu)$ is the distribution of actions given iid draws according to μ and play according to $\sigma^*(\cdot)$

Construction: Preliminaries

- Definition: A **Limit Bayes-Nash Equilibrium** at a given prior $\mu \in \Delta T$ is a strategy $\sigma^* : T \to \Delta A$ such that for all t_i, a_i

$$u_{t_i}[\phi^\infty(\sigma^*(t_i), \sigma^*(\mu))] \geq u_{t_i}[\phi^\infty(a_i, \sigma^*(\mu))]$$

  - where $\sigma^*(\mu)$ is the distribution of actions given iid draws according to μ and play according to $\sigma^*(\cdot)$

- Definition: A **family of limit equilibria** $(\sigma^*_\mu)_{\mu \in \Delta T}$ specifies a limit BNE for each prior $\mu \in \Delta T$.

Construction: Preliminaries

- Definition: A **Limit Bayes-Nash Equilibrium** at a given prior $\mu \in \Delta T$ is a strategy $\sigma^* : T \to \Delta A$ such that for all t_i, a_i

$$u_{t_i}[\phi^\infty(\sigma^*(t_i), \sigma^*(\mu))] \geq u_{t_i}[\phi^\infty(a_i, \sigma^*(\mu))]$$

  - where $\sigma^*(\mu)$ is the distribution of actions given iid draws according to μ and play according to $\sigma^*(\cdot)$

- Definition: A **family of limit equilibria** $(\sigma^*_\mu)_{\mu \in \Delta T}$ specifies a limit BNE for each prior $\mu \in \Delta T$.
- Notation
  - linearly extend the definition of a mechanism from action vectors to distributions of action vectors. Given $\bar{m} \in \Delta(A^n)$, let

$$\Phi^n(\bar{m}) = \sum_a \Phi^n(a) \cdot \bar{m}(a).$$

  - given a type vector t, let $\text{emp}[t]$ denote its empirical distribution on T.

Construction

- Input: a mechanism $\{(\Phi_n)_{\mathbb{N}}, A\}$ and a family of limit equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$

Construction

- Input: a mechanism $\{(\Phi_n)_{\mathbb{N}}, A\}$ and a family of limit equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$
- Construct a new direct mechanism, $\{(F_n)_{\mathbb{N}}, T\}$ according to:

$$F^n(t) = \Phi^n(\sigma^*_{\text{emp}[t]}(t)).$$

- In words: $F^n(\cdot)$ acts as a proxy playing the original mechanism on each agent's behalf, plays $\sigma^*_{\text{emp}[t]}(t_i)$ on behalf of t_i

Construction

- Input: a mechanism $\{(\Phi_n)_{\mathbb{N}}, A\}$ and a family of limit equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$
- Construct a new direct mechanism, $\{(F_n)_{\mathbb{N}}, T\}$ according to:

$$F^n(t) = \Phi^n(\sigma^*_{\text{emp}[t]}(t)).$$

- In words: $F^n(\cdot)$ acts as a proxy playing the original mechanism on each agent's behalf, plays $\sigma^*_{\text{emp}[t]}(t_i)$ on behalf of t_i
- Key feature: rather than use the Bayes-Nash equilibrium strategy associated with the "true prior", which need not be known by the designer, it uses the strategy, $\sigma^*_{\text{emp}[t]}(\cdot)$ associated with the empirical distribution of reports
  - Proxy: "I do not know the distribution of preferences, and presumably neither do you. But whatever the distribution happens to be, I will play the Bayes-Nash strategy on your behalf."
  - Note: in finite markets, i's report affects emp$[t]$, and hence what BNE gets "activated"

Quasi-Continuity

- We will show that $\{(F_n)_{\mathbb{N}}, T\}$ constructed according to $F^n(t) = \Phi^n(\sigma^*_{\text{emp}[t]}(t))$:
  - Is SP-L
  - Gives agents approximately the same outcomes as the BNE of the original mechanism $\{(\Phi_n)_{\mathbb{N}}, A\}$
- This result requires a continuity condition on $\{(\Phi_n)_{\mathbb{N}}, A\}$ and its equilibria, that we turn to next

Quasi-Continuity

A family of equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$ of mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ is **quasi-continuous** at full support prior $\mu_0 \in \Delta T$ if for every $\epsilon > 0$, there exists a neighborhood $\mathcal{N}$ of μ_0 that can be decomposed as $\mathcal{N} = \cup_{1 \leq k \leq K} \mathcal{A}_k \cup \mathcal{B}$ with each $\mathcal{A}_k$ open, such that:

Quasi-Continuity

A family of equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$ of mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ is **quasi-continuous** at full support prior $\mu_0 \in \Delta T$ if for every $\epsilon > 0$, there exists a neighborhood $\mathcal{N}$ of μ_0 that can be decomposed as $\mathcal{N} = \cup_{1 \leq k \leq K} \mathcal{A}_k \cup \mathcal{B}$ with each $\mathcal{A}_k$ open, such that:

1. If types are drawn iid according to μ_0, then the probability that the empirical distribution of types lands within distance $1/n$ of $\mathcal{B}$ goes to zero as n grows large. Formally,

$$\lim_{n \to \infty} \Pr\{\text{distance}(\text{emp}[t], \mathcal{B}) \leq 1/n | t \in T^n, t \sim \text{iid}(\mu_0)\} = 0.$$

Quasi-Continuity

2. Within each set $\mathcal{A}_k$, in a large enough market, agents' outcomes are continuous with respect to changes in the empirical distribution of opponents' types and the strategy that agents use.

   Formally, for each $\mathcal{A}_k$, there exists n_0 such that for any $n > n_0$, and any $\mu, \mu', \text{emp}[t_i, t_{-i}], \text{emp}[t_i, t'_{-i}] \in \mathcal{A}_k$, we have:

$$|\Phi_i^n(\sigma_\mu^*(t_i), \sigma_\mu^*(t_{-i})) - \Phi_i^n(\sigma_{\mu'}^*(t_i), \sigma_{\mu'}^*(t'_{-i}))| < \epsilon.$$

Quasi-Continuity

2. Within each set $\mathcal{A}_k$, in a large enough market, agents' outcomes are continuous with respect to changes in the empirical distribution of opponents' types and the strategy that agents use.

   Formally, for each $\mathcal{A}_k$, there exists n_0 such that for any $n > n_0$, and any $\mu, \mu', \text{emp}[t_i, t_{-i}], \text{emp}[t_i, t'_{-i}] \in \mathcal{A}_k$, we have:

$$|\Phi_i^n(\sigma_\mu^*(t_i), \sigma_\mu^*(t_{-i})) - \Phi_i^n(\sigma_{\mu'}^*(t_i), \sigma_{\mu'}^*(t'_{-i}))| < \epsilon.$$

- The family is **continuous** at μ_0 if $\mathcal{B} = \emptyset$ and $K = 1$.
- The family is (quasi-)continuous if it is (quasi-)continuous at every prior in $\bar{\Delta} T$.

Quasi-Continuity

2. Within each set $\mathcal{A}_k$, in a large enough market, agents' outcomes are continuous with respect to changes in the empirical distribution of opponents' types and the strategy that agents use.

   Formally, for each $\mathcal{A}_k$, there exists n_0 such that for any $n > n_0$, and any $\mu, \mu', \text{emp}[t_i, t_{-i}], \text{emp}[t_i, t'_{-i}] \in \mathcal{A}_k$, we have:

$$|\Phi_i^n(\sigma^*_\mu(t_i), \sigma^*_\mu(t_{-i})) - \Phi_i^n(\sigma^*_{\mu'}(t_i), \sigma^*_{\mu'}(t'_{-i}))| < \epsilon.$$

- The family is **continuous** at μ_0 if $\mathcal{B} = \emptyset$ and $K = 1$.
- The family is (quasi-)continuous if it is (quasi-)continuous at every prior in $\bar{\Delta} T$.
- N.B. our notion of continuous is stronger than e.g. Kalai's (2004): continuous w/r/t both reports *and strategies*

Quasi-Continuity

- Why quasi-continuity? Many discrete-goods allocation mechanisms have knife-edge discontinuities, and it is important that our main theorem include such mechanisms.

Quasi-Continuity

- Why quasi-continuity? Many discrete-goods allocation mechanisms have knife-edge discontinuities, and it is important that our main theorem include such mechanisms.
- Consider the uniform-price auction:

Quasi-Continuity

- Why quasi-continuity? Many discrete-goods allocation mechanisms have knife-edge discontinuities, and it is important that our main theorem include such mechanisms.
- Consider the uniform-price auction:
- Generic case
  - Expected demand at μ_0 strictly exceeds supply at the limit market-clearing price: $E[D(\rho^*(\mu_0); t_i)|t_i \sim \mu_0] > k$
  - Equilibrium is locally continuous

Quasi-Continuity

- Why quasi-continuity? Many discrete-goods allocation mechanisms have knife-edge discontinuities, and it is important that our main theorem include such mechanisms.
- Consider the uniform-price auction:
- Generic case
  - Expected demand at μ_0 strictly exceeds supply at the limit market-clearing price: $E[D(\rho^*(\mu_0); t_i)|t_i \sim \mu_0] > k$
  - Equilibrium is locally continuous
- Knife-edge case
  - Expected demand at μ_0 is exactly equal to supply at the limit market-clearing price: $E[D(\rho^*(\mu_0); t_i)|t_i \sim \mu_0] = k$.

Quasi-Continuity

- Why quasi-continuity? Many discrete-goods allocation mechanisms have knife-edge discontinuities, and it is important that our main theorem include such mechanisms.
- Consider the uniform-price auction:
- Generic case
  - Expected demand at μ_0 strictly exceeds supply at the limit market-clearing price: $E[D(\rho^*(\mu_0); t_i)|t_i \sim \mu_0] > k$
  - Equilibrium is locally continuous
- Knife-edge case
  - Expected demand at μ_0 is exactly equal to supply at the limit market-clearing price: $E[D(\rho^*(\mu_0); t_i)|t_i \sim \mu_0] = k$.
  - Take a small neighborhood $\mathcal{N}$ of μ_0 with:
    - $\mathcal{A}_1 = \{\mu \in \mathcal{N} : E[D(\rho^*(\mu_0); t_i)|t_i \sim \mu] > k\}$
    - $\mathcal{A}_2 = \{\mu \in \mathcal{N} : E[D(\rho^*(\mu_0); t_i)|t_i \sim \mu] < k\}$
    - $\mathcal{B} = \mathcal{N} \backslash (\mathcal{A}_1 \cup \mathcal{A}_2)$
  - Outcomes are continuous within the sets $\mathcal{A}_1$ and $\mathcal{A}_2$, and the likelihood of landing in the knife-edge set $\mathcal{B}$ goes to zero as $n \to \infty$.

Quasi-Continuity

- So, the uniform-price auction's family of equilibria is not continuous, but is quasi-continuous
- We show the same for pay-as-bid (see Appendix D)

Other Examples:

- Boston mechanism
  - Potential discontinuity: if a school reaches capacity exactly at the end of some round
- College admissions model
  - Potential discontinuity: students who are right at the cutoff for a particular college
- Bidding Points Auction
  - Potential discontinuity: students whose bid is right at the cutoff for a particular course

Construction Theorem

Construction Theorem

- Suppose we are given a non-SP-L mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ with a quasi-continuous family of Bayes-Nash equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$. Fix an arbitrary full support prior $\mu_0 \in \Delta T$
- Construct $\{(F_n)_{\mathbb{N}}, T\}$ according to:

$$F^n(t) = \Phi^n(\sigma^*_{\text{emp}[t]}(t)).$$

The constructed mechanism has the following properties:

Construction Theorem

- Suppose we are given a non-SP-L mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ with a quasi-continuous family of Bayes-Nash equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$. Fix an arbitrary full support prior $\mu_0 \in \Delta T$
- Construct $\{(F_n)_{\mathbb{N}}, T\}$ according to:

$$F^n(t) = \Phi^n(\sigma^*_{\mathrm{emp}[t]}(t)).$$

The constructed mechanism has the following properties:

1. $\{(F^n)_{\mathbb{N}}, T\}$ is SP-L.

Construction Theorem

- Suppose we are given a non-SP-L mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ with a quasi-continuous family of Bayes-Nash equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$. Fix an arbitrary full support prior $\mu_0 \in \Delta T$
- Construct $\{(F_n)_{\mathbb{N}}, T\}$ according to:

$$F^n(t) = \Phi^n(\sigma^*_{\text{emp}[t]}(t)).$$

The constructed mechanism has the following properties:

1. $\{(F^n)_{\mathbb{N}}, T\}$ is SP-L.
2. If $\{(\Phi^n)_{\mathbb{N}}, A\}$ is continuous at μ_0, then, in the limit as $n \to \infty$, truthful play of $\{(F^n)_{\mathbb{N}}, T\}$ gives agents the same outcomes as Bayes-Nash eqm play of $\{(\Phi^n)_{\mathbb{N}}, A\}$ for prior μ_0.

Construction Theorem

- Suppose we are given a non-SP-L mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ with a quasi-continuous family of Bayes-Nash equilibria $(\sigma^*_\mu)_{\mu\in\Delta T}$. Fix an arbitrary full support prior $\mu_0 \in \Delta T$
- Construct $\{(F_n)_{\mathbb{N}}, T\}$ according to:

$$F^n(t) = \Phi^n(\sigma^*_{\mathrm{emp}[t]}(t)).$$

The constructed mechanism has the following properties:

1. $\{(F^n)_{\mathbb{N}}, T\}$ is SP-L.
2. If $\{(\Phi^n)_{\mathbb{N}}, A\}$ is continuous at μ_0, then, in the limit as $n \to \infty$, truthful play of $\{(F^n)_{\mathbb{N}}, T\}$ gives agents the same outcomes as Bayes-Nash eqm play of $\{(\Phi^n)_{\mathbb{N}}, A\}$ for prior μ_0.
3. If $\{(\Phi^n)_{\mathbb{N}}, A\}$ is not continuous at μ_0 then, in the limit, $\{(F^n)_{\mathbb{N}}, T\}$ coincides with a convex combination of BNE outcomes of $\{(\Phi^n)_{n\in\mathbb{N}}, A\}$, for priors in an arbitrarily small neighborhood of μ_0.

Construction Theorem

- Suppose we are given a non-SP-L mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ with a quasi-continuous family of Bayes-Nash equilibria $(\sigma^*_\mu)_{\mu \in \Delta T}$. Fix an arbitrary full support prior $\mu_0 \in \Delta T$
- Construct $\{(F_n)_{\mathbb{N}}, T\}$ according to:

$$F^n(t) = \Phi^n(\sigma^*_{\mathrm{emp}[t]}(t)).$$

The constructed mechanism has the following properties:

1. $\{(F^n)_{\mathbb{N}}, T\}$ is SP-L.
2. If $\{(\Phi^n)_{\mathbb{N}}, A\}$ is continuous at μ_0, then, in the limit as $n \to \infty$, truthful play of $\{(F^n)_{\mathbb{N}}, T\}$ gives agents the same outcomes as Bayes-Nash eqm play of $\{(\Phi^n)_{\mathbb{N}}, A\}$ for prior μ_0.
3. If $\{(\Phi^n)_{\mathbb{N}}, A\}$ is not continuous at μ_0 then, in the limit, $\{(F^n)_{\mathbb{N}}, T\}$ coincides with a convex combination of BNE outcomes of $\{(\Phi^n)_{n \in \mathbb{N}}, A\}$, for priors in an arbitrarily small neighborhood of μ_0.

Takeaway: SP-L is approximately costless relative to Bayes-Nash

Sketch of Proof: Step 1

- Suppose that the other agents report truthfully, according to the true prior μ_0

Sketch of Proof: Step 1

- Suppose that the other agents report truthfully, according to the true prior μ_0
- In a finite market there will be sampling error, so the realized empirical will be, say $\hat{\mu}$

Sketch of Proof: Step 1

- Suppose that the other agents report truthfully, according to the true prior μ_0
- In a finite market there will be sampling error, so the realized empirical will be, say $\hat{\mu}$
- As $n \to \infty$, the realized empirical $\hat{\mu}$ converges to μ_0, by the law of large numbers.

Sketch of Proof: Step 1

- Suppose that the other agents report truthfully, according to the true prior μ_0
- In a finite market there will be sampling error, so the realized empirical will be, say $\hat{\mu}$
- As $n \to \infty$, the realized empirical $\hat{\mu}$ converges to μ_0, by the law of large numbers.
- Assume for now that the original mechanism is continuous at μ_0.

Sketch of Proof: Step 1

- Suppose that the other agents report truthfully, according to the true prior μ_0
- In a finite market there will be sampling error, so the realized empirical will be, say $\hat{\mu}$
- As $n \to \infty$, the realized empirical $\hat{\mu}$ converges to μ_0, by the law of large numbers.
- Assume for now that the original mechanism is continuous at μ_0.
- Then agent i's allocation, $\Phi_i^n(\sigma_{\hat{\mu}}^*(t_i), \sigma_{\hat{\mu}}^*(t_{-i}))$ is converging to $\phi^\infty(\sigma_{\mu_0}^*(t_i), \sigma_{\mu_0}^*(\mu_0))$, exactly what he receives under the limit Bayes-Nash equilibrium of the original mechanism.

Sketch of Proof: Step 1

- Suppose that the other agents report truthfully, according to the true prior μ_0
- In a finite market there will be sampling error, so the realized empirical will be, say $\hat{\mu}$
- As $n \to \infty$, the realized empirical $\hat{\mu}$ converges to μ_0, by the law of large numbers.
- Assume for now that the original mechanism is continuous at μ_0.
- Then agent i's allocation, $\Phi_i^n(\sigma_{\hat{\mu}}^*(t_i), \sigma_{\hat{\mu}}^*(t_{-i}))$ is converging to $\phi^\infty(\sigma_{\mu_0}^*(t_i), \sigma_{\mu_0}^*(\mu_0))$, exactly what he receives under the limit Bayes-Nash equilibrium of the original mechanism.
- Thus, if all agents report truthfully, our mechanism coincides with the original mechanism in the limit, as required.

Sketch of Proof: Step 2

- Now, suppose that the agents other than i misreport their preferences, according to some distribution $m \in \Delta T$.

Sketch of Proof: Step 2

- Now, suppose that the agents other than i misreport their preferences, according to some distribution $m \in \Delta T$.
- As before, in a finite market of size n, there will be sampling error, so the realized empirical will be, say, $\hat{m}$.

Sketch of Proof: Step 2

- Now, suppose that the agents other than i misreport their preferences, according to some distribution $m \in \Delta T$.
- As before, in a finite market of size n, there will be sampling error, so the realized empirical will be, say, $\hat{m}$.
- As $n \to \infty$, $\hat{m} \to m$. Assume continuity at m.

Sketch of Proof: Step 2

- Now, suppose that the agents other than i misreport their preferences, according to some distribution $m \in \Delta T$.
- As before, in a finite market of size n, there will be sampling error, so the realized empirical will be, say, $\hat{m}$.
- As $n \to \infty$, $\hat{m} \to m$. Assume continuity at m.
- Then agent i's allocation, $\Phi_i^n(\sigma_{\hat{m}}^*(t_i), \sigma_{\hat{m}}^*(t_{-i}))$ is converging to $\phi^\infty(\sigma_m^*(t_i), \sigma_m^*(m))$, exactly what he receives under the original mechanism *in the Bayes-Nash equilibrium corresponding to prior m.*

Sketch of Proof: Step 2

- Now, suppose that the agents other than i misreport their preferences, according to some distribution $m \in \Delta T$.
- As before, in a finite market of size n, there will be sampling error, so the realized empirical will be, say, $\hat{m}$.
- As $n \to \infty$, $\hat{m} \to m$. Assume continuity at m.
- Then agent i's allocation, $\Phi_i^n(\sigma_{\hat{m}}^*(t_i), \sigma_{\hat{m}}^*(t_{-i}))$ is converging to $\phi^\infty(\sigma_m^*(t_i), \sigma_m^*(m))$, exactly what he receives under the original mechanism *in the Bayes-Nash equilibrium corresponding to prior m.*
- Key point: even though the other agents are systematically misreporting their preferences, our agent i remains happy to tell the truth!
  - Our agent doesn't care that the others are misreporting: it's as if he's in the m world not the μ_0 world. Still wants to report truthfully, hence SP-L
  - Note role of private values assumption

Sketch of Proof: Step 3

- The last step is to describe what happens in the event that the equilibrium of the original mechanism is not continuous at μ_0.
- This requires a technical lemma, informally:
  - For any arbitrary full support prior $m \in \Delta T$, the allocation an agent receives under $\{(F^n)_{\mathbb{N}}, T\}$ can be approximated by a convex combination of the allocations he would receive in the limit Bayes-Nash equilibria of $\{(\Phi^n)_{\mathbb{N}}, A\}$, for priors arbitrarily close to m.
- Key to the proof: in a large enough market, a single agent cannot appreciably change the probability that the aggregate profile lands in each region $\mathcal{A}_k$
  - This allows us to exploit the continuity *within* each region $\mathcal{A}_k$, and the vanishing likelihood that the aggregate profile lands near the discontinuity region $\mathcal{B}$.

Roadmap

- Introduction
- Environment
- Strategyproof in the Large
- Classification of non-SP Mechanisms
- Constructing SP-L Mechanisms
- **Discussion and Extensions**
- Conclusion

Discussion of Theorem 1

1. Extensions of the theorem
2. Relation to the revelation principle
3. Relation to previous BNE-SP equivalences
4. Application to the ongoing debate re the Boston mechanism

Extensions of Theorem 1

- Semi-Anonymity
  - Punchline: results can be generalized from anonymous to semi-anonymous mechanisms (Kalai, 2004)
- Can use complete-information Nash equilibria as our input, instead of Bayes-Nash equilibria
  - Constructed mechanism becomes: agents report their types, then compute and execute the CINE associated with the reported types
  - Neat feature: exactly coincides with the original mechanism if everyone reports truthfully
    - N.B.: not a Nash eqm to report truthfully in finite markets: i's report affects emp[t] which affects j's strategy as activated by the proxy
- Can use finite economy Bayes-Nash equilibria as our input, instead of limit BNE

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i
  - Reporting truthfully is a BNE

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i
  - Reporting truthfully is a BNE
- Our mechanism

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i
  - Reporting truthfully is a BNE
- Our mechanism
  - Mechanism does not know the prior

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i
  - Reporting truthfully is a BNE
- Our mechanism
  - Mechanism does not know the prior
  - Instead *infers* a prior from the empirical

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i
  - Reporting truthfully is a BNE
- Our mechanism
  - Mechanism does not know the prior
  - Instead *infers* a prior from the empirical
  - If agents report truthfully, inference is correct in the limit

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i
  - Reporting truthfully is a BNE
- Our mechanism
  - Mechanism does not know the prior
  - Instead *infers* a prior from the empirical
  - If agents report truthfully, inference is correct in the limit
  - If agents misreport, so empirical $\hat{m}$ is very different from prior μ_0, our mechanism automatically adjusts to play $\sigma^*_{\hat{m}}$

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i
  - Reporting truthfully is a BNE
- Our mechanism
  - Mechanism does not know the prior
  - Instead *infers* a prior from the empirical
  - If agents report truthfully, inference is correct in the limit
  - If agents misreport, so empirical $\hat{m}$ is very different from prior μ_0, our mechanism automatically adjusts to play $\sigma^*_{\hat{m}}$
  - Hence, we get SP-L not just Bayes Nash

Relation to the Revelation Principle

- Traditional Revelation Principle (Myerson, 1979)
  - Mechanism knows the prior, μ_0
  - Announces the BNE associated with that prior, $\sigma^*_{\mu_0}(\cdot)$.
  - Mechanism plays $\sigma^*_{\mu_0}(t_i)$ on behalf of agent who reports t_i
  - Reporting truthfully is a BNE
- Our mechanism
  - Mechanism does not know the prior
  - Instead *infers* a prior from the empirical
  - If agents report truthfully, inference is correct in the limit
  - If agents misreport, so empirical $\hat{m}$ is very different from prior μ_0, our mechanism automatically adjusts to play $\sigma^*_{\hat{m}}$
  - Hence, we get SP-L not just Bayes Nash
  - And, our mechanism is prior free, consistent with Wilson doctrine

Relation to Previous BNE-SP Equivalences

- Our theorem can be understood in relation to Manelli and Vincent (2010) and Gershkov et al (forth.), who find striking *exact* equivalences between Bayes-Nash and SP in finite markets
  - Myerson (1981): in optimal auction problem, no gap between BNE and SP
  - Manelli and Vincent (2010): in Myerson's environment, equivalence between BNE and SP obtains for any BNIC mechanism, not just revenue maximizing
  - Gershkov et al. (forth.): broaden to other mechanism design environments, still with 1-D types and quasi-linear utility
- Limitation of these results: one-dimensional types, quasi-linear utility.
  - Rules out all mechanisms in Table 1
- Our Theorem 1 recovers approximate equivalence between BNE and SP, for a rich enough class environments to include multi-object auctions, matching, assignment, school choice, etc.

Discussion of Theorem 1: Boston Mechanism

Generation 1:

Discussion of Theorem 1: Boston Mechanism

Generation 1:

- Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al (2006) criticized the Boston Mechanism on the grounds that it is not strategyproof

Discussion of Theorem 1: Boston Mechanism

Generation 1:

- Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al (2006) criticized the Boston Mechanism on the grounds that it is not strategyproof
- Proposed that strategyproof Gale-Shapley deferred acceptance be used instead

Discussion of Theorem 1: Boston Mechanism

Generation 1:

- Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al (2006) criticized the Boston Mechanism on the grounds that it is not strategyproof
- Proposed that strategyproof Gale-Shapley deferred acceptance be used instead
- GS eventually adopted for use in practice, in Boston, NYC, others

Discussion of Theorem 1: Boston Mechanism

Generation 1:

- Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al (2006) criticized the Boston Mechanism on the grounds that it is not strategyproof
- Proposed that strategyproof Gale-Shapley deferred acceptance be used instead
- GS eventually adopted for use in practice, in Boston, NYC, others

Generation 2:

- Abdulkadiroğlu, Che and Yasuda (2011), Miralles (2008), Featherstone and Niederle (2009) argue that maybe Gen 1 was too quick to dismiss the Boston mechanism

Discussion of Theorem 1: Boston Mechanism

Generation 1:

- Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al (2006) criticized the Boston Mechanism on the grounds that it is not strategyproof
- Proposed that strategyproof Gale-Shapley deferred acceptance be used instead
- GS eventually adopted for use in practice, in Boston, NYC, others

Generation 2:

- Abdulkadiroğlu, Che and Yasuda (2011), Miralles (2008), Featherstone and Niederle (2009) argue that maybe Gen 1 was too quick to dismiss the Boston mechanism
- Boston has (family of) Bayes Nash equilibria that yield greater welfare than does the dominant strategy equilibrium of Gale-Shapley

Discussion of Theorem 1: Boston Mechanism

Generation 1:

- Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al (2006) criticized the Boston Mechanism on the grounds that it is not strategyproof
- Proposed that strategyproof Gale-Shapley deferred acceptance be used instead
- GS eventually adopted for use in practice, in Boston, NYC, others

Generation 2:

- Abdulkadiroğlu, Che and Yasuda (2011), Miralles (2008), Featherstone and Niederle (2009) argue that maybe Gen 1 was too quick to dismiss the Boston mechanism
- Boston has (family of) Bayes Nash equilibria that yield greater welfare than does the dominant strategy equilibrium of Gale-Shapley
- Interpretation: strategyproofness has a cost

Discussion of Theorem 1: Boston Mechanism

Our Paper (Gen 3):

Discussion of Theorem 1: Boston Mechanism

Our Paper (Gen 3):

- Bayes-Nash equilibria have costs too

Discussion of Theorem 1: Boston Mechanism

Our Paper (Gen 3):

- Bayes-Nash equilibria have costs too
    - Students must have common knowledge of preference distribution
    - Coordinate on a specific equilibrium
    - Make very precise calculations to determine whether to risk asking for a popular school
    - Empirical record suggests these costs are important in practice

Discussion of Theorem 1: Boston Mechanism

Our Paper (Gen 3):

- Bayes-Nash equilibria have costs too
  - Students must have common knowledge of preference distribution
  - Coordinate on a specific equilibrium
  - Make very precise calculations to determine whether to risk asking for a popular school
  - Empirical record suggests these costs are important in practice
- Our main result says that all of this is unnecessary in the large market limit: there must exist yet another mechanism that implements the same outcomes as the attractive BNE equilibria, yet with dominant strategy incentives

Roadmap

- Introduction
- Environment
- Strategyproof in the Large
- Classification of non-SP Mechanisms
- Constructing SP-L Mechanisms
- Discussion and Extensions
- **Conclusion**

Summary

Summary

- This paper proposes SP-L as a second-best alternative to SP:

Summary

- This paper proposes 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

Summary

- This paper proposes 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
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

Summary

- This paper proposes 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
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

Summary

- This paper proposes 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
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

- Our analysis provides formal justification for focusing on SP-L when confronting a new market design problem for which there are no good SP solutions

Summary

- This paper proposes 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
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

- Our analysis provides 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 (real-life implementation at Wharton planned for Fall 2013)

Caveat

Caveat

- No simple bright-line answer to the question of how large is large enough
  - This is true even in most studies of the convergence properties of specific mechanisms
  - Convergence is often slow, or has a large constant term
  - Exception: double auctions (e.g., Rustichini, Satterthwaite and Williams 1994)

Caveat

- No simple bright-line answer to the question of how large is large enough
  - This is true even in most studies of the convergence properties of specific mechanisms
  - Convergence is often slow, or has a large constant term
  - Exception: double auctions (e.g., Rustichini, Satterthwaite and Williams 1994)
- We view the limit as a frequently useful (always imperfect) abstraction
  - Just as the assumption of price-taking behavior is a useful if imperfect abstraction in some other parts of economics
  - In any specific context, the analyst's case that the market is large can be enhanced with computational simulations, empirical analysis, etc.

Caveat

- No simple bright-line answer to the question of how large is large enough
  - This is true even in most studies of the convergence properties of specific mechanisms
  - Convergence is often slow, or has a large constant term
  - Exception: double auctions (e.g., Rustichini, Satterthwaite and Williams 1994)
- We view the limit as a frequently useful (always imperfect) abstraction
  - Just as the assumption of price-taking behavior is a useful if imperfect abstraction in some other parts of economics
  - In any specific context, the analyst's case that the market is large can be enhanced with computational simulations, empirical analysis, etc.
- In environments where this abstraction is compelling: **consider designing a mechanism that is SP-L!**