CSCI 357: Algorithmic Game Theory

Lecture 14: Voting & Social Choice 1

Shikha Singh

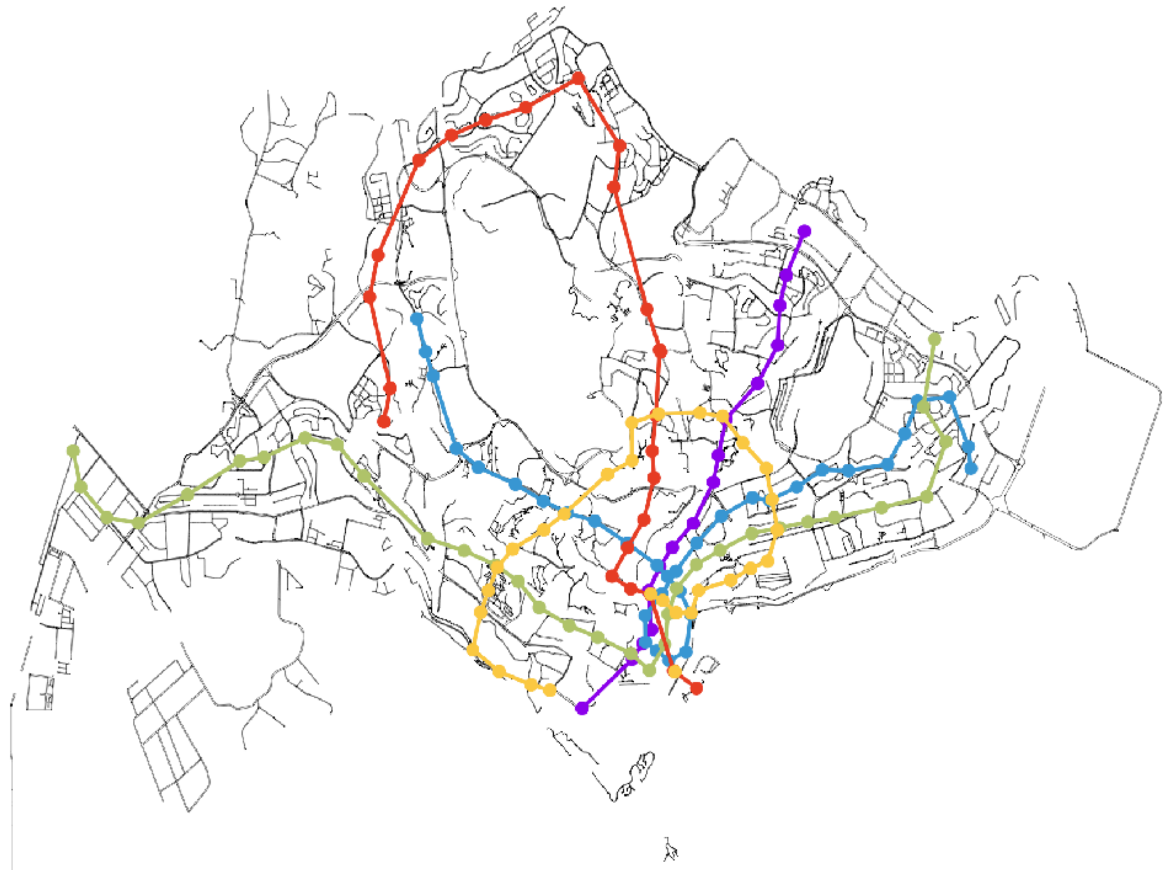

Announcements and Logistics

- Welcome back!
- Homework 6 is due this Thursday at 11 pm
  - Topics from Lecture 11, 12 and 13:
    - Competitive equilibrium
    - One-sided matching and stable matchings
      - Stability definition is different for one-sided & two-sided markets
- Midterm 1 feedback returned March 18
  - Median: 90%, Mean: 88%
  - If you have questions, let me know

Questions?

Plan for Second Half

	Monday	Tuesday	Wednesday	Thursday	Friday
Week 7	Voting & Social Choice			HW 6 due	
Week 8	Extensive-form games, SPNE and Repeated games				
Week 9	Incentives in Distributed systems		HW 8 due	HW 7 due; Project proposal due	
Week 10	Midterm 2			Misc	Project checkpoint 1
Week 11	Misc		Misc	Project checkpoint 2	
Week 12	Project presentations: Checkpoint 3				
Finals period	Final Project Report Due				

Two-Sided Markets: Recap

- Input: A set H of n hospitals, a set S of n students and rankings from each:
  - Each hospital ranks all students
  - Each student ranks all hospitals
- **Goal**: Find a **perfect matching** M (one where each student is matched to exactly one hospital and vice versa) that is **stable** (has no blocking pairs)
- A hospital h and student s form a **blocking pair** (h, s) in a matching M if
  - h prefers s to its current match in M, and s prefers h to its current match in M
- **Algorithm**: Deferred acceptance (DA) by Gale and Shapley
  - Each hospital makes offers to their most favorite that has not rejected them
  - Each student holds on the best offer received and trade up as the algorithm proceeds

Two-Sided Markets: Recap

- **Theorem**. Gale Shapley outputs a stable matching
- A given instance can have many possible stable matchings
  - **Lemma**. Hospital proposing algorithm is hospital optimal (and student pessimal.)
- **Theorem**. Hospital proposing DA algorithm is strategyproof for hospitals but not for students
  - Students can misreport and get a better match!
- **Theorem**. No mechanism for two-sided matching is both stable and strategyproof.
  - Proof developed in Homework 6
  - Uses incomplete lists with stability defined analogously

Project Ideas

- Time to start thinking about what topic you want to do a project on
- Also think about potential project partners and start discussing
- Will share suggested projects but encourage you to explore your interest
- Topics/themes:
  - **Game theory**: evolutionary, sequential games, game theory & AI
  - **Auctions & mechanism design with money**: price of anarchy of auctions, sponsored search, etc
  - **Matching markets**: TTC, stable matchings, school choice, etc
  - **Voting**: strategic issues, rank aggregation etc
  - **Distributed systems**: BitTorrent, network routing, blockchains

Today

- Discuss some generalizations of stable matching
  - With an eye towards project ideas
- Discuss applications of matching markets to school choice
  - Some of this in HW 6
- Move on to voting theory
  - Discuss basic voting algorithms and their properties

Research on Matching Markets

Strategic Behavior in DA

- Truncation strategy: in hospital-proposing DA, a student can truncate their list at their best achievable partner and ensure they are matched to them
- Optimal cheating strategy when complete lists are required?
- How susceptible is the algorithm to manipulation?
    - If the number of stable partners is low, manipulation has little bite

Gale-Shapley Stable Marriage Problem Revisited: Strategic Issues and Applications

Chung-Piaw Teo • Jay Sethuraman • Wee-Peng Tan

Stable Husbands

Donald E. Knuth, Rajeev Motwani, and Boris Pittel
Computer Science Department, Stanford University

Abstract. Suppose n boys and n girls rank each other at random. We show that any particular girl has at least $(\frac{1}{2} - \epsilon)\ln n$ and at most $(1+\epsilon)\ln n$ different husbands in the set of all Gale/Shapley stable matchings defined by these rankings, with probability approaching 1 as $n \to \infty$, if ϵ is any positive constant. The proof emphasizes general methods that appear to be useful for the analysis of many other combinatorial algorithms.

Marriage, Honesty, and Stability

Nicole Immorlica* Mohammad Mahdian*

Abstract

Many centralized two-sided markets form a matching between participants by running a stable marriage algorithm. It is a well-known fact that no matching mechanism based on a stable marriage algorithm can guarantee truthfulness as a dominant strategy for participants. However, as we will show in this paper, in a probabilistic setting where the preference lists of one side of the market are composed of only a constant (independent of the the size of the market) number of entries, each drawn from an *arbitrary* distribution, the number of participants that have more than one stable partner is vanishingly small. This proves (and generalizes) a conjecture of Roth and Peranson [23]. As a corollary of this result, we show that, with high probability, the truthful strategy is the best response for a given player when the other players are truthful. We also analyze equilibria of the deferred acceptance stable marriage game. We show that the game with complete information has an equilibrium in which a $(1-o(1))$ fraction of the strategies are truthful in expectation. In the more realistic setting of a game of incomplete information, we will show that the set of truthful strategies form a $(1+o(1))$-approximate Bayesian-Nash equilibrium. Our results have implications in many practical settings and were inspired by the work of Roth and Peranson [23] on the National Residency Matching Program.

Stable Matching Generalizations

- Many to one matching:
  - Hospitals have a capacity c and can accept that many students
  - Stability defined similarly
- Similar deferred acceptance generalizes
- Many results carry over but no longer strategproof even on one side
  - No stable matching is strategyproof for hospitals in hospital-proposing DA
- If graph is general (not bipartite): **stable roommates problem**
  - No stable matching exists!
  - Approximately stable matchings are studied

Incomplete Preferences & Ties

- If preferences are incomplete and no ties:
  - Set of unmatched people stays the same in all stable matchings
  - Slight modification of Gale Shapley computes stable matching
  - **Open**: how does size of matching relate to size of preference lists?
    - **KMQ 2021**: For random matching markets if lists are at least $\Omega(\log^2 n)$ in size then matching is perfect whp, for $o(\log^2 n)$ size lists, not perfect w.h.p
    - A tight bound on size of matching not known even for random markets
- Incompletions **and** ties: the problem of finding the max matching is **NP hard**
  - Several approximations studied, best known approximation ratio 1.5
  - Most recent (**LM 2021** result) shows $1 + 1/e$ approximation for one-sided ties

Effect of Balance (Competition)

- What if there are n hospitals and $m < n$ students?
- **[AKL '13]** Size of core is a knife edge, and short side enjoys significant advantage. Follow up **[KMQ '21]** extends to incomplete lists.

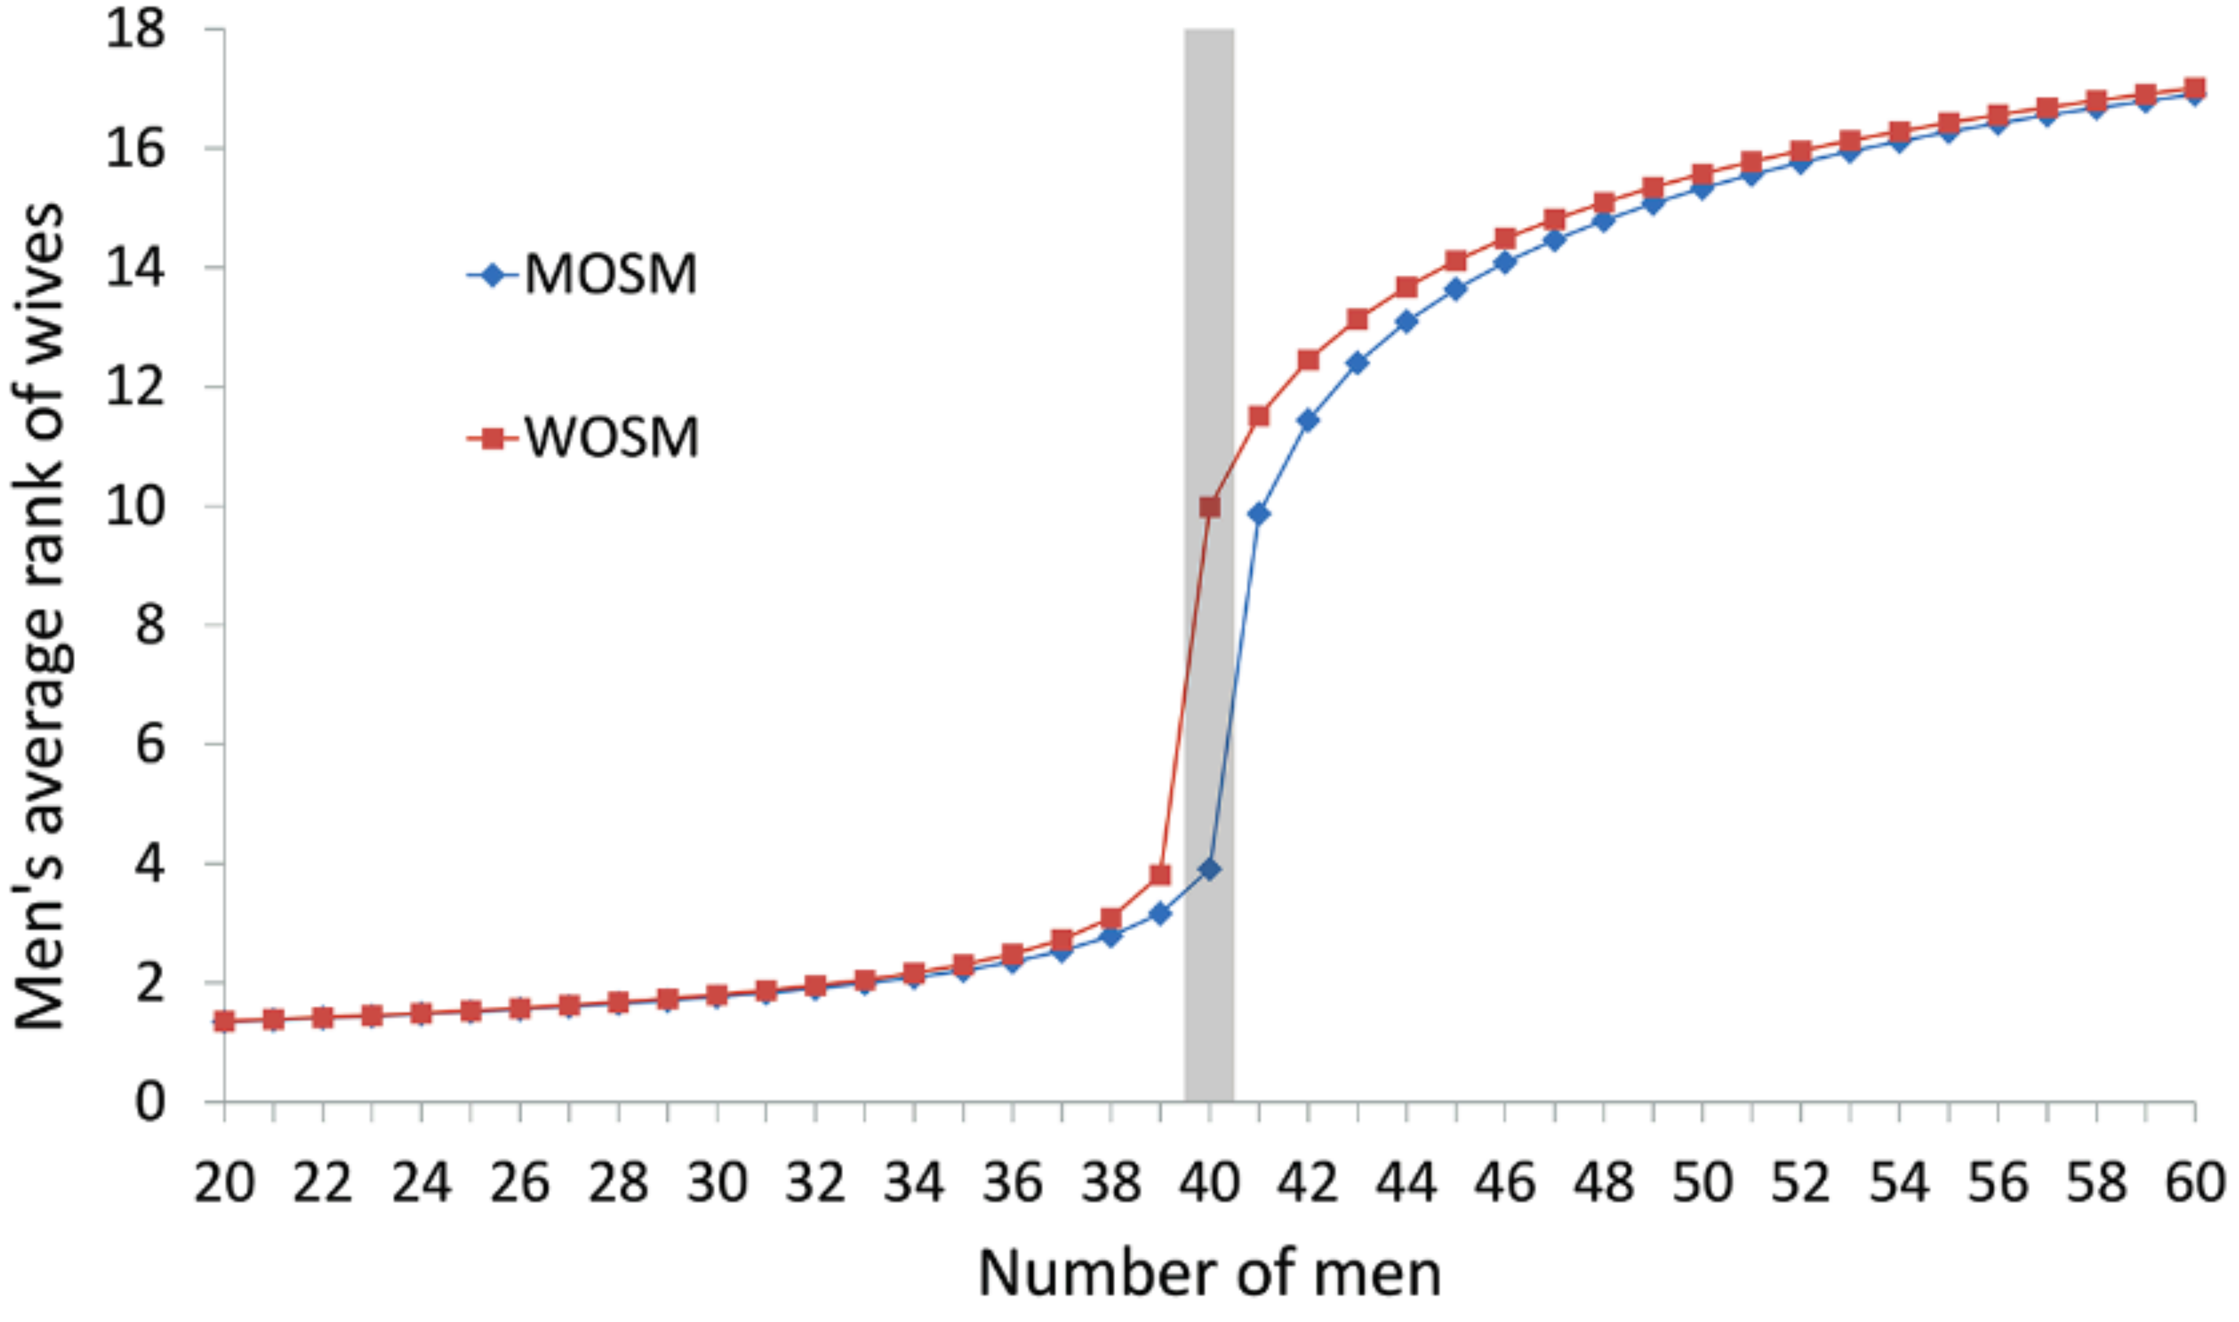

FIG. 2.—Men's average rank of wives under MOSM and WOSM in random markets with 40 women and a varying number of men. The lines indicate the average over 10,000 realizations.

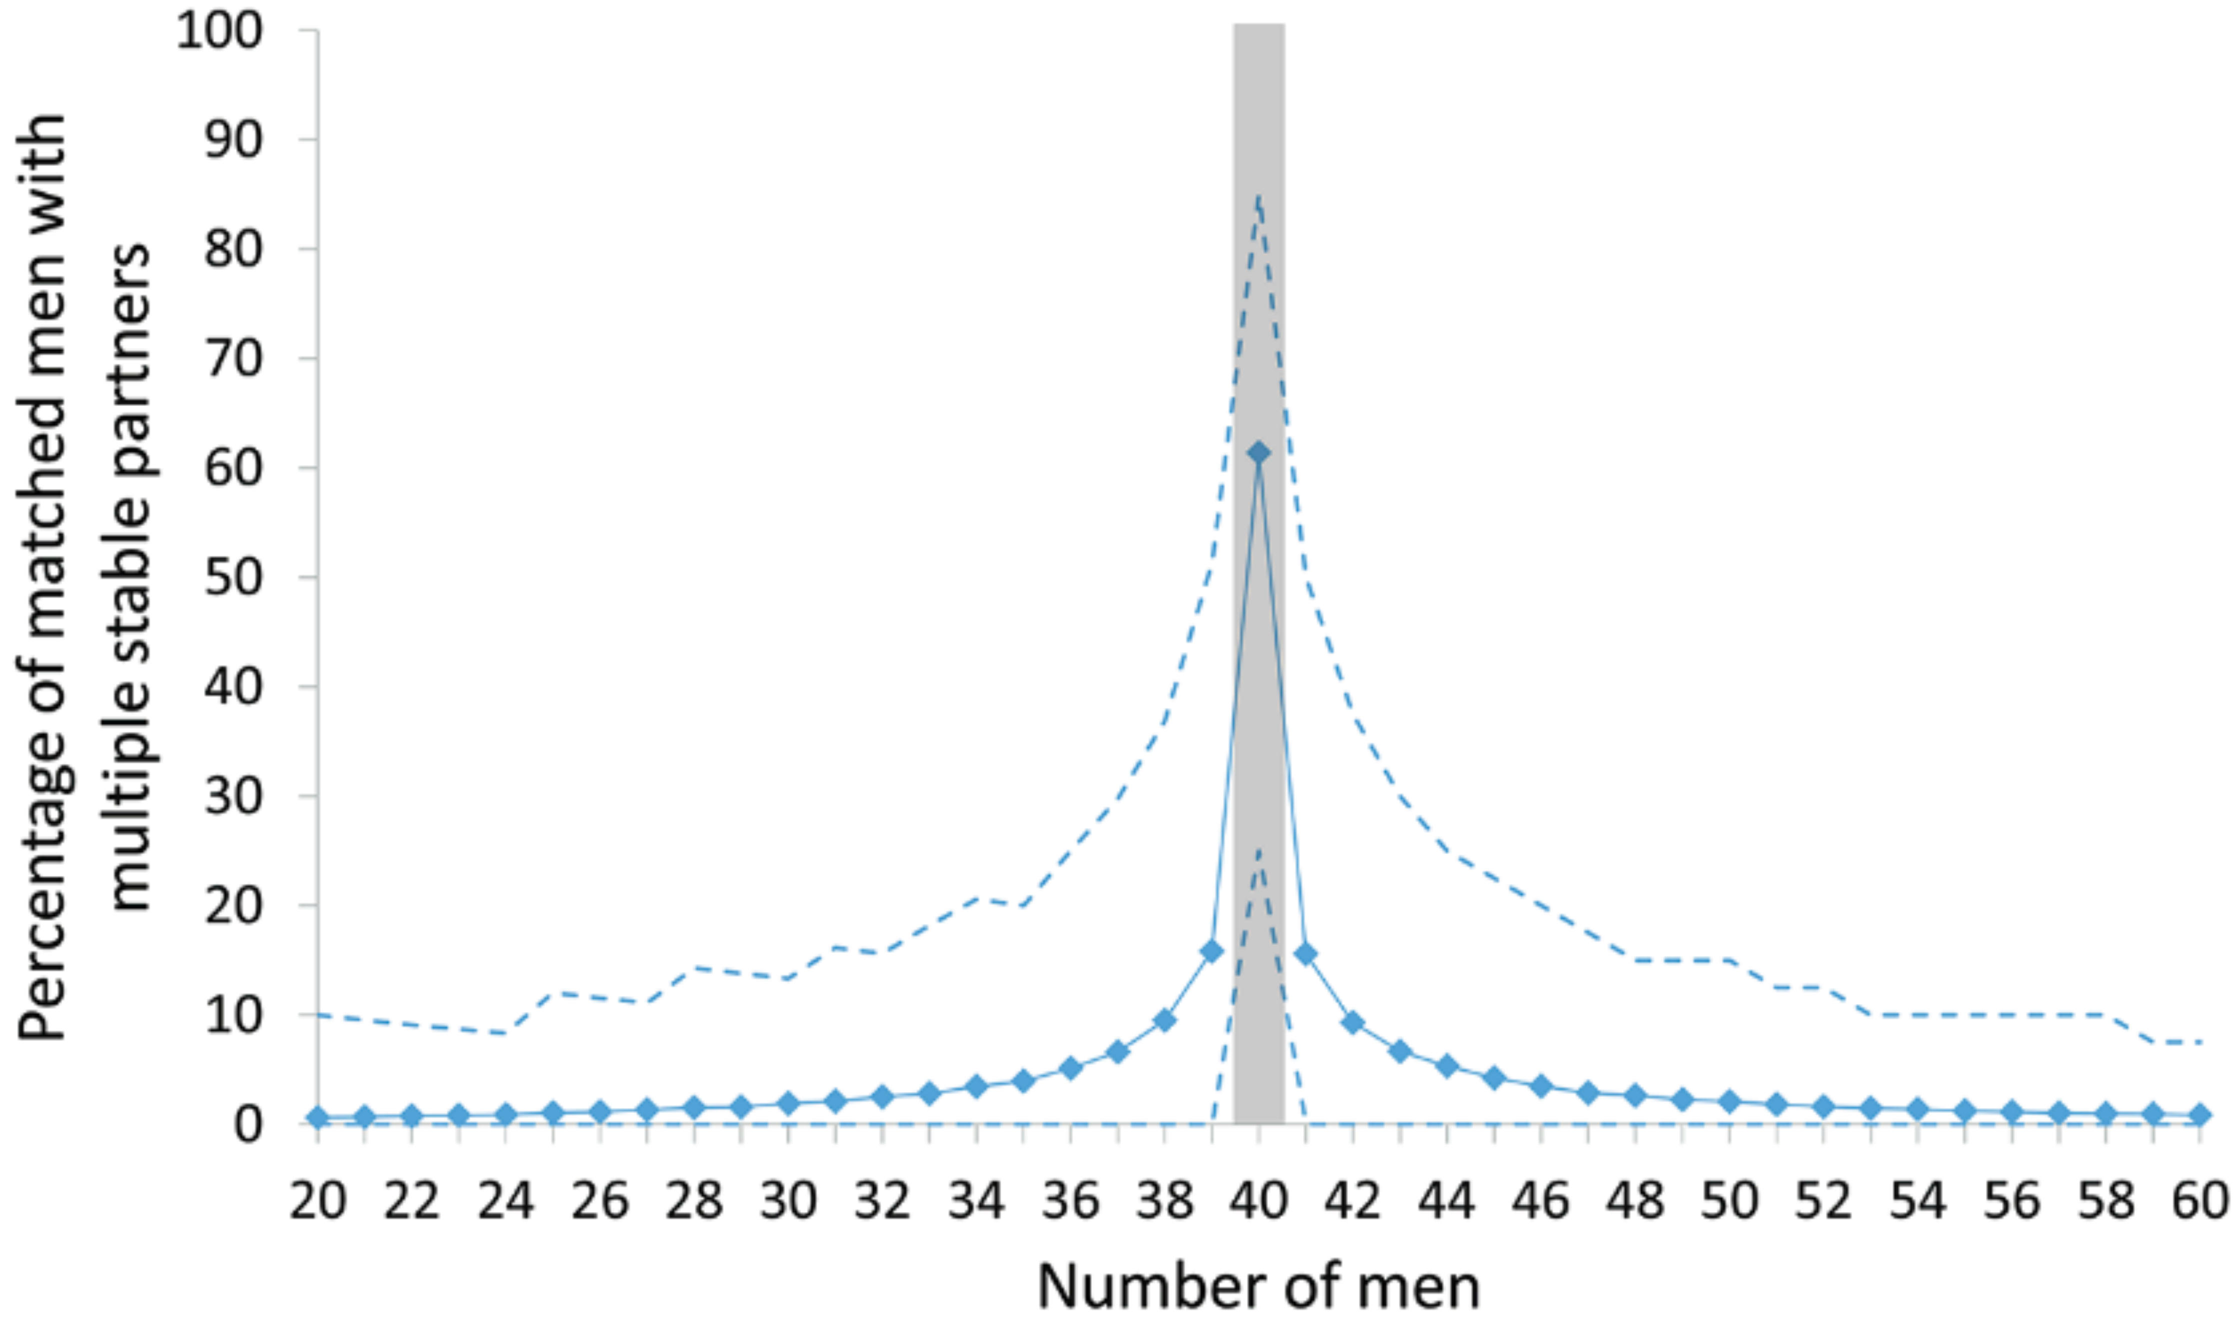

FIG. 1.—Percentage of men with multiple stable partners, in random markets with 40 women and a varying number of men. The main line indicates the average over 10,000 realizations. The dotted lines indicate the top and bottom 2.5th percentiles.

Many More Recent Results

Jan 2022

The Short-Side Advantage in Random Matching Markets

Linda Cai * Clayton Thomas †

Two-sided matching markets with correlated random preferences have few stable pairs

HUGO GIMBERT, Université de Bordeaux, LaBRI, CNRS, F-33400 Talence, France
CLAIRE MATHIEU, Université de Paris, IRIF, CNRS, F-75013 Paris, France
SIMON MAURAS, Université de Paris, IRIF, CNRS, F-75013 Paris, France

2021

2021

Tiered Random Matching Markets: Rank Is Proportional to Popularity

Itai Ashlagi
Department of Management Science and Engineering, Stanford University, CA, USA
iashlagi@stanford.edu
Mark Braverman
Department of Computer Science, Princeton University, NJ, USA
mbraverm@cs.princeton.edu
Amin Saberi
Department of Management Science and Engineering, Stanford University, CA, USA
saberi@stanford.edu
Clayton Thomas
Department of Computer Science, Princeton University, NJ, USA
claytont@cs.princeton.edu
Geng Zhao
Department of Computer Science, Stanford University, CA, USA
gengz@stanford.edu

In which matching markets does the short side enjoy an advantage?*

Yash Kanoria† Seungki Min‡ Pengyu Qian§

2021

2018

Stable matching mechanisms are not obviously strategy-proof ☆

Itai Ashlagi [a,*], Yannai A. Gonczarowski [b,c,*]

[a] *Management Science & Engineering, Stanford University, United States of America*
[b] *Einstein Institute of Mathematics, Rachel & Selim Benin School of Computer Science & Engineering, and the Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel*
[c] *Microsoft Research, Israel*

Received 30 July 2017; final version received 14 March 2018; accepted 9 July 2018

2022

On Fairness and Stability in Two-Sided Matchings

Gili Karni
Weizmann Institute of Science, Rehovot, Israel
Guy N. Rothblum
Weizmann Institute of Science, Rehovot, Israel
Gal Yona
Weizmann Institute of Science, Rehovot, Israel

2019

A $(1+1/e)$-Approximation Algorithm for Maximum Stable Matching with One-Sided Ties and Incomplete Lists*

Chi-Kit Lam C. Gregory Plaxton

Applications to School Choice

- Historically in US, children go to neighborhood schools
- Recently, many cities (NYC, Boston, San Francisco, etc) have adopted school choice programs to give families flexibility and to create competition between schools
- **Goals**: Pareto efficiency, stability to reduce envy, fairness
- Many mechanism used currently are flawed:
    - (Boston choice: Problem 3 on HW 6) is not strategyproof and unfair
    - **[AS '03]** show that up to a quarter of the students ended up being unassigned
- Columbus (OH): essentially a lottery with a short fuse for acceptance
- NYC (90,000+ students applying for high school) submit up to 5 applications
    - Roughly 30,000 students end up being unassigned and then manually assigned

TTC and Student-Proposing DA

- Boston and NYC have both adopted student-proposing DA
- Downside: when school's preferences are disregarded, outcome can be inefficient!
  - HW 6, Problem 4a
- TTC on the other hand is efficient

The Cutoff Structure of Top Trading Cycles in School Choice*

Jacob D. Leshno Irene Lo†

November 7, 2017

"One of the main barriers to using TTC in practice is the difficulty designers had in communicating it to parents and school boards. The standard explanation of TTC is an algorithmic description in terms of sequentially clearing trade cycles, from which it is not directly apparent how priorities and preferences determine assignment. This makes it difficult for school boards to evaluate the effects of policy decisions on the TTC assignment and resulting welfare. It is also difficult for students to verify they were correctly assigned and be convinced that the mechanism is strategyproof."

Takeaway: School Choice

- Active area of research in matching markets
- Lots of avenue for projects

MINIMIZING JUSTIFIED ENVY IN SCHOOL CHOICE:
THE DESIGN OF NEW ORLEANS' ONEAPP

Atila Abdulkadiroglu
Yeon-Koo Che
Parag A. Pathak
Alvin E. Roth
Olivier Tercieux

ABSTRACT

In 2012, New Orleans Recovery School District (RSD) became the first U.S. district to unify charter and traditional public school admissions in a single-offer assignment mechanism known as OneApp. The RSD also became the first district to use a mechanism based on Top Trading Cycles (TTC) in a real-life allocation problem. Since TTC was originally devised for settings in which agents have endowments, there is no formal rationale for TTC in school choice. In particular, TTC is a Pareto efficient and strategy-proof mechanism, but so are other mechanisms. We show that TTC is constrained-optimal in the following sense: TTC minimizes justified envy among all Pareto efficient and strategy-proof mechanisms when each school has one seat. When schools have more than one seat, there are multiple possible implementations of TTC. Data from New Orleans and Boston indicate that there is little difference across these versions of TTC, but significantly less justified envy compared to a serial dictatorship.

Playing on a Level Field:
Sincere and Sophisticated Players in the Boston Mechanism with a Coarse Priority Structure

Moshe Babaioff   Yannai A. Gonczarowski   Assaf Romm*

June 9, 2020

Abstract

Who gains and who loses from a manipulable school-choice mechanism? Studying the outcomes of sincere and sophisticated students under the manipulable Boston Mechanism as compared with the strategy-proof Deferred Acceptance, we provide robust "anything-goes" theorems for large random markets with coarse priority structures. I.e., there are many sincere and sophisticated students who prefer the Boston Mechanism to Deferred Acceptance, and vice versa. Some populations may even benefit from being sincere (if also perceived as such). Our findings reconcile qualitative differences between previous theory and known empirical results. We conclude by studying market forces that can influence the choice between these mechanisms.

Voting and Social Choice

Social Choice

- In social choice theory, we focus on the following question:
  how to aggregate preferences and make decisions that is representative of the collective interests of a group of agents
- Includes topics like
  - Voting to elect a winner or to aggregate preferences and select a ranking
  - Participatory democracy: budgeting decisions
  - Fair division: how to divide indivisible goods fairly (cake-cutting problems)
- No money or transfers involved
  - Mechanism design without money

Social Choice: Model

- A set A of **alternatives**, e.g. different webpages for a search engine to rank or candidates in an election
- A set $N = \{1,2,\ldots,n\}$ of agents or voters
- Each agent $i \in N$ has a strict preference order alternatives: given by a ranked list L_i that agent submits to the mechanism
    - No guarantee of truthful reporting
- Most of what we discuss can be extended to handle non-strict preference orders or truncated preferences orders:
    - We stick to complete strictly ranked lists for simplicity
- Voting rules can have two forms: either output is a **single alternative** (the "winner") or a **full ranked list**

Voting Rules

- **Social-choice function** selects a single alternative for a given preferences profile, that is, $L_1, L_2, \ldots, L_n \mapsto a^*$ where $a^* \in A$
  - Selecting a winner in an election
- **Social-ranking function** selections a rank order of alternatives for a given preference profile, that is, $L_1, L_2, \ldots, L_n \mapsto L^*$ where L^* is a ranking of A
  - Sometimes also referred to as a **social-welfare** function
  - Computing ranking of job candidates
  - Rank aggregation in crowdsourcing applications

Voting vs Matchings

- **Similarity**: Each participant submits a ranked preferences list and the mechanism choose an outcome
- Alternative set A in matching problems: set of all possible matchings
- How is the mechanism design problem of matching different from voting?
  - Social choice framework is general enough to capture matching markets
  - Matching problems had additional "nice structure": agents only cared about their own allocation, not others
  - In contrast, in an election the outcome affects everyone
- Turns out that such a restriction on the possible preferences is key to designing strategyproof mechanisms!

Common Voting Algorithms

Majority Voting

- Suppose there are only two alternates ($|A| = 2$)
- An obvious voting rule is majority vote:
  - Elect the alternative that appears first in the largest number of voters' lists (to avoid ties say n is odd)
  - If outputting a ranking, output the most preferred candidate followed by the second
- Is this majority rule strategyproof?
  - Suppose your preference is $a \succ b$ and you submit $b \succ a$
  - Can only cause the less favored candidate b to be chosen
- Is the story so simple for more than two alternatives?

Plurality Rule

- Suppose there are at least three alternatives ($|A| \geq 3$)
- Suppose we care only about electing a winner, what is the analog of majority rule?
    - If some candidate appears first in more than half of the voters' list, then it is clear that she should be the winner
    - However with 3 or more candidates, this may not occur
    - E.g., you may get a $40/35/25$ split
- In most countries (including US), you use the **plurality rule**: elect the candidate with the most first-place votes
    - Thus, all voters only need to give their 1st preference
- **Questions**. Is this a good voting rule? Is it strategyproof?

2000 US Presidential Election

- To consider the problems with plurality rule, we look back to the 2000 US Presidential election (Bush vs. Gore)
- The race was very close and the outcome came down to the state of Florida
  - Final vote tallies in FL (ignoring other candidates):

Candidate	Party	Vote Total
Bush	Republican	2,912,790
Gore	Democrat	2,912,253
Nader	Green	97,488

- Only a 500 vote difference between Bush and Gore
- It is generally assumed that most voters who viewed Nader as their 1st choice, preferred Gore to Bush
  - Nader was a "spoiler" candidate: his presence flipped the election result even though he couldn't possibly have won
- This example also shows why plurality rule is not strategyproof
- Can you see why?

Plurality Rule Pathologies

- For winner selection, plurality tends to be biased towards "extreme candidates"
- For example, suppose there are 10 "mainstream" candidates (all very similar viewpoints) and 1 "extreme candidate"
  - Suppose 90% of the voters prefer a mainstream candidate to the extreme candidate, 10% prefer the extreme choice
  - If the mainstream candidates manage to split the 90% of the vote equally, they each get 9% of first-place votes
  - This makes the extreme candidate the winner, even though in "pairwise" comparisons, the person would never win
- This is the reason voting theorists are not a fan of Plurality rule

Fairness Criterion: Condorcet

Condorcet Criterion

- An alternate a **beats** b if a majority of voters prefer a to b in a pairwise comparison
- **Condorcet winner:** an alternative that defeats every other alternative
- A social choice function f satisfies the **Condorcet criterion** (is Condorcet consistent) if f selects a Condorcet winner (whenever one exists)
- Does a Condorcet winner always exist?
  - Consider $A = \{a, b, c\}$ and following ballots:
  - Voter 1: a, b, c, Voter 2: b, c, a, Voter 3: c, a, b
  - b defeats c, c defeats a, and a defeats b
- Considered to be a fairness criterion in voting theory
- Plurality does not satisfy Condorcet criterion !

Ranked-Choice Voting

- Alternative to **plurality**: also called single-transferable vote (STV) or instant-runoff voting

YES ON 2
MORE VOICE

The New York Times

After New York Tests a New Way of Voting, Other Cities May Do the Same

Elected leaders and voters in New York remain split over the ranked-choice system, but officials in Washington and elsewhere like the results.

The New York Times

THE MORNING NEWSLETTER

A Guide to Ranked-Choice Voting

The New York mayor's race is the latest example of a ranked-choice election. We offer a strategic explainer.

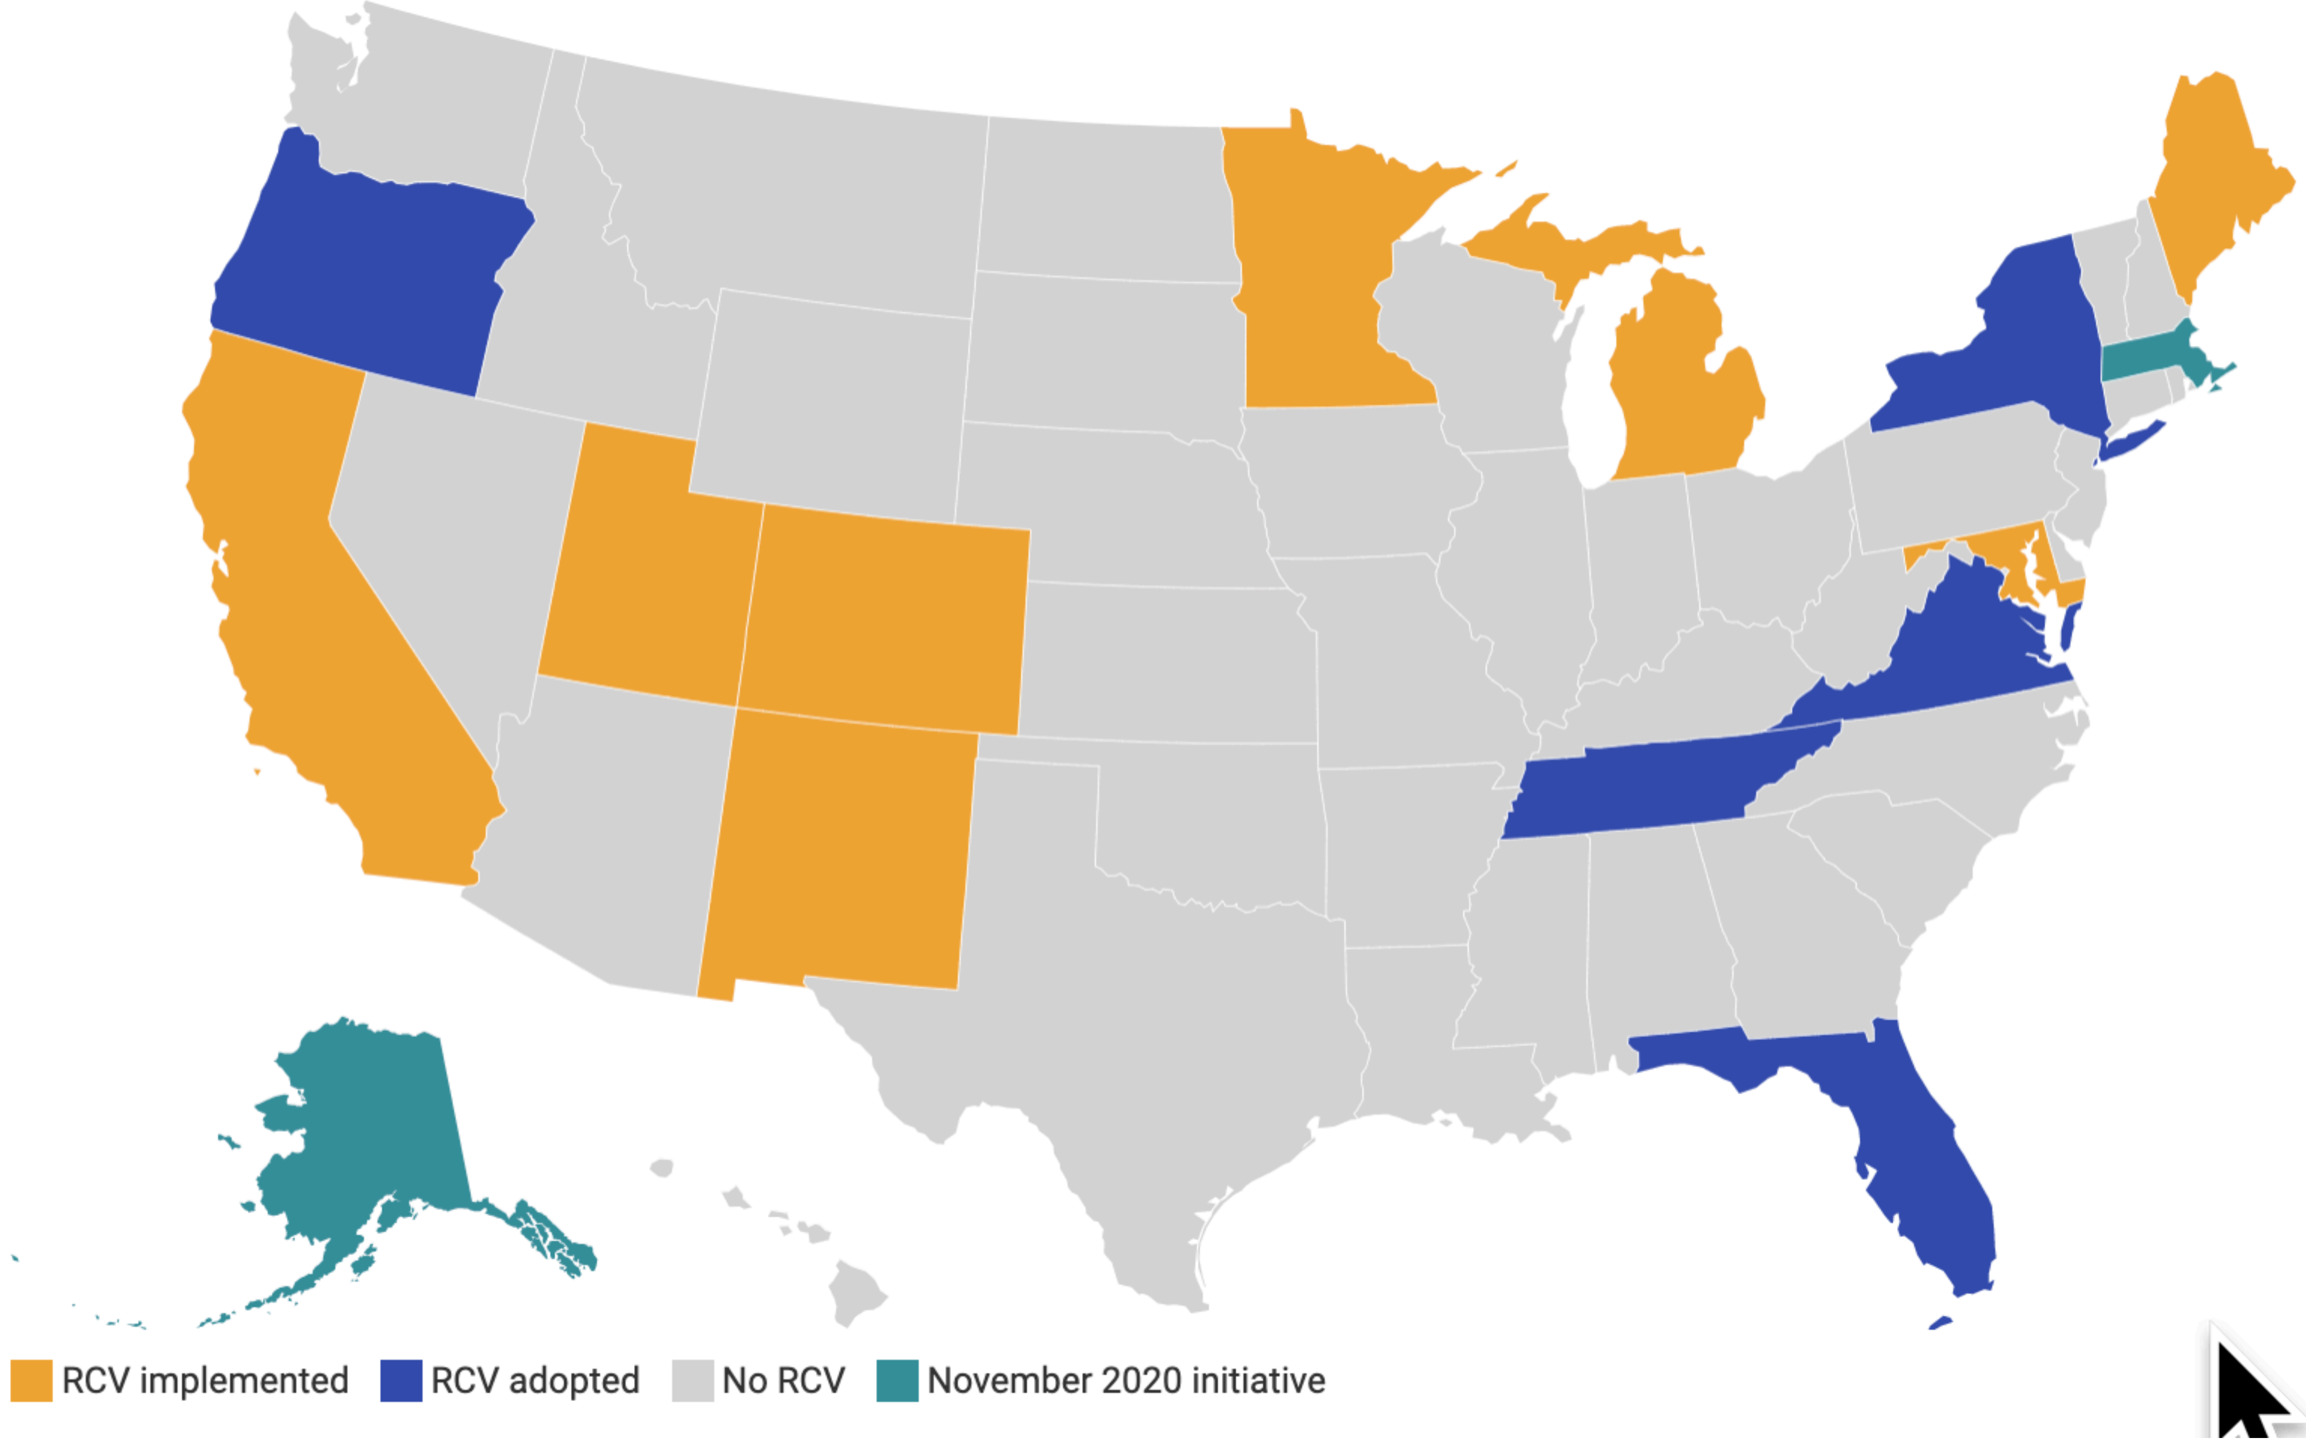

Source: https://www.southcoasttoday.com/news/20201003/in-massachusetts-yes-on-2-pitches-ranked-choice-voting

Ranked-Choice Voting

- Alternative to **plurality**: also called single-transferable vote (STV) or instant-runoff voting
- Voters submit a full ranked list (not just their first choice)
- (**Majority rule**) If there is some alternative a^* that receives more than 50% of the first-place voters, then a^* is the winner
- Otherwise, the alternative with the fewest first-place votes is deleted and the winner is computed recursively on the rest
    - Base case: only two alternatives left, use majority rule
- Notice that this rule is not biased towards "extreme candidates"
- Various tie-breaking rules used in case of ties

Ranked-Choice Voting

- For example, consider $A = \{1,2,3,4\}$ and 5 voters s.t.

	Voters #1,2	Voters #3,4	Voter #5
1st Choice	a	b	c
2nd choice	d	a	d
3rd choice	c	d	b
4th choice	b	c	a

- Which alternative is eliminated in round 1?
    - d: has zero first-place votes

	Voters #1,2	Voters #3,4	Voter #5
1st Choice	a	b	c
2nd choice	c	a	b
3rd choice	b	c	a

Ranked-Choice Voting

- After c is eliminated in round 2:

	Voters #1,2	Voters #3,4,5
1st Choice	a	b
2nd choice	b	a

- a is eliminated in round 3, so b wins
- Should we be happy with this outcome?
  - Condorcet winner?
- Is this rule strategyproof?
  - Can we see this in our example?

	Voters #1,2	Voters #3,4	Voter #5
1st Choice	a	b	c
2nd choice	d	a	d
3rd choice	c	d	b
4th choice	b	c	a

Ranked-Choice Voting

- Ranked-choice voting is not strategyproof
- **Intuition**: there can be an incentive to influence who gets eliminated early on, so that your preferred candidate gets more favored matchups in later rounds
- Compared to plurality, it seems trickier to figure out a profitable manipulation
    - In fact, even if you know everyone else's vote, the problem of finding a profitable manipulation is **NP hard**
- This is why many voting theorists prefer ranked-choice voting

Single transferable vote resists strategic voting

John J. Bartholdi III[1] and James B. Orlin[2]

[1] School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA
[2] Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Received December 24, 1990 / Accepted May 12, 1991

Abstract. We give evidence that Single Tranferable Vote (STV) is computationally resistant to manipulation: It is NP-complete to determine whether there exists a (possibly insincere) preference that will elect a favored candiate, even in an election for a single seat. Thus strategic voting under STV is qualitatively more difficult than under other commonly-used voting schemes. Furthermore, this resistance to manipulation is inherent to STV and does not depend on hopeful extraneous assumptions like the presumed difficulty of learning the preferences of the other voters. We also prove that it is NP-complete to recognize when an STV election violates monotonicity. This suggests that non-monotonicity in STV elections might be perceived as less threatening since it is in effect "hidden" and hard to exploit for strategic advantage.

Digging Deeper

- Since ranked-choice voting is now being used in elections, there is a need to understand its properties better
- How does it perform under practical (non-worst case) distributions?
    - Random preferences
    - Mallow model of generating real world preferences?
    - Is it still difficult to find a profitable manipulation?
- How robust is the voting rule to perturbations?
- NYC Mayoral data is now public and can be used for analysis

Borda Count

- Well known voting rule: often used in sports, also used in Eurovision song contest
- Voters submit their full ranked lists: an alternate gets $|A|$ for each first-choice vote, $|A| - 1$ points for each second-choice vote, and so on and 1 point for each last-choice vote
- Example:
  - a gets 15 points, b gets 12 points
  - c gets 10 points, d gets 13 points
- Borda count would elect a
  - In contrast to ranked-choice b
- Is Borda count Cordorcet consistent? Show in HW 7.

	Voters #1,2	Voters #3,4	Voter #5
1st Choice	a	b	c
2nd choice	d	a	d
3rd choice	c	d	b
4th choice	b	c	a

Positional Scoring Rules

- In general, you can have different ways to score each position
- For each vote, a **positional-scoring rule** on $m = |A|$ alternatives assigns a score of α_j to the alternative ranked in jth place. The alternative with maximum total score (across all votes) is selected.
    - Assume $\alpha_1 \geq \alpha_2 \geq \ldots \alpha_m$ and $\alpha_1 > \alpha_m$
    - E.g., plurality gives 1 point for first-choice, zero for others
- Many positional scoring rules have been studied
    - You might see some on the homework/ papers you read

Borda Count

- Is Borda count strategyproof?
  - **Idea**: incentive to rank closest competitor to preferred candidate last
- In example, what is the Borda score of a and b?
  - a's score: $2 \cdot 3 + 4 = 10$
  - b's score: $2 * 4 + 3 = 11$
- If voter 3 moves b to the last place
  - b's score: $8 + 1 = 9$

Winner
b

1	2	3
b	b	a
a	a	b
c	c	c
d	d	d

1	2	3
b	b	a
a	a	c
c	c	d
d	d	b

Winner
a

Strategyproof Voting

- There are some trivial strategy proof voting rules:
- A voting rule is **dictatorial** if it has dictator voter i and always elects i's first choice (regardless of others' preferences)
    - Defeats the purpose of voting
- Are there any “reasonable” strategy proof voting rules?

Strategyproof Voting

- Changing the voting rule changes the outcome of the mechanism
- Leads to contention on which voting rule is the “best”
- A desirable property that everyone wants out of voting mechanism: inability to manipulate the outcome
    - We will also see other desirable properties we’d want from a voting mechanism
- None of the voting rules we have seen are strategy proof
- Are we just not being smart enough or is there is a barrier to achieving this property?

Gibbard-Satterthwaite & Arrow's Impossibility Theorems

Impossibility Result

- **Gibbard-Satterthwaite theorem**.
  Any voting/social-choice rule with at least 3 alternatives that is strategyproof and onto must be dictatorial.
  - We will prove this next lecture
- *"The GS theorem seems to quash any hope of designing incentive-compatible social-choice functions. The whole field of Mechanism Design attempts escaping from this impossibility result using various modifications.": Nisan*
- How AGT escapes this sweeping impossibility:
  - Restricting what preference lists are possible (e.g. matching markets)
  - Money!   Thats why we focused so much on mechanism design with money

Arrow's Impossibility Theorem

- The GS theorem is closely related to and can be derived from an even more famous impossibility result: Arrow's theorem
- **Arrow's impossibility theorem**. With three or more alternatives, no social-rank function satisfies the following three properties:
    - Non-dictatorship
    - Unanimity
    - Independence of irrelevant alternatives (IIA)
- Unanimity means if every voter ranks a over b, then the social-rank function should rank a over b
- IIA means that, for every pair a, b of alternatives, the relative order of a over b in the output ranking should be a function of only the relative order of a, b in each voter's list and not depend on the position of any "irrelevant" alternative c in anyone's preferences

Plurality does not satisfy IIA (e.g., Bush vs Gore outcome was affected by Nader)

Arrow's and GS

- One can also derive the Gibbard-Satterthwaite theorem from Arrow's theorem, using a reduction argument
- Suppose we have a non-trivial and strategyproof voting rule
    - Use it to construct a a voting rule that satisfies the three conditions in Arrow's theorem
- Intuitively, not satisfying IIA can lead to opportunities for strategic manipulation
- You also need to ensure technicalities like Arrow's theorem is a result about social-ranking functions (voting rules that produce a full ranked list) while the GS theorem holds even for social choice functions (voting rules that elect a winner)

Single-Peaked Preferences

- Just like computational hardness forces us to compromise our goals (find approximation algorithms, focus on special cases), impossibility results in voting theory force us to be creative
- Most common restriction on preferences considered in the voting landscape:
    - Single-peaked preferences
- Imagine that the candidates are points on a real line
    - Line could represent the political spectrum
- A voter i has single-peaked preferences if there is a "peak" $p_i \in \mathbb{R}$ such that the voters prefers candidates closer to her peak
- Idea is that single-peaked preferences are a reasonable approximation of voter's preferences

Single-Peaked Preferences

- Imagine that the candidates are points on a real line
  - Line could represent the political spectrum
- A voter i has single-peaked preferences if there is a "peak" $p_i \in \mathbb{R}$ such that the voters prefers candidates closer to her peak
- Given single-peaked preferences, how do we select a candidate?
  - Average rule?
  - Median rule?
    - Individual and group strategyproof!
- Median rule is also Pareto optimal and satisfies the Condorcet criterion (will discuss this soon)

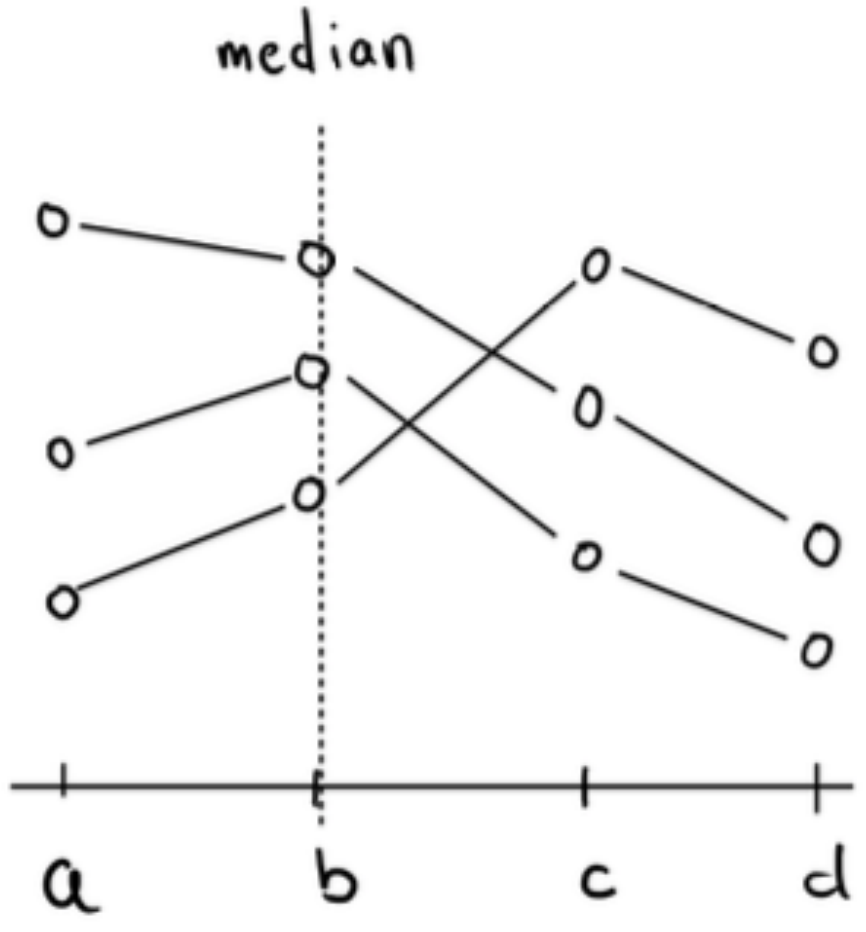

Single-peaked preferences

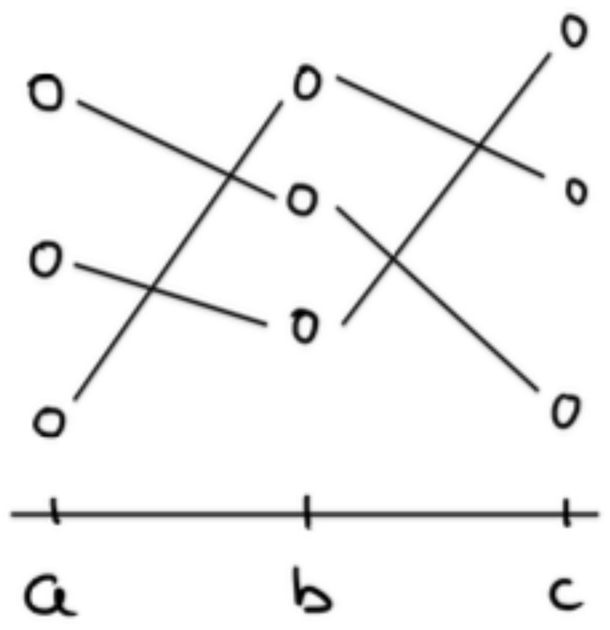

Not single-peaked

Hardness of Manipulation

- Figuring out whether there is a profitable manipulation is intractable for ranked-choice voting (even in the presence of complete information)
    - However, this result holds when the number of alternatives grow (in contrast to voters)
- Unfortunately, NP-hardness just says it is hard for some worst-case instances
- What if it is actually easy for most practical instances?
- Many rules admit polynomial time manipulation algorithms for fixed #alternatives
- Many rules admit polynomial time algorithms that find a successful manipulation on almost all profiles!

Hardness of Manipulation

- Interesting open problem to design voting rules that are hard to manipulate on average
- Very nice and readable article about manipulation in voting (on GLOW)

AI's War on Manipulation: Are We Winning?

Piotr Faliszewski and Ariel D. Procaccia

"The most controversial part of the approach is that it relies on NP-hardness as a measure of computational difficulty.The issue is that NP-hardness is a worst-case notion and the fact that a problem is NP-hard simply means that it has some dif]cult instances and not that necessarily the ones typically occurring in practice are hard to solve. "

Approximate Approaches

- In the vein of approximate solutions in algorithms, one can try to relax the strategyproofness conditions
    - Consider "milder" notions of incentive compatibility

Approximate Strategyproofness

Benjamin Lubin
School of Management
Boston University
blubin@bu.edu

David C. Parkes
School of Engineering and Applied Sciences
Harvard University
parkes@eecs.harvard.edu

July 24, 2012

Abstract

The standard approach of mechanism design theory insists on equilibrium behavior by participants. This assumption is captured by imposing *incentive constraints* on the design space. But in bridging from theory to practice, it often becomes necessary to relax incentive constraints in order to allow tradeoffs with other desirable properties. This paper surveys a number of different options that can be adopted in relaxing incentive constraints, providing a current view of the state-of-the-art.