Matching “versus” Mechanism Design

Eric Budish
University of Chicago, Booth School of Business

Matching in Practice Conference
ULB, May 2013

Matching “versus” mechanism design

Matching "versus" mechanism design

Mechanism design approach

- Max objective s.t. constraints (technology, incentives)
- Vickrey auction
- Myerson auction

Matching “versus” mechanism design

Mechanism design approach

- Max objective s.t. constraints (technology, incentives)
- Vickrey auction
- Myerson auction

Matching approach

- Seek a mechanism that satisfies "good properties"
- Gale-Shapley deferred acceptance algorithm
- Gale’s Top Trading Cycles algorithm

Matching “versus” mechanism design

These two approaches aren’t always so different

Matching “versus” mechanism design

These two approaches aren’t always so different

- Vickrey’s auction has an axiomatic formulation
  - unique mechanism that is Pareto efficient and strategyproof

Matching "versus" mechanism design

These two approaches aren't always so different

- Vickrey's auction has an axiomatic formulation
  - unique mechanism that is Pareto efficient and strategyproof
- Deferred acceptance can be formulated as a constrained optimization problem
  - Maximize proposer-side welfare s.t. stability constraints
  - G-S have a section in their paper on "optimality" that explicitly makes this point

Matching "versus" mechanism design

These two approaches aren't always so different

- Vickrey's auction has an axiomatic formulation
  - unique mechanism that is Pareto efficient and strategyproof
- Deferred acceptance can be formulated as a constrained optimization problem
  - Maximize proposer-side welfare s.t. stability constraints
  - G-S have a section in their paper on "optimality" that explicitly makes this point

But, of course, often times these approaches end up looking quite different. (Else, Rakesh wouldn't have suggested this topic!)

Matching "versus" mechanism design

And the winner is

Matching "versus" mechanism design

And the winner is ... MATCHING

Matching “versus” mechanism design

And the winner is ... MATCHING

Actually, if anything, the reverse.

Matching “versus” mechanism design

And the winner is ... MATCHING

Actually, if anything, the reverse. Of course what we want to do as economists is maximize design objectives subject to constraints.

Matching "versus" mechanism design

And the winner is ... MATCHING

Actually, if anything, the reverse. Of course what we want to do as economists is maximize design objectives subject to constraints.

So, why don't all matching papers look like mechanism design papers? A few reasons

Matching “versus” mechanism design

And the winner is ... MATCHING

Actually, if anything, the reverse. Of course what we want to do as economists is maximize design objectives subject to constraints.

So, why don’t all matching papers look like mechanism design papers? A few reasons

1. It is sometimes difficult to pin down *the* objective.

Matching "versus" mechanism design

And the winner is ... MATCHING

Actually, if anything, the reverse. Of course what we want to do as economists is maximize design objectives subject to constraints.

So, why don't all matching papers look like mechanism design papers? A few reasons

1. It is sometimes difficult to pin down *the* objective.
2. Sometimes difficulty to pin down the true constraints.

Matching “versus” mechanism design

And the winner is ... MATCHING

Actually, if anything, the reverse. Of course what we want to do as economists is maximize design objectives subject to constraints.

So, why don’t all matching papers look like mechanism design papers? A few reasons

1. It is sometimes difficult to pin down *the* objective.
2. Sometimes difficulty to pin down the true constraints.
3. Lack of tools. Main difficulty: *all* objects in the economy are indivisible, no numeraire

Matching "versus" mechanism design

And the winner is ... MATCHING

Actually, if anything, the reverse. Of course what we want to do as economists is maximize design objectives subject to constraints.

So, why don't all matching papers look like mechanism design papers? A few reasons

1. It is sometimes difficult to pin down *the* objective.
2. Sometimes difficulty to pin down the true constraints.
3. Lack of tools. Main difficulty: *all* objects in the economy are indivisible, no numeraire

For all of these reasons it can be useful to find one or more good solutions to the problem, rather than pursue the optimal solution.

Keep in mind: Myerson, Vickrey ... these are the ones that worked!

- If only all problems had such elegant and compelling solutions.

Matching “versus” mechanism design: plan of talk

Matching "versus" mechanism design: plan of talk

1. Introductory Example: School Choice
   - Simple win for the good properties approach

Matching "versus" mechanism design: plan of talk

1. Introductory Example: School Choice
   - Simple win for the good properties approach
2. Main example: Course Allocation (aka combinatorial assignment)
   - Paints a more complicated picture

Matching "versus" mechanism design: plan of talk

1. Introductory Example: School Choice
   - Simple win for the good properties approach
2. Main example: Course Allocation (aka combinatorial assignment)
   - Paints a more complicated picture
3. A few methodological observations

Matching “versus” mechanism design: plan of talk

1. Introductory Example: School Choice
   - Simple win for the good properties approach
2. Main example: Course Allocation (aka combinatorial assignment)
   - Paints a more complicated picture
3. A few methodological observations
4. From theory to practice: course allocation at Wharton

Introductory example: school choice

Introductory example: school choice

In a seminal paper, Abdulkadiroglu and Sonmez (2003) initiate the market design literature on the school choice problem.

Introductory example: school choice

In a seminal paper, Abdulkadiroglu and Sonmez (2003) initiate the market design literature on the school choice problem.

They propose two mechanisms that satisfy attractive properties:

Introductory example: school choice

In a seminal paper, Abdulkadiroglu and Sonmez (2003) initiate the market design literature on the school choice problem.

They propose two mechanisms that satisfy attractive properties:

1. Gale-Shapley variant, adapted for school choice
   - Stable (i.e. no justified envy)
   - Strategyproof for students

Introductory example: school choice

In a seminal paper, Abdulkadiroglu and Sonmez (2003) initiate the market design literature on the school choice problem.

They propose two mechanisms that satisfy attractive properties:

1. Gale-Shapley variant, adapted for school choice
    - Stable (i.e. no justified envy)
    - Strategyproof for students
2. Gale Top Trading Cycles variant
    - Pareto efficient for students
    - Strategyproof for students

Introductory example: school choice

Introductory example: school choice

AS then guide policy makers on how to choose between these two approaches

- "In some applications, policy makers may rank complete elimination of justified envy before full [student] efficiency, and Gale-Shapley student optimal stable mechanism can be used in those cases. . . "

Introductory example: school choice

AS then guide policy makers on how to choose between these two approaches

- "In some applications, policy makers may rank complete elimination of justified envy before full [student] efficiency, and Gale-Shapley student optimal stable mechanism can be used in those cases. . . "
- "In other applications, the top trading cycles mechanism may be more appealing. . . ."

Introductory example: school choice

AS then guide policy makers on how to choose between these two approaches

- "In some applications, policy makers may rank complete elimination of justified envy before full [student] efficiency, and Gale-Shapley student optimal stable mechanism can be used in those cases. . . "
- "In other applications, the top trading cycles mechanism may be more appealing. . . ."
- "In other cases the choice between the two mechanisms may be less clear and it depends on the policy priorities of the policy makers"

Introductory example: school choice

Not as definitive a conclusion as Vickrey, Myerson ...

But a hugely important paper, with big policy successes associated with it.

Introductory example: school choice

Not as definitive a conclusion as Vickrey, Myerson ...

But a hugely important paper, with big policy successes associated with it.

Why? Mechanisms that came before AS (2003) had serious flaws

- "Boston" mechanism (incentive problems)
- "Non mechanisms"

Introductory example: school choice

Not as definitive a conclusion as Vickrey, Myerson ...

But a hugely important paper, with big policy successes associated with it.

Why? Mechanisms that came before AS (2003) had serious flaws

- "Boston" mechanism (incentive problems)
- "Non mechanisms"

Thanks to AS we now have two mechanisms that satisfy attractive properties like Pareto efficiency, strategyproofness, stability, etc.

Introductory example: school choice

Not as definitive a conclusion as Vickrey, Myerson ...

But a hugely important paper, with big policy successes associated with it.

Why? Mechanisms that came before AS (2003) had serious flaws

- "Boston" mechanism (incentive problems)
- "Non mechanisms"

Thanks to AS we now have two mechanisms that satisfy attractive properties like Pareto efficiency, strategyproofness, stability, etc.

Moreover, subsequent empirical work on NYC suggests that the good properties approach was a reliable guide to welfare

The fact that we don't know the "optimal" school choice mechanism doesn't mean that we shouldn't discuss "good" school choice mechanisms!

Main example: course allocation

Main example: course allocation

- School choice is a win for the "mechanism with good properties" approach.

Main example: course allocation

- School choice is a win for the "mechanism with good properties" approach.
- I now want to turn to a problem where the story is a bit more complicated: assignment with multi-unit demand

Main example: course allocation

- School choice is a win for the "mechanism with good properties" approach.
- I now want to turn to a problem where the story is a bit more complicated: assignment with multi-unit demand

Specific instance: course allocation at universities

- The indivisible objects are seats in courses
- Each student requires a bundle of courses
- Exogenous restriction against monetary transfers (even at Chicago!)

Main example: course allocation

- School choice is a win for the "mechanism with good properties" approach.
- I now want to turn to a problem where the story is a bit more complicated: assignment with multi-unit demand

Specific instance: course allocation at universities

- The indivisible objects are seats in courses
- Each student requires a bundle of courses
- Exogenous restriction against monetary transfers (even at Chicago!)

Other examples: assigning interchangeable workers to tasks or shifts; leads to salespeople; takeoff and landing slots to airlines; shared scientific resources amongst scientists; players to teams

Environment

Environment

- Set of N students $\mathcal{S}$ (s_i)

Environment

- Set of N students $\mathcal{S}$ (s_i)
- Set of M courses $\mathcal{C}$ (c_j) with integral capacities $\mathbf{q} = (q_1, ..., q_M)$. No other goods in the economy.

Environment

- Set of N students $\mathcal{S}$ (s_i)
- Set of M courses $\mathcal{C}$ (c_j) with integral capacities $\mathbf{q} = (q_1, ..., q_M)$. No other goods in the economy.
- Each student s_i has a set of permissible schedules $\Psi_i \subseteq 2^{\mathcal{C}}$, and a utility function $u_i : 2^{\mathcal{C}} \to \mathbb{R}_+$
  - Impermissible schedules have utility of zero.
  - No peer effects.
  - Will sometimes make additional assumptions about preferences (e.g. responsiveness)

Environment

- Set of N students $\mathcal{S}$ (s_i)
- Set of M courses $\mathcal{C}$ (c_j) with integral capacities $\mathbf{q} = (q_1, ..., q_M)$. No other goods in the economy.
- Each student s_i has a set of permissible schedules $\Psi_i \subseteq 2^{\mathcal{C}}$, and a utility function $u_i : 2^{\mathcal{C}} \to \mathbb{R}_+$
  - Impermissible schedules have utility of zero.
  - No peer effects.
  - Will sometimes make additional assumptions about preferences (e.g. responsiveness)
- An allocation $\mathbf{x} = (x_i)_{i=1}^N$ is feasible if each $x_i \in 2^{\mathcal{C}}$ and $\sum_{i=1}^N x_{ij} \leq q_j$ for each j

Efficiency Criteria

Three notions of efficiency

Efficiency Criteria

Three notions of efficiency

1. Ex-post Pareto efficiency

   - A feasible allocation is ex-post Pareto efficient if there is no other such allocation weakly preferred by all, strictly by some

Efficiency Criteria

Three notions of efficiency

1. Ex-post Pareto efficiency
   - A feasible allocation is ex-post Pareto efficient if there is no other such allocation weakly preferred by all, strictly by some
2. Ex-ant Pareto efficiency
   - A lottery over feasible allocations is ex-ante efficient if there is no other such lottery weakly preferred by all, strictly by some
   - Note: ex-ante implies ex-post, converse false

Efficiency Criteria

Three notions of efficiency

1. Ex-post Pareto efficiency
   - A feasible allocation is ex-post Pareto efficient if there is no other such allocation weakly preferred by all, strictly by some
2. Ex-ant Pareto efficiency
   - A lottery over feasible allocations is ex-ante efficient if there is no other such lottery weakly preferred by all, strictly by some
   - Note: ex-ante implies ex-post, converse false
3. Max social welfare.
   - Allocation $\mathbf{x}$ maximizes $\sum_{i=1}^{N} u_i(x_i)$ subject to feasibility
   - Requires inter-personal comparability, and taking a stand on welfare weights / utility normalization

Efficiency Criteria

Three notions of efficiency

1. Ex-post Pareto efficiency
   - A feasible allocation is ex-post Pareto efficient if there is no other such allocation weakly preferred by all, strictly by some
2. Ex-ant Pareto efficiency
   - A lottery over feasible allocations is ex-ante efficient if there is no other such lottery weakly preferred by all, strictly by some
   - Note: ex-ante implies ex-post, converse false
3. Max social welfare.
   - Allocation $\mathbf{x}$ maximizes $\sum_{i=1}^{N} u_i(x_i)$ subject to feasibility
   - Requires inter-personal comparability, and taking a stand on welfare weights / utility normalization

In NTU assignment: Max social welfare $\subset$ Ex-ante Pareto efficient $\subset$ Ex-post Pareto efficient

Efficiency Criteria

Three notions of efficiency

1. Ex-post Pareto efficiency
   - A feasible allocation is ex-post Pareto efficient if there is no other such allocation weakly preferred by all, strictly by some
2. Ex-ant Pareto efficiency
   - A lottery over feasible allocations is ex-ante efficient if there is no other such lottery weakly preferred by all, strictly by some
   - Note: ex-ante implies ex-post, converse false
3. Max social welfare.
   - Allocation $\mathbf{x}$ maximizes $\sum_{i=1}^{N} u_i(x_i)$ subject to feasibility
   - Requires inter-personal comparability, and taking a stand on welfare weights / utility normalization

In NTU assignment: Max social welfare $\subset$ Ex-ante Pareto efficient $\subset$ Ex-post Pareto efficient

By contrast, in TU settings the three concepts tend to exactly coincide (e.g. Vickrey auction)

Impossibility Theorems

Impossibility Theorems

- Ex-ante Pareto efficiency. There is no symmetric mechanism that is ex-ante Pareto efficient and strategyproof (Zhou, 1990)
  - Note contrast to setting with monetary transfers; VCG maximizes social welfare and is strategyproof

Impossibility Theorems

- Ex-ante Pareto efficiency. There is no symmetric mechanism that is ex-ante Pareto efficient and strategyproof (Zhou, 1990)
  - Note contrast to setting with monetary transfers; VCG maximizes social welfare and is strategyproof
- Ex-post Pareto efficiency. Essentially, the only mechanisms that are ex-post Pareto efficient and strategyproof are serial/sequential dictatorships (Papai (2001), Ehlers and Klaus (2003), Hatfield (2009))

Impossibility Theorems

- Ex-ante Pareto efficiency. There is no symmetric mechanism that is ex-ante Pareto efficient and strategyproof (Zhou, 1990)
  - Note contrast to setting with monetary transfers; VCG maximizes social welfare and is strategyproof
- Ex-post Pareto efficiency. Essentially, the only mechanisms that are ex-post Pareto efficient and strategyproof are serial/sequential dictatorships (Papai (2001), Ehlers and Klaus (2003), Hatfield (2009))
- Essentially no progress on the "constrained Max SWF" problem, for either Bayesian IC or dominant strategy IC

What should we make of this?

What should we make of this?

In a sense, we are in a position that is similar to school choice after AS (2003).

What should we make of this?

In a sense, we are in a position that is similar to school choice after AS (2003).

We have a mechanism that is SP + ex-post Pareto efficient, don't know much about ex-ante efficiency, and don't know much about Max SWF s.t. constraints.

What should we make of this?

In a sense, we are in a position that is similar to school choice after AS (2003).

We have a mechanism that is SP + ex-post Pareto efficient, don't know much about ex-ante efficiency, and don't know much about Max SWF s.t. constraints.

- Papai (2001, p. 270): "[t]he implications are clear (...) if strategic manipulation is an issue, one should seriously consider using a serial dictatorship, however restrictive it may seem."

What should we make of this?

In a sense, we are in a position that is similar to school choice after AS (2003).

We have a mechanism that is SP + ex-post Pareto efficient, don't know much about ex-ante efficiency, and don't know much about Max SWF s.t. constraints.

- Papai (2001, p. 270): "[t]he implications are clear (...) if strategic manipulation is an issue, one should seriously consider using a serial dictatorship, however restrictive it may seem."
- Ehlers and Klaus (2003, p. 266): "[a] practical advantage of dictatorships is that they are simple and can be implemented easily. Furthermore, they are efficient, strategyproof (...). They can be considered to be 'fair' if the ordering of the agents is fairly determined; for instance by queuing, seniority, or randomization."

What should we make of this?

In a sense, we are in a position that is similar to school choice after AS (2003).

We have a mechanism that is SP + ex-post Pareto efficient, don't know much about ex-ante efficiency, and don't know much about Max SWF s.t. constraints.

- Papai (2001, p. 270): "[t]he implications are clear (...) if strategic manipulation is an issue, one should seriously consider using a serial dictatorship, however restrictive it may seem."
- Ehlers and Klaus (2003, p. 266): "[a] practical advantage of dictatorships is that they are simple and can be implemented easily. Furthermore, they are efficient, strategyproof (...). They can be considered to be 'fair' if the ordering of the agents is fairly determined; for instance by queuing, seniority, or randomization."
- Hatfield (2009, p. 514): "[the] results have shown that the only acceptable mechanisms for allocation problems of this sort is a sequential dictatorship, even when we restrict preferences to be responsive (...). Although unfortunate, it seems that in many of these applications, the best procedure (...) may well be a random serial dictatorship."

A Worry

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

But does the dictatorship stray too far from the underlying problem of maximizing social welfare s.t. constraints?

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

But does the dictatorship stray too far from the underlying problem of maximizing social welfare s.t. constraints? That is, does it stray too far from the problem that we would like to solve, but don't know how to solve?

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

But does the dictatorship stray too far from the underlying problem of maximizing social welfare s.t. constraints? That is, does it stray too far from the problem that we would like to solve, but don't know how to solve?

In NTU assignment there are a lot of ex-post Pareto efficient allocations, some of which seem quite different from Max SWF.

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

But does the dictatorship stray too far from the underlying problem of maximizing social welfare s.t. constraints? That is, does it stray too far from the problem that we would like to solve, but don't know how to solve?

In NTU assignment there are a lot of ex-post Pareto efficient allocations, some of which seem quite different from Max SWF. Example:

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

But does the dictatorship stray too far from the underlying problem of maximizing social welfare s.t. constraints? That is, does it stray too far from the problem that we would like to solve, but don't know how to solve?

In NTU assignment there are a lot of ex-post Pareto efficient allocations, some of which seem quite different from Max SWF. Example:

- 2 students who require 10 courses each.

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

But does the dictatorship stray too far from the underlying problem of maximizing social welfare s.t. constraints? That is, does it stray too far from the problem that we would like to solve, but don't know how to solve?

In NTU assignment there are a lot of ex-post Pareto efficient allocations, some of which seem quite different from Max SWF. Example:

- 2 students who require 10 courses each.
- 20 course seats: 10 have "good" professors, 10 have "bad" professors

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

But does the dictatorship stray too far from the underlying problem of maximizing social welfare s.t. constraints? That is, does it stray too far from the problem that we would like to solve, but don't know how to solve?

In NTU assignment there are a lot of ex-post Pareto efficient allocations, some of which seem quite different from Max SWF. Example:

- 2 students who require 10 courses each.
- 20 course seats: 10 have "good" professors, 10 have "bad" professors
- Both students agree that any "good" class is better than any "bad" class, and have responsive preferences

A Worry

Strategyproofness and ex-post Pareto efficiency are certainly attractive properties.

But does the dictatorship stray too far from the underlying problem of maximizing social welfare s.t. constraints? That is, does it stray too far from the problem that we would like to solve, but don't know how to solve?

In NTU assignment there are a lot of ex-post Pareto efficient allocations, some of which seem quite different from Max SWF. Example:

- 2 students who require 10 courses each.
- 20 course seats: 10 have "good" professors, 10 have "bad" professors
- Both students agree that any "good" class is better than any "bad" class, and have responsive preferences
- Among the many ex-post Pareto efficient allocations are those in which one student gets all 10 good courses, while the other gets all 10 bad courses.

A mechanism from practice: the "draft"

Budish and Cantillon:"The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard" (*AER*, 2012)

A mechanism from practice: the “draft”

Budish and Cantillon:"The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard" (*AER*, 2012)

- In practice we rarely observe dictatorships, in which agents take turns choosing their entire bundle of objects.

A mechanism from practice: the “draft”

Budish and Cantillon:"The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard" (*AER*, 2012)

- In practice we rarely observe dictatorships, in which agents take turns choosing their entire bundle of objects.
- But we frequently observe "drafts", in which agents take turns choosing *one object at a time*, over a series of rounds

A mechanism from practice: the “draft”

Budish and Cantillon:"The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard" (*AER*, 2012)

- In practice we rarely observe dictatorships, in which agents take turns choosing their entire bundle of objects.
- But we frequently observe "drafts", in which agents take turns choosing *one object at a time*, over a series of rounds
- Harvard Business School's course draft

A mechanism from practice: the “draft”

Budish and Cantillon:"The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard" (*AER*, 2012)

- In practice we rarely observe dictatorships, in which agents take turns choosing their entire bundle of objects.
- But we frequently observe "drafts", in which agents take turns choosing *one object at a time*, over a series of rounds
- Harvard Business School's course draft
    1. Students submit preferences, in the form of an ROL over courses (implicit assumption: preferences are responsive)

A mechanism from practice: the “draft”

Budish and Cantillon:"The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard" (*AER*, 2012)

- In practice we rarely observe dictatorships, in which agents take turns choosing their entire bundle of objects.
- But we frequently observe "drafts", in which agents take turns choosing *one object at a time*, over a series of rounds
- Harvard Business School's course draft
  1. Students submit preferences, in the form of an ROL over courses (implicit assumption: preferences are responsive)
  2. Students are randomly ordered by the computer

A mechanism from practice: the “draft”

Budish and Cantillon:"The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard" (*AER*, 2012)

- In practice we rarely observe dictatorships, in which agents take turns choosing their entire bundle of objects.
- But we frequently observe "drafts", in which agents take turns choosing *one object at a time*, over a series of rounds
- Harvard Business School's course draft
  1. Students submit preferences, in the form of an ROL over courses (implicit assumption: preferences are responsive)
  2. Students are randomly ordered by the computer
  3. Students are allocated courses one at a time, based on their reported preferences and remaining availability.
     - Rounds 1, 3, 5, ...: ascending priority order
     - Rounds 2, 4, 6, ...: descending priority order

A mechanism from practice: the “draft”

A mechanism from practice: the “draft”

It is easy to show that the draft is not strategyproof

- Incentive to overreport "popular courses", underreport "unpopular courses"
- Intuition: don’t waste early round draft picks on courses that will sell out much later

A mechanism from practice: the “draft”

It is easy to show that the draft is not strategyproof

- Incentive to overreport "popular courses", underreport "unpopular courses"
- Intuition: don't waste early round draft picks on courses that will sell out much later

It is also straightforward to show that the draft is not ex-post Pareto efficient in Nash equilibrium

- Similar results in Brams and Straffin (1979), Manea (2007), for slightly different game forms

A mechanism from practice: the "draft"

It is easy to show that the draft is not strategyproof

- Incentive to overreport "popular courses", underreport "unpopular courses"
- Intuition: don't waste early round draft picks on courses that will sell out much later

It is also straightforward to show that the draft is not ex-post Pareto efficient in Nash equilibrium

- Similar results in Brams and Straffin (1979), Manea (2007), for slightly different game forms

So, on the properties emphasized by Abdulkadiroglu and Sonmez (2003),Papai (2001), Ehlers and Klaus (2003), and Hatfield (2009):

A mechanism from practice: the “draft”

It is easy to show that the draft is not strategyproof

- Incentive to overreport "popular courses", underreport "unpopular courses"
- Intuition: don't waste early round draft picks on courses that will sell out much later

It is also straightforward to show that the draft is not ex-post Pareto efficient in Nash equilibrium

- Similar results in Brams and Straffin (1979), Manea (2007), for slightly different game forms

So, on the properties emphasized by Abdulkadiroglu and Sonmez (2003),Papai (2001), Ehlers and Klaus (2003), and Hatfield (2009):

$$\text{dictatorship} > \text{draft}$$

A different efficiency question

A different efficiency question

We ask a different question about efficiency: how well does the draft do at the problem of maximizing ex-ante social welfare?

A different efficiency question

We ask a different question about efficiency: how well does the draft do at the problem of maximizing ex-ante social welfare?

- All we know from the failure of ex-post Pareto efficiency is that the draft doesn't achieve the unconstrained maximum.

A different efficiency question

We ask a different question about efficiency: how well does the draft do at the problem of maximizing ex-ante social welfare?

- All we know from the failure of ex-post Pareto efficiency is that the draft doesn't achieve the unconstrained maximum.
- And we know that RSD doesn't achieve the unconstrained maximum either, from Hylland and Zeckhauser (1979)

A different efficiency question

We ask a different question about efficiency: how well does the draft do at the problem of maximizing ex-ante social welfare?

- All we know from the failure of ex-post Pareto efficiency is that the draft doesn't achieve the unconstrained maximum.
- And we know that RSD doesn't achieve the unconstrained maximum either, from Hylland and Zeckhauser (1979)

Data (from 2005-2006 academic year)

- Students' actual submitted ROLs (potentially strategic)
- Students' underlying truthful ROLs, from an administration survey (caveats / robustness in paper)

A different efficiency question

A different efficiency question

- Key feature of the data: because we have truthful and strategic preferences, we can look directly at how well the HBS draft does at the "Max SWF s.t. constraints" problem.

A different efficiency question

- Key feature of the data: because we have truthful and strategic preferences, we can look directly at how well the HBS draft does at the "Max SWF s.t. constraints" problem.
- We can also use the truthful preferences to simulate equilibrium play of the counterfactual of interest, RSD

A different efficiency question

- Key feature of the data: because we have truthful and strategic preferences, we can look directly at how well the HBS draft does at the "Max SWF s.t. constraints" problem.
- We can also use the truthful preferences to simulate equilibrium play of the counterfactual of interest, RSD
- On some simple measures of ex-ante welfare, the draft looks better than the dictatorship:

A different efficiency question

- Key feature of the data: because we have truthful and strategic preferences, we can look directly at how well the HBS draft does at the "Max SWF s.t. constraints" problem.
- We can also use the truthful preferences to simulate equilibrium play of the counterfactual of interest, RSD
- On some simple measures of ex-ante welfare, the draft looks better than the dictatorship:

	E(Avg Rank)	% Who Get #1 Choice	% Who Get All Top 10

A different efficiency question

- Key feature of the data: because we have truthful and strategic preferences, we can look directly at how well the HBS draft does at the "Max SWF s.t. constraints" problem.
- We can also use the truthful preferences to simulate equilibrium play of the counterfactual of interest, RSD
- On some simple measures of ex-ante welfare, the draft looks better than the dictatorship:

	E(Avg Rank)	% Who Get #1 Choice	% Who Get All Top 10
No Scarcity	5.50	100%	100%

A different efficiency question

- Key feature of the data: because we have truthful and strategic preferences, we can look directly at how well the HBS draft does at the "Max SWF s.t. constraints" problem.
- We can also use the truthful preferences to simulate equilibrium play of the counterfactual of interest, RSD
- On some simple measures of ex-ante welfare, the draft looks better than the dictatorship:

	E(Avg Rank)	% Who Get #1 Choice	% Who Get All Top 10
No Scarcity	5.50	100%	100%
HBS-Truthful	7.67	82%	1.3%

A different efficiency question

- Key feature of the data: because we have truthful and strategic preferences, we can look directly at how well the HBS draft does at the "Max SWF s.t. constraints" problem.
- We can also use the truthful preferences to simulate equilibrium play of the counterfactual of interest, RSD
- On some simple measures of ex-ante welfare, the draft looks better than the dictatorship:

	E(Avg Rank)	% Who Get #1 Choice	% Who Get All Top 10
No Scarcity	5.50	100%	100%
HBS-Truthful	7.67	82%	1.3%
HBS-Strategic	8.01	63%	1.9%

A different efficiency question

- Key feature of the data: because we have truthful and strategic preferences, we can look directly at how well the HBS draft does at the "Max SWF s.t. constraints" problem.
- We can also use the truthful preferences to simulate equilibrium play of the counterfactual of interest, RSD
- On some simple measures of ex-ante welfare, the draft looks better than the dictatorship:

	E(Avg Rank)	% Who Get #1 Choice	% Who Get All Top 10
No Scarcity	5.50	100%	100%
HBS-Truthful	7.67	82%	1.3%
HBS-Strategic	8.01	63%	1.9%
RSD-Truthful	8.74	49%	29.4%

A different efficiency question

Comparison of the societal average rank distribution under HBS-actual to RSD-truthful

A different efficiency question

Comparison of the societal average rank distribution under HBS-actual to RSD-truthful

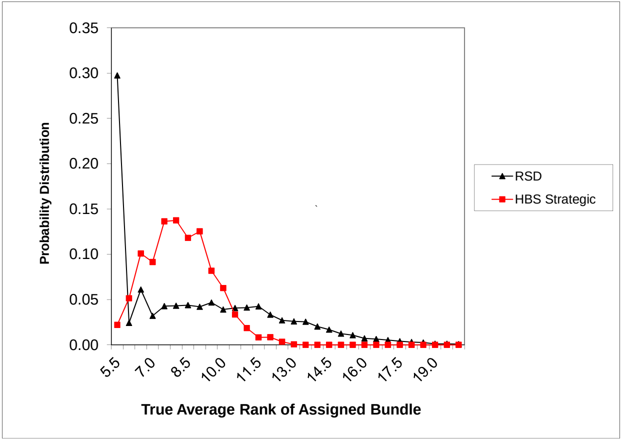

A different efficiency question

Comparison of the societal average rank distribution under HBS-actual to RSD-truthful

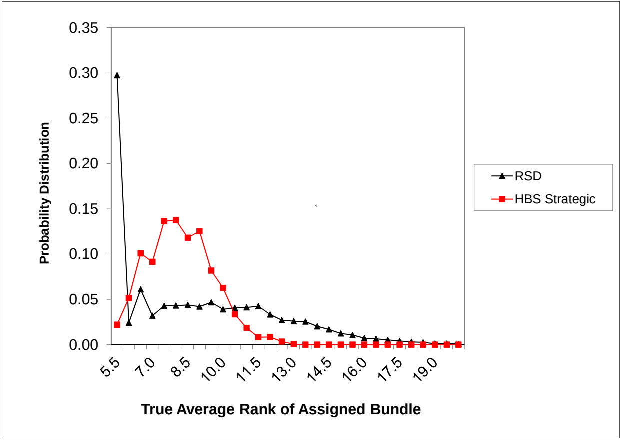

- HBS Second-Order Stochastically Dominates RSD

A different efficiency question

Comparison of the societal average rank distribution under HBS-actual to RSD-truthful

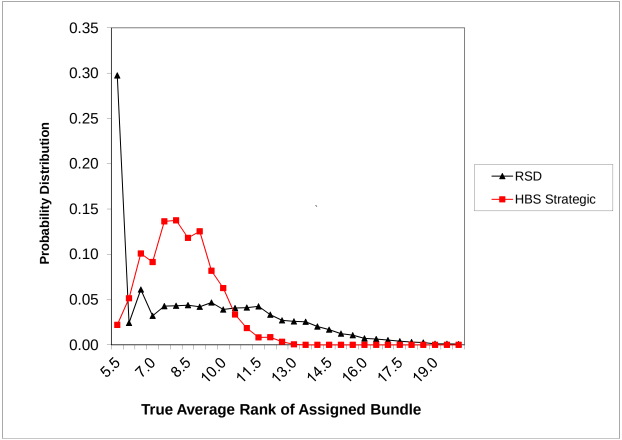

- HBS Second-Order Stochastically Dominates RSD
- Implication: social planner prefers HBS to RSD if students have average-rank preferences and are weakly risk-averse

Why is RSD so unattractive ex-ante? Example

Suppose there are 4 courses with capacity of $\frac{1}{2}N$ seats each. Students require 2 courses each. Preferences are as follows:

$\frac{N}{2}$ students are $P_1 : a, b, c, d$
$\frac{N}{2}$ students are $P_2 : b, a, d, c$

Why is RSD so unattractive ex-ante? Example

Suppose there are 4 courses with capacity of $\frac{1}{2}N$ seats each. Students require 2 courses each. Preferences are as follows:

$\frac{N}{2}$ students are $P_1 : a, b, c, d$
$\frac{N}{2}$ students are $P_2 : b, a, d, c$

- What happens under RSD?

Why is RSD so unattractive ex-ante? Example

Suppose there are 4 courses with capacity of $\frac{1}{2}N$ seats each. Students require 2 courses each. Preferences are as follows:

$\frac{N}{2}$ students are $P_1 : a, b, c, d$
$\frac{N}{2}$ students are $P_2 : b, a, d, c$

- What happens under RSD?
  - Pr $\frac{1}{2}$: get 1st and 2nd favorites ($\{a, b\}$)

Why is RSD so unattractive ex-ante? Example

Suppose there are 4 courses with capacity of $\frac{1}{2}N$ seats each. Students require 2 courses each. Preferences are as follows:

$\frac{N}{2}$ students are $P_1 : a, b, c, d$
$\frac{N}{2}$ students are $P_2 : b, a, d, c$

- What happens under RSD?
  - Pr $\frac{1}{2}$: get 1st and 2nd favorites ($\{a, b\}$)
  - Pr $\frac{1}{2}$: get 3rd and 4th favorites ($\{c, d\}$)

Why is RSD so unattractive ex-ante? Example

Suppose there are 4 courses with capacity of $\frac{1}{2}N$ seats each. Students require 2 courses each. Preferences are as follows:

$\frac{N}{2}$ students are $P_1 : a, b, c, d$
$\frac{N}{2}$ students are $P_2 : b, a, d, c$

- What happens under RSD?
  - Pr $\frac{1}{2}$: get 1st and 2nd favorites ($\{a, b\}$)
  - Pr $\frac{1}{2}$: get 3rd and 4th favorites ($\{c, d\}$)
- What happens under HBS?

Why is RSD so unattractive ex-ante? Example

Suppose there are 4 courses with capacity of $\frac{1}{2}N$ seats each. Students require 2 courses each. Preferences are as follows:

$\frac{N}{2}$ students are $P_1 : a, b, c, d$
$\frac{N}{2}$ students are $P_2 : b, a, d, c$

- What happens under RSD?
  - Pr $\frac{1}{2}$: get 1st and 2nd favorites ($\{a, b\}$)
  - Pr $\frac{1}{2}$: get 3rd and 4th favorites ($\{c, d\}$)
- What happens under HBS?
  - Always get 1st and 3rd favorites

Why is RSD so unattractive ex-ante? Example

Suppose there are 4 courses with capacity of $\frac{1}{2}N$ seats each. Students require 2 courses each. Preferences are as follows:

$\frac{N}{2}$ students are $P_1 : a, b, c, d$
$\frac{N}{2}$ students are $P_2 : b, a, d, c$

- What happens under RSD?
  - Pr $\frac{1}{2}$: get 1st and 2nd favorites ($\{a, b\}$)
  - Pr $\frac{1}{2}$: get 3rd and 4th favorites ($\{c, d\}$)
- What happens under HBS?
  - Always get 1st and 3rd favorites
  - P_1 types always get $\{a, c\}$, P_2 types get $\{b, d\}$

Why is RSD so unattractive ex-ante? Example

Suppose there are 4 courses with capacity of $\frac{1}{2}N$ seats each. Students require 2 courses each. Preferences are as follows:

$\frac{N}{2}$ students are $P_1 : a, b, c, d$
$\frac{N}{2}$ students are $P_2 : b, a, d, c$

- What happens under RSD?
  - Pr $\frac{1}{2}$: get 1st and 2nd favorites ($\{a, b\}$)
  - Pr $\frac{1}{2}$: get 3rd and 4th favorites ($\{c, d\}$)
- What happens under HBS?
  - Always get 1st and 3rd favorites
  - P_1 types always get $\{a, c\}$, P_2 types get $\{b, d\}$
  - Note: truthful play is an eqm

Why is RSD so unattractive ex-ante? Callousness

- In RSD, lucky students with good random draws make their last choices independently of whether these courses would be some unlucky students' first choices.

Why is RSD so unattractive ex-ante? Callousness

- In RSD, lucky students with good random draws make their last choices independently of whether these courses would be some unlucky students' first choices.
  - In the example, lucky P_1's take their second-choice b which is some unlucky P_2's first choice. (and vice versa)

Why is RSD so unattractive ex-ante? Callousness

- In RSD, lucky students with good random draws make their last choices independently of whether these courses would be some unlucky students' first choices.
  - In the example, lucky P_1's take their second-choice b which is some unlucky P_2's first choice. (and vice versa)
- Students "callously disregard" the preferences of those who choose after them

Why is RSD so unattractive ex-ante? Callousness

- In RSD, lucky students with good random draws make their last choices independently of whether these courses would be some unlucky students' first choices.
  - In the example, lucky P_1's take their second-choice b which is some unlucky P_2's first choice. (and vice versa)
- Students "callously disregard" the preferences of those who choose after them
- Ex-post, since there are no transfers, RSD is Pareto efficient

Why is RSD so unattractive ex-ante? Callousness

- In RSD, lucky students with good random draws make their last choices independently of whether these courses would be some unlucky students' first choices.
  - In the example, lucky P_1's take their second-choice b which is some unlucky P_2's first choice. (and vice versa)
- Students "callously disregard" the preferences of those who choose after them
- Ex-post, since there are no transfers, RSD is Pareto efficient
- But ex-ante, this behavior is bad for welfare:
  - benefit to lucky is small
  - harm to unlucky is large

Why is RSD so unattractive ex-ante? Callousness

- In RSD, lucky students with good random draws make their last choices independently of whether these courses would be some unlucky students' first choices.
  - In the example, lucky P_1's take their second-choice b which is some unlucky P_2's first choice. (and vice versa)
- Students "callously disregard" the preferences of those who choose after them
- Ex-post, since there are no transfers, RSD is Pareto efficient
- But ex-ante, this behavior is bad for welfare:
  - benefit to lucky is small
  - harm to unlucky is large
- Important note: unattractiveness of RSD does not depend on risk preferences. Even risk-neutral agents regard a "win a little, lose a lot" lottery as unappealing.

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

- A sensible prior is that switching from the all-at-once / ex-post efficient RSD to the one-at-a-time / ex-post inefficient HBS would be good for fairness but bad for welfare

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

- A sensible prior is that switching from the all-at-once / ex-post efficient RSD to the one-at-a-time / ex-post inefficient HBS would be good for fairness but bad for welfare
- But in NTU settings there are many Pareto efficient allocations; and the lottery over efficient allocations induced by RSD is very unattractive when assessed ex-ante.

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

- A sensible prior is that switching from the all-at-once / ex-post efficient RSD to the one-at-a-time / ex-post inefficient HBS would be good for fairness but bad for welfare
- But in NTU settings there are many Pareto efficient allocations; and the lottery over efficient allocations induced by RSD is very unattractive when assessed ex-ante.
- So much so that the HBS lottery over *in*efficient allocations looks more attractive ex-ante than RSD.

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

- A sensible prior is that switching from the all-at-once / ex-post efficient RSD to the one-at-a-time / ex-post inefficient HBS would be good for fairness but bad for welfare
- But in NTU settings there are many Pareto efficient allocations; and the lottery over efficient allocations induced by RSD is very unattractive when assessed ex-ante.
- So much so that the HBS lottery over *in*efficient allocations looks more attractive ex-ante than RSD.
  - No efficiency-fairness tradeoff

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

- A sensible prior is that switching from the all-at-once / ex-post efficient RSD to the one-at-a-time / ex-post inefficient HBS would be good for fairness but bad for welfare
- But in NTU settings there are many Pareto efficient allocations; and the lottery over efficient allocations induced by RSD is very unattractive when assessed ex-ante.
- So much so that the HBS lottery over *in*efficient allocations looks more attractive ex-ante than RSD.
  - No efficiency-fairness tradeoff
  - Ex-post efficiency need not even proxy for ex-ante efficiency

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

- A sensible prior is that switching from the all-at-once / ex-post efficient RSD to the one-at-a-time / ex-post inefficient HBS would be good for fairness but bad for welfare
- But in NTU settings there are many Pareto efficient allocations; and the lottery over efficient allocations induced by RSD is very unattractive when assessed ex-ante.
- So much so that the HBS lottery over *in*efficient allocations looks more attractive ex-ante than RSD.
  - No efficiency-fairness tradeoff
  - Ex-post efficiency need not even proxy for ex-ante efficiency
- Punchline: if the "real" problem is Max SWF s.t. constraints, the HBS draft appears to be a better solution than the RSD

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

- A sensible prior is that switching from the all-at-once / ex-post efficient RSD to the one-at-a-time / ex-post inefficient HBS would be good for fairness but bad for welfare
- But in NTU settings there are many Pareto efficient allocations; and the lottery over efficient allocations induced by RSD is very unattractive when assessed ex-ante.
- So much so that the HBS lottery over *in*efficient allocations looks more attractive ex-ante than RSD.
  - No efficiency-fairness tradeoff
  - Ex-post efficiency need not even proxy for ex-ante efficiency
- Punchline: if the "real" problem is Max SWF s.t. constraints, the HBS draft appears to be a better solution than the RSD
- "Mistake" in the prior literature was to conclude that because we can't get exact ex-ante efficiency, we should settle for exact ex-post efficiency...

What do we learn from the HBS draft?

Lesson 1: Fairness and Welfare

- A sensible prior is that switching from the all-at-once / ex-post efficient RSD to the one-at-a-time / ex-post inefficient HBS would be good for fairness but bad for welfare
- But in NTU settings there are many Pareto efficient allocations; and the lottery over efficient allocations induced by RSD is very unattractive when assessed ex-ante.
- So much so that the HBS lottery over *in*efficient allocations looks more attractive ex-ante than RSD.
  - No efficiency-fairness tradeoff
  - Ex-post efficiency need not even proxy for ex-ante efficiency
- Punchline: if the "real" problem is Max SWF s.t. constraints, the HBS draft appears to be a better solution than the RSD
- "Mistake" in the prior literature was to conclude that because we can't get exact ex-ante efficiency, we should settle for exact ex-post efficiency...better to admit that we want ex-ante efficiency but don't know how to maximize it yet!

What do we learn from the HBS draft?

Lesson 2: Strategy-proofness in practical market design

What do we learn from the HBS draft?

Lesson 2: Strategy-proofness in practical market design

- Our field data allow us to directly document that students at HBS – real-life participants in a one-shot high-stakes setting – figure out how to manipulate the non-strategyproof HBS mechanism

What do we learn from the HBS draft?

Lesson 2: Strategy-proofness in practical market design

- Our field data allow us to directly document that students at HBS – real-life participants in a one-shot high-stakes setting – figure out how to manipulate the non-strategyproof HBS mechanism
- Further, we show that this manipulability harms welfare, and that the magnitudes are large

What do we learn from the HBS draft?

Lesson 2: Strategy-proofness in practical market design

- Our field data allow us to directly document that students at HBS – real-life participants in a one-shot high-stakes setting – figure out how to manipulate the non-strategyproof HBS mechanism
- Further, we show that this manipulability harms welfare, and that the magnitudes are large
- These findings are strongly consistent with the view that SP is an important desideratum in practical market design

What do we learn from the HBS draft?

Lesson 2: Strategy-proofness in practical market design

- Our field data allow us to directly document that students at HBS – real-life participants in a one-shot high-stakes setting – figure out how to manipulate the non-strategyproof HBS mechanism
- Further, we show that this manipulability harms welfare, and that the magnitudes are large
- These findings are strongly consistent with the view that SP is an important desideratum in practical market design
- However, constraints often have costs ...

What do we learn from the HBS draft?

Lesson 2: Strategy-proofness in practical market design

- Our field data allow us to directly document that students at HBS – real-life participants in a one-shot high-stakes setting – figure out how to manipulate the non-strategyproof HBS mechanism
- Further, we show that this manipulability harms welfare, and that the magnitudes are large
- These findings are strongly consistent with the view that SP is an important desideratum in practical market design
- However, constraints often have costs ...
- And we also find that the welfare costs of using a strategyproof dictatorship appear to be much larger than the welfare costs of manipulability

What do we learn from the HBS draft?

Lesson 2: Strategy-proofness in practical market design

- Our field data allow us to directly document that students at HBS – real-life participants in a one-shot high-stakes setting – figure out how to manipulate the non-strategyproof HBS mechanism
- Further, we show that this manipulability harms welfare, and that the magnitudes are large
- These findings are strongly consistent with the view that SP is an important desideratum in practical market design
- However, constraints often have costs ...
- And we also find that the welfare costs of using a strategyproof dictatorship appear to be much larger than the welfare costs of manipulability
- Overall, suggests a nuanced view of the role of strategyproofness in design, and the need for second-best alternatives to exact SP (eg "strategy-proofness in the large", Azevedo and Budish, 2013)

What do we learn from the HBS draft?

Lesson 3: "Where to look" for multi-unit assignment mechanisms that are better still

What do we learn from the HBS draft?

Lesson 3: "Where to look" for multi-unit assignment mechanisms that are better still

- Seek an incentives middle ground between strict strategyproofness (RSD) and simple-to-manipulate (HBS).

What do we learn from the HBS draft?

Lesson 3: "Where to look" for multi-unit assignment mechanisms that are better still

- Seek an incentives middle ground between strict strategyproofness (RSD) and simple-to-manipulate (HBS).
- Seek a rmechanism that yields a relatively equal distribution of outcomes, like the draft and unlike the dictatorship

Approximate CEEI

Budish: "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes" (*JPE*, 2011)

Approximate CEEI

Budish: "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes" (*JPE*, 2011)

Approach: seek a mechanism that is attractive with respect to

- Ex-post Pareto efficiency
- Ex-post fairness
- Incentive compatibility

Approximate CEEI

Budish: “The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes” (*JPE*, 2011)

Approach: seek a mechanism that is attractive with respect to

- Ex-post Pareto efficiency
- Ex-post fairness
- Incentive compatibility

That is, the "good properties" approach but with ex-post fairness an explicit concern alongside ex-post efficiency and strategyproofness.

Approximate CEEI

Budish: "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes" (*JPE*, 2011)

Approach: seek a mechanism that is attractive with respect to

- Ex-post Pareto efficiency
- Ex-post fairness
- Incentive compatibility

That is, the "good properties" approach but with ex-post fairness an explicit concern alongside ex-post efficiency and strategyproofness. The dictatorship theorems imply that this will involve compromises.

Approximate CEEI

Budish: "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes" (*JPE*, 2011)

Approach: seek a mechanism that is attractive with respect to

- Ex-post Pareto efficiency
- Ex-post fairness
- Incentive compatibility

That is, the "good properties" approach but with ex-post fairness an explicit concern alongside ex-post efficiency and strategyproofness. The dictatorship theorems imply that this will involve compromises.

No restrictions on preferences: students allowed to have arbitrary preferences over schedules. Allows for scheduling constraints, complementarities, etc.

Approximate CEEI

- Starting point is an old idea from general equilibrium theory: competitive equilibrium from equal incomes (CEEI).

Approximate CEEI

- Starting point is an old idea from general equilibrium theory: competitive equilibrium from equal incomes (CEEI).
- Well known to be a both efficient and fair solution to the problem of allocating divisible goods

Approximate CEEI

- Starting point is an old idea from general equilibrium theory: competitive equilibrium from equal incomes (CEEI).
- Well known to be a both efficient and fair solution to the problem of allocating divisible goods
- What would CEEI mean in our environment?

Approximate CEEI

- Starting point is an old idea from general equilibrium theory: competitive equilibrium from equal incomes (CEEI).
- Well known to be a both efficient and fair solution to the problem of allocating divisible goods
- What would CEEI mean in our environment?

1. Agents report preferences over bundles

Approximate CEEI

- Starting point is an old idea from general equilibrium theory: competitive equilibrium from equal incomes (CEEI).
- Well known to be a both efficient and fair solution to the problem of allocating divisible goods
- What would CEEI mean in our environment?

1. Agents report preferences over bundles
2. Agents are given equal budgets b^* of an artificial currency

Approximate CEEI

- Starting point is an old idea from general equilibrium theory: competitive equilibrium from equal incomes (CEEI).
- Well known to be a both efficient and fair solution to the problem of allocating divisible goods
- What would CEEI mean in our environment?

1. Agents report preferences over bundles
2. Agents are given equal budgets b^* of an artificial currency
3. We find an item price vector $\mathbf{p}^*$ such that, when each agent is allocated his favorite affordable bundle, the market clears

Approximate CEEI

- Starting point is an old idea from general equilibrium theory: competitive equilibrium from equal incomes (CEEI).
- Well known to be a both efficient and fair solution to the problem of allocating divisible goods
- What would CEEI mean in our environment?

1. Agents report preferences over bundles
2. Agents are given equal budgets b^* of an artificial currency
3. We find an item price vector $\mathbf{p}^*$ such that, when each agent is allocated his favorite affordable bundle, the market clears
4. We allocate each agent their demand at $\mathbf{p}^*$

Approximate CEEI

- Starting point is an old idea from general equilibrium theory: competitive equilibrium from equal incomes (CEEI).
- Well known to be a both efficient and fair solution to the problem of allocating divisible goods
- What would CEEI mean in our environment?

1. Agents report preferences over bundles
2. Agents are given equal budgets b^* of an artificial currency
3. We find an item price vector $\mathbf{p}^*$ such that, when each agent is allocated his favorite affordable bundle, the market clears
4. We allocate each agent their demand at $\mathbf{p}^*$

It is easy to see that existence is problematic.

Approximate CEEI

- I prove existence of an approximation to CEEI in which

Approximate CEEI

- I prove existence of an approximation to CEEI in which
  - Agents are given approximately equal as opposed to exactly equal budgets of an artificial currency (e.g. budgets distributed on $[1000, 1000 + \varepsilon]$)

Approximate CEEI

- I prove existence of an approximation to CEEI in which
  - Agents are given approximately equal as opposed to exactly equal budgets of an artificial currency (e.g. budgets distributed on $[1000, 1000 + \varepsilon]$)
  - The market clears approximately instead of exactly

Approximate CEEI

- I prove existence of an approximation to CEEI in which
  - Agents are given approximately equal as opposed to exactly equal budgets of an artificial currency (e.g. budgets distributed on $[1000, 1000 + \varepsilon]$)
  - The market clears approximately instead of exactly
    - Worst-case market-clearing error is "small", as measured in Euclidean distance of excess demand vector (cf. Starr, 1969)

Approximate CEEI

- I prove existence of an approximation to CEEI in which
  - Agents are given approximately equal as opposed to exactly equal budgets of an artificial currency (e.g. budgets distributed on $[1000, 1000 + \varepsilon]$)
  - The market clears approximately instead of exactly
    - Worst-case market-clearing error is "small", as measured in Euclidean distance of excess demand vector (cf. Starr, 1969)
    - Average-case performance on real data smaller still (+/- one seat in six courses, out of 4500 seats allocated)

Approximate CEEI

- I prove existence of an approximation to CEEI in which
  - Agents are given approximately equal as opposed to exactly equal budgets of an artificial currency (e.g. budgets distributed on $[1000, 1000 + \varepsilon]$)
  - The market clears approximately instead of exactly
    - Worst-case market-clearing error is "small", as measured in Euclidean distance of excess demand vector (cf. Starr, 1969)
    - Average-case performance on real data smaller still (+/- one seat in six courses, out of 4500 seats allocated)
- Equal budgets: market-clearing error could be arbitrarily large

Approximate CEEI

- I prove existence of an approximation to CEEI in which
  - Agents are given approximately equal as opposed to exactly equal budgets of an artificial currency (e.g. budgets distributed on $[1000, 1000 + \varepsilon]$)
  - The market clears approximately instead of exactly
    - Worst-case market-clearing error is "small", as measured in Euclidean distance of excess demand vector (cf. Starr, 1969)
    - Average-case performance on real data smaller still (+/- one seat in six courses, out of 4500 seats allocated)
- Equal budgets: market-clearing error could be arbitrarily large
- Other extreme: dictatorships can be interpreted as exact CE, but from arbitrarily unequal budgets

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.
- Dictatorship?

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.
- Dictatorship?
  - Fairness problems: whoever's first gets both Diamonds.

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.
- Dictatorship?
  - Fairness problems: whoever's first gets both Diamonds.
- CEEI?

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.
- Dictatorship?
  - Fairness problems: whoever's first gets both Diamonds.
- CEEI?
  - Existence problems: at any price vector, for any object, either both agents demand it or neither does.

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.
- Dictatorship?
  - Fairness problems: whoever's first gets both Diamonds.
- CEEI?
  - Existence problems: at any price vector, for any object, either both agents demand it or neither does.
- Approximate CEEI?

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.
- Dictatorship?
  - Fairness problems: whoever's first gets both Diamonds.
- CEEI?
  - Existence problems: at any price vector, for any object, either both agents demand it or neither does.
- Approximate CEEI?
  - Randomly assign budgets of 1 and $1 + \beta$, for $\beta \gtrapprox 0$

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.
- Dictatorship?
  - Fairness problems: whoever's first gets both Diamonds.
- CEEI?
  - Existence problems: at any price vector, for any object, either both agents demand it or neither does.
- Approximate CEEI?
  - Randomly assign budgets of 1 and $1 + \beta$, for $\beta \gtrapprox 0$
  - Set the price of the Big Diamond strictly greater than 1

A Simple Example: Two Diamonds, Two Rocks

- Two agents. Four objects: two valuable Diamonds (Big, Small) and two ordinary Rocks (Pretty, Ugly). At most two objects per agent.
- Dictatorship?
  - Fairness problems: whoever's first gets both Diamonds.
- CEEI?
  - Existence problems: at any price vector, for any object, either both agents demand it or neither does.
- Approximate CEEI?
  - Randomly assign budgets of 1 and $1 + \beta$, for $\beta \gtrapprox 0$
  - Set the price of the Big Diamond strictly greater than 1
    - Set other prices such that the poorer agent can afford {Small Diamond, Pretty Rock}, wealthier agent gets {Big Diamond, Ugly Rock}

Properties of the Approximate CEEI Mechanism

Efficiency

Ex-post efficient, but for small error

Fairness

N+1 Maximin Share Guarantee

Envy Bounded by a Single Good

Incentives

Strategyproof in the Large

Approximate CEEI and Matching "versus" Mechanism Design

Two possible interpretations of the role of ex-post fairness in A-CEEI

Approximate CEEI and Matching “versus” Mechanism Design

Two possible interpretations of the role of ex-post fairness in A-CEEI

1. Ex-post fairness as an explicit design objective, alongside efficiency and incentive compatibility

Approximate CEEI and Matching "versus" Mechanism Design

Two possible interpretations of the role of ex-post fairness in A-CEEI

1. Ex-post fairness as an explicit design objective, alongside efficiency and incentive compatibility
2. Ex-post fairness as a means to an end: ex-ante welfare.

Approximate CEEI and Matching “versus” Mechanism Design

Two possible interpretations of the role of ex-post fairness in A-CEEI

1. Ex-post fairness as an explicit design objective, alongside efficiency and incentive compatibility
2. Ex-post fairness as a means to an end: ex-ante welfare.

Approximate CEEI is attractive relative to alternatives under either intrepretation. (As I stated at the outset, these approaches aren't always so different!)

Table 2: Comparison of Alternative Mechanisms

Mechanism	Efficiency (Truthful Play)	Outcome Fairness (Truthful Play)	Procedural Fairness	Incentives	Preference Language
Approximate CEEI Mechanism (A-CEEI)	Pareto Efficient w/r/t Allocated Goods Allocation error is small for practice and goes to zero in the limit	N+1 – Maximin Share Guaranteed Envy Bounded by a Single Good	Symmetric	Strategyproof in the Large	Ordinal over Schedules
A-CEEI v2: Competitive Equilibrium from Equal-as-Possible Incomes (Sec 6.1)	Pareto Efficient	Worst Case: coincides with dictatorship	Symmetric	Strategyproof in the Large	Ordinal over Schedules
A-CEEI v3: A-CEEI with a Pareto-Improving Secondary Market (Sec 6.1)	Pareto Efficient	A bit weaker than N+1 – Maximin Share Guarantee, because prices in the initial allocation may be outside of P(δ,b'). Initial allocation is Envy Bounded by a Single Good. The Pareto-improvement stage may exacerbate envy.	Symmetric	Manipulable in the Large	Ordinal over Schedules
Random Serial Dictatorship (Sec 8.1)	Pareto Efficient	Worst Case: Get k worst Objects	Symmetric	Strategyproof	Ordinal over Schedules
Multi-unit generalization of Hylland Zeckhauser Mechanism (Sec 8.2)	If vNM preferences are described by assignment messages, ex-ante Pareto efficient	If preferences are additive separable, envy bounded by the value of two goods Worst Case: Get Zero Objects	Symmetric	If vNM preferences are described by assignment messages, Strategyproof in the Large	Assignment messages
Bidding Points Mechanism (Sec 8.3)	If preferences are additive-separable, Pareto Efficient but for quota issues described in Unver and Sonmez (forth.)	Worst Case: Get Zero Objects	Symmetric	Manipulable in the Large	Cardinal over Items

Mechanism	Efficiency (Truthful Play)	Outcome Fairness (Truthful Play)	Procedural Fairness	Incentives	Preference Language
Sonmez-Unver (forth.) Enhancement to Bidding Points Mechanism	If preferences are additive-separable, Pareto Efficient	Worst Case: Get Zero Objects	Symmetric	Bidding Phase: Manipulable in the Large Allocation Phase: Strategyproof in the Large	Bidding Phase: Cardinal over Items Allocation Phase: Ordinal over Items
HBS Draft Mechanism (Sec 9.2)	If preferences are responsive, Pareto Efficient with respect to the reported information (i.e., Pareto Possible)	If preferences are responsive and k=2, Maximin Share Guaranteed If preferences are responsive, Envy Bounded by a Single Good	Symmetric	Manipulable in the Large	Ordinal over Items
Bezakova and Dani (2005) Maximin Utility Algorithm	If preferences are additive-separable, ideal fractional allocation is Pareto efficient. Realized integer allocation is close to the fractional ideal.	Worst Case: Get approximately zero objects (if a hedonist and all other agents are depressives)	Symmetric	Manipulable in the Large	Cardinal over items
Brams and Taylor (1996) Adjusted Winner	If preferences are additive-separable, Pareto Efficient	Worst Case: Get Zero Objects	Symmetric	Manipulable in the Large	Cardinal over Items
Herreiner and Puppe (2002) Descending Demand Procedure	Pareto Efficient	Does not satisfy Maximin Share Guarantee or Envy Bounded by a Single Object	Symmetric	Manipulable in the Large	Ordinal over Schedules
Lipton et al (2004) Fair Allocation Mechanism	Algorithm ignores efficiency	If preferences are additive separable, Envy Bounded by a Single Good	Symmetric	Manipulable in the Large	Cardinal over items
UChicago Primal-Dual Linear Programming Mechanism (Graves et al 1993)	Pareto Efficient when preference-reporting limits don't bind	Worst Case: Get Zero Objects	Symmetric	Manipulable in the Large	Cardinal over a Limited Number of Schedules

Figure 3: Ex-Ante Efficiency Comparison

Approximate CEEI Mechanism vs. HBS Draft Mechanism

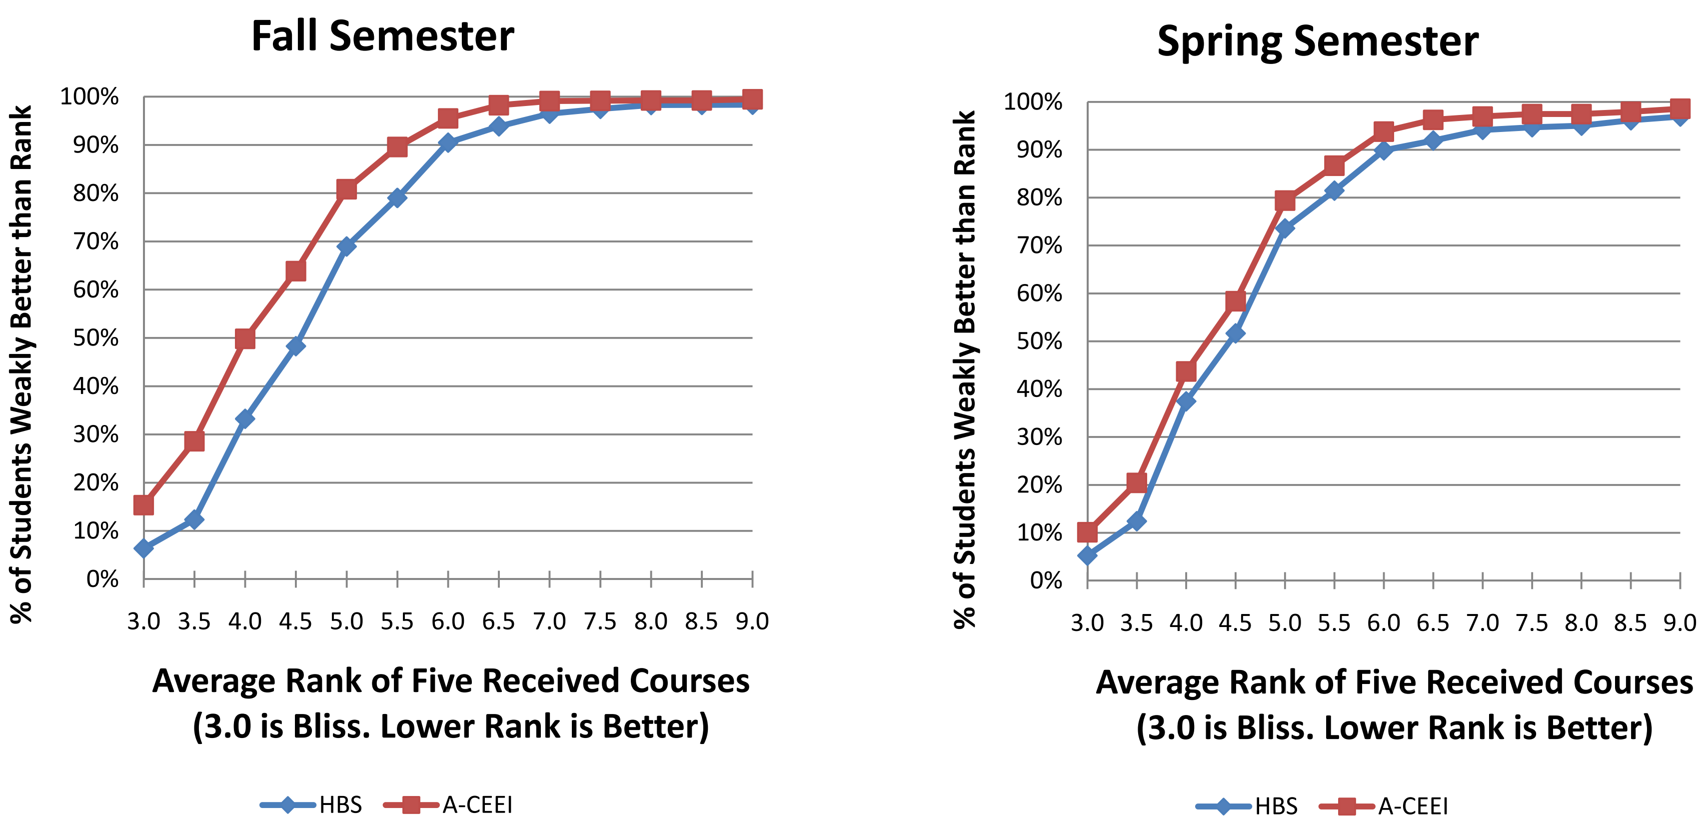

Description: The Othman, Budish and Sandholm (2010) Approximate CEEI algorithm is run 100 times for each semester of the Harvard Business School course allocation data (456 students, ~50 courses, 5 courses per student). Each run uses randomly generated budgets. For each random budget ordering I also run the HBS Draft Mechanism, using the random budget order as the draft order. The HBS Draft Mechanism is run using students' actual strategic reports under that mechanism. The Approximate CEEI algorithm is run using students' truthful preferences. This table reports the cumulative distribution of outcomes, as measured by average rank, over the 456*100 = 45,600 student-trial pairs. Average rank is calculated based on the student's true preferences. For instance, a student who receives her 1,2,3,4 and 5^{th} favorite courses has an average rank of (1+2+3+4+5)/5 = 3.

Matching “versus” mechanism design: some reflections

Observation 1: new tools are needed!

Matching “versus” mechanism design: some reflections

Observation 1: new tools are needed!

What makes "max SWF s.t. constraints" difficult here?

Matching “versus” mechanism design: some reflections

Observation 1: new tools are needed!

What makes "max SWF s.t. constraints" difficult here?

- Basic difficulty: no money. Somehow need to avoid incentivizing agents to report that their utility from their favorite bundle is $+\infty$

Matching “versus” mechanism design: some reflections

Observation 1: new tools are needed!

What makes "max SWF s.t. constraints" difficult here?

- Basic difficulty: no money. Somehow need to avoid incentivizing agents to report that their utility from their favorite bundle is $+\infty$
  - AS (2003), Approximate CEEI: overcome this by asking only for ordinal preference information (over individual objects and bundles, resp.)

Matching "versus" mechanism design: some reflections

Observation 1: new tools are needed!

What makes "max SWF s.t. constraints" difficult here?

- Basic difficulty: no money. Somehow need to avoid incentivizing agents to report that their utility from their favorite bundle is $+\infty$
  - AS (2003), Approximate CEEI: overcome this by asking only for ordinal preference information (over individual objects and bundles, resp.)
  - Hylland-Zeckhauser (1979), as well as its multi-unit generalization (Budish et al, 2013): overcome this by asking only for marginal rates of substitution across individual objects

Matching “versus” mechanism design: some reflections

Observation 1: new tools are needed!

What makes "max SWF s.t. constraints" difficult here?

- Basic difficulty: no money. Somehow need to avoid incentivizing agents to report that their utility from their favorite bundle is $+\infty$
  - AS (2003), Approximate CEEI: overcome this by asking only for ordinal preference information (over individual objects and bundles, resp.)
  - Hylland-Zeckhauser (1979), as well as its multi-unit generalization (Budish et al, 2013): overcome this by asking only for marginal rates of substitution across individual objects
- Perhaps there is a better way?

Matching "versus" mechanism design: some reflections

Observation 1: new tools are needed!

What makes "max SWF s.t. constraints" difficult here?

- Basic difficulty: no money. Somehow need to avoid incentivizing agents to report that their utility from their favorite bundle is $+\infty$
  - AS (2003), Approximate CEEI: overcome this by asking only for ordinal preference information (over individual objects and bundles, resp.)
  - Hylland-Zeckhauser (1979), as well as its multi-unit generalization (Budish et al, 2013): overcome this by asking only for marginal rates of substitution across individual objects
- Perhaps there is a better way?
  - Recent work by Nguyen and Vohra seems quite promising in this regard
  - Intead of IC constraints, use envy-freeness constraints, and rely on large-market relationship between EF and SP-L (cf. Azevedo and Budish, 2013)

Matching "versus" mechanism design: some reflections

Observation 2: know thy objective

Matching "versus" mechanism design: some reflections

Observation 2: know thy objective

- Positive design should always be clear on the true objectives and the true constraints

Matching “versus” mechanism design: some reflections

Observation 2: know thy objective

- Positive design should always be clear on the true objectives and the true constraints
- "Mistake" in the axiomatic literature on multi-unit assigment was to conclude that, because we can’t get exact ex-ante efficiency, we should settle for exact ex-post efficiency

Matching “versus” mechanism design: some reflections

Observation 2: know thy objective

- Positive design should always be clear on the true objectives and the true constraints
- "Mistake" in the axiomatic literature on multi-unit assigment was to conclude that, because we can’t get exact ex-ante efficiency, we should settle for exact ex-post efficiency
  - In school choice, ex-post efficiency was a reliable guide to welfare and good market design, but not so for course allocation

Matching “versus” mechanism design: some reflections

Observation 2: know thy objective

- Positive design should always be clear on the true objectives and the true constraints
- "Mistake" in the axiomatic literature on multi-unit assigment was to conclude that, because we can't get exact ex-ante efficiency, we should settle for exact ex-post efficiency
  - In school choice, ex-post efficiency was a reliable guide to welfare and good market design, but not so for course allocation
- Sometimes we don't know how to maximize the true objective subject to the true constraints because of limitations of the theory.

Matching "versus" mechanism design: some reflections

Observation 2: know thy objective

- Positive design should always be clear on the true objectives and the true constraints
- "Mistake" in the axiomatic literature on multi-unit assigment was to conclude that, because we can't get exact ex-ante efficiency, we should settle for exact ex-post efficiency
  - In school choice, ex-post efficiency was a reliable guide to welfare and good market design, but not so for course allocation
- Sometimes we don't know how to maximize the true objective subject to the true constraints because of limitations of the theory.
- Such situations call for attacking the true objective using both theory and either data or computational simulations (Roth, 2002)

Matching "versus" mechanism design: some reflections

Observation 2: know thy objective

- Positive design should always be clear on the true objectives and the true constraints
- "Mistake" in the axiomatic literature on multi-unit assigment was to conclude that, because we can't get exact ex-ante efficiency, we should settle for exact ex-post efficiency
  - In school choice, ex-post efficiency was a reliable guide to welfare and good market design, but not so for course allocation
- Sometimes we don't know how to maximize the true objective subject to the true constraints because of limitations of the theory.
- Such situations call for attacking the true objective using both theory and either data or computational simulations (Roth, 2002)
  - Approximate CEEI: use theory to obtain reasonable worst-case performance guarantees re efficiency and envy-freeness, use data to study average case

Matching “versus” mechanism design: some reflections

Observation 2: know thy objective

- Positive design should always be clear on the true objectives and the true constraints
- "Mistake" in the axiomatic literature on multi-unit assigment was to conclude that, because we can't get exact ex-ante efficiency, we should settle for exact ex-post efficiency
  - In school choice, ex-post efficiency was a reliable guide to welfare and good market design, but not so for course allocation
- Sometimes we don't know how to maximize the true objective subject to the true constraints because of limitations of the theory.
- Such situations call for attacking the true objective using both theory and either data or computational simulations (Roth, 2002)
  - Approximate CEEI: use theory to obtain reasonable worst-case performance guarantees re efficiency and envy-freeness, use data to study average case
  - We still don't know the “optimal” solution ...

Matching “versus” mechanism design: some reflections

Observation 3: know thy constraints

Matching “versus” mechanism design: some reflections

Observation 3: know thy constraints

- Sometimes it is tolerable to satisfy constraints approximately instead of exactly. Such approximations represent a challenge for both methodologies

Matching “versus” mechanism design: some reflections

Observation 3: know thy constraints

- Sometimes it is tolerable to satisfy constraints approximately instead of exactly. Such approximations represent a challenge for both methodologies
- Mechanism design approach

Matching “versus” mechanism design: some reflections

Observation 3: know thy constraints

- Sometimes it is tolerable to satisfy constraints approximately instead of exactly. Such approximations represent a challenge for both methodologies
- Mechanism design approach
  - "Max objective s.t. constraints" treats constraints as lexicographically more important than the objective

Matching “versus” mechanism design: some reflections

Observation 3: know thy constraints

- Sometimes it is tolerable to satisfy constraints approximately instead of exactly. Such approximations represent a challenge for both methodologies
- Mechanism design approach
  - "Max objective s.t. constraints" treats constraints as lexicographically more important than the objective
- Good properties approach

Matching “versus” mechanism design: some reflections

Observation 3: know thy constraints

- Sometimes it is tolerable to satisfy constraints approximately instead of exactly. Such approximations represent a challenge for both methodologies
- Mechanism design approach
  - "Max objective s.t. constraints" treats constraints as lexicographically more important than the objective
- Good properties approach
  - Stated axioms / properties imposed as lexicographically more important than other properties

Matching "versus" mechanism design: some reflections

Observation 3: know thy constraints

- Sometimes it is tolerable to satisfy constraints approximately instead of exactly. Such approximations represent a challenge for both methodologies
- Mechanism design approach
  - "Max objective s.t. constraints" treats constraints as lexicographically more important than the objective
- Good properties approach
  - Stated axioms / properties imposed as lexicographically more important than other properties
  - E.g. tendency to impose strategy-proofness inflexibly in parts of matching, social choice, algorithmic game theory

Matching “versus” mechanism design: some reflections

Observation 3: know thy constraints

- Sometimes it is tolerable to satisfy constraints approximately instead of exactly. Such approximations represent a challenge for both methodologies
- Mechanism design approach
  - "Max objective s.t. constraints" treats constraints as lexicographically more important than the objective
- Good properties approach
  - Stated axioms / properties imposed as lexicographically more important than other properties
  - E.g. tendency to impose strategy-proofness inflexibly in parts of matching, social choice, algorithmic game theory
  - E.g. in the dictatorship papers, getting exact ex-post Pareto efficiency was treated as more important than having even a modicum of ex-post fairness.

Matching "versus" mechanism design: some reflections

Observation 3: know thy constraints

- Sometimes it is tolerable to satisfy constraints approximately instead of exactly. Such approximations represent a challenge for both methodologies
- Mechanism design approach
  - "Max objective s.t. constraints" treats constraints as lexicographically more important than the objective
- Good properties approach
  - Stated axioms / properties imposed as lexicographically more important than other properties
  - E.g. tendency to impose strategy-proofness inflexibly in parts of matching, social choice, algorithmic game theory
  - E.g. in the dictatorship papers, getting exact ex-post Pareto efficiency was treated as more important than having even a modicum of ex-post fairness.
- We know from Micro 101 that we don't expect most preferences in the world to be lexicographic ... Perhaps we need new tools to make our preferences over mechanism designs a bit less lexicographic as well.

Matching “versus” mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

Matching “versus” mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

As theorists we often ask for "perfect" solutions to problems. Temptation to find *the* solution.

Matching “versus” mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

As theorists we often ask for "perfect" solutions to problems. Temptation to find *the* solution.

- Mechanism design: find the optimum

Matching "versus" mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

As theorists we often ask for "perfect" solutions to problems. Temptation to find *the* solution.

- Mechanism design: find the optimum
- Axiomatic approach: get a unique characterization

Matching “versus” mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

As theorists we often ask for "perfect" solutions to problems. Temptation to find *the* solution.

- Mechanism design: find the optimum
- Axiomatic approach: get a unique characterization

Danger: we end up only paying attention to the problems we can solve beautifully. "Keys under the lamppost" joke.

Matching "versus" mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

As theorists we often ask for "perfect" solutions to problems. Temptation to find *the* solution.

- Mechanism design: find the optimum
- Axiomatic approach: get a unique characterization

Danger: we end up only paying attention to the problems we can solve beautifully. "Keys under the lamppost" joke.

I think there is value in designing mechanisms that are "better" even if not "perfect".

Matching “versus” mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

As theorists we often ask for "perfect" solutions to problems. Temptation to find *the* solution.

- Mechanism design: find the optimum
- Axiomatic approach: get a unique characterization

Danger: we end up only paying attention to the problems we can solve beautifully. "Keys under the lamppost" joke.

I think there is value in designing mechanisms that are "better" even if not "perfect". Data can play an important role in such theorizing to yield a “sense of magnitudes” (Roth, 2002)

Matching "versus" mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

As theorists we often ask for "perfect" solutions to problems. Temptation to find *the* solution.

- Mechanism design: find the optimum
- Axiomatic approach: get a unique characterization

Danger: we end up only paying attention to the problems we can solve beautifully. "Keys under the lamppost" joke.

I think there is value in designing mechanisms that are "better" even if not "perfect". Data can play an important role in such theorizing to yield a "sense of magnitudes" (Roth, 2002)

- Comparing the good but but imperfect mechanism to existing mechanisms gives a sense of the magnitude of the achievement

Matching “versus” mechanism design: some reflections

Observation 4: is "perfect" the enemy of the "good"?

As theorists we often ask for "perfect" solutions to problems. Temptation to find *the* solution.

- Mechanism design: find the optimum
- Axiomatic approach: get a unique characterization

Danger: we end up only paying attention to the problems we can solve beautifully. "Keys under the lamppost" joke.

I think there is value in designing mechanisms that are "better" even if not "perfect". Data can play an important role in such theorizing to yield a “sense of magnitudes” (Roth, 2002)

- Comparing the good but but imperfect mechanism to existing mechanisms gives a sense of the magnitude of the achievement
- Comparing the good but imperfect mechanism to a performance upper bound gives a sense of how much work there is left to do!