CSCI 357: Algorithmic Game Theory

Lecture 14: Voting & Social Choice I

Shikha Singh

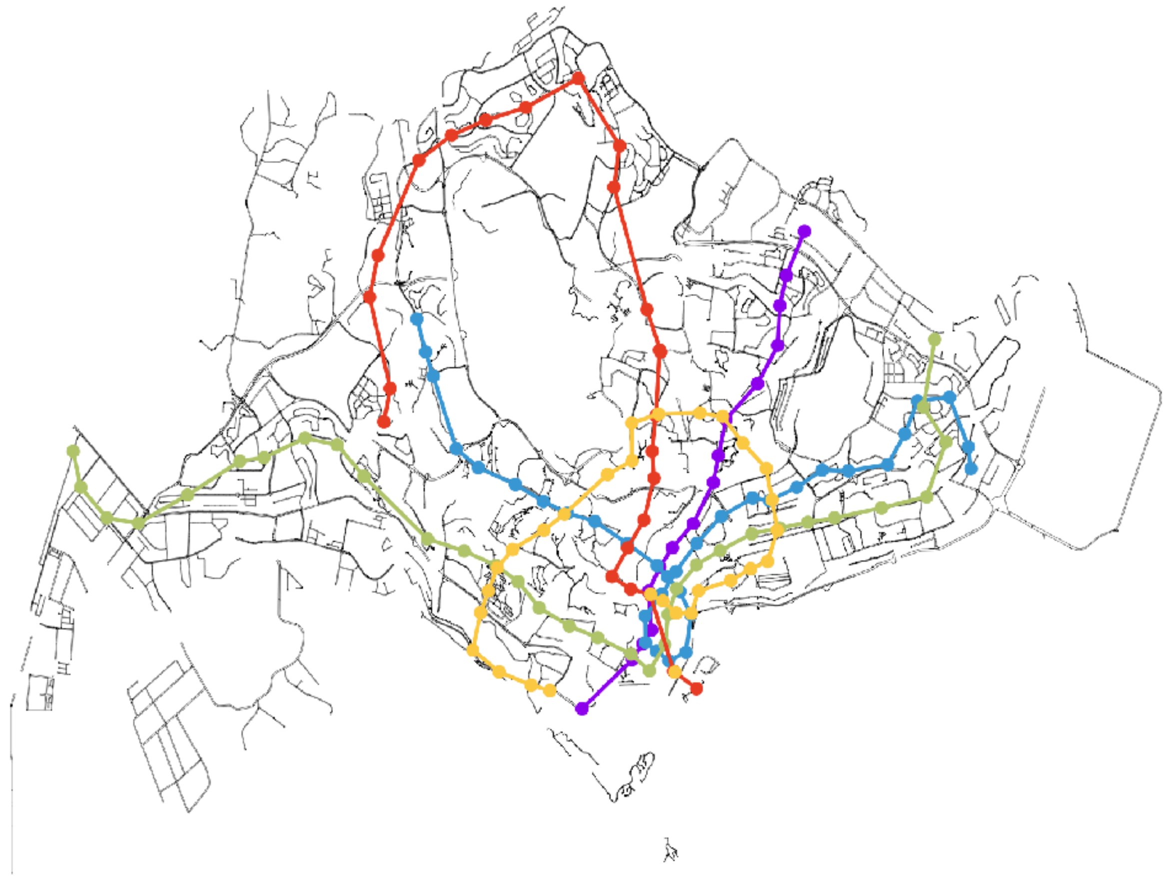

Announcements and Logistics

- Welcome back!
- Assignment 3 is due this Friday at noon
  - Submit jointly with your partner
- Looking ahead: see calendar
- I will be traveling week of April 28
  - Midterm on Tuesday
  - Required project meetings in lieu of Friday's lecture

APRIL 2025

SUN	MON	TUE	WED	THU	FRI	SAT
		1			4	5
	We're here	8		Assignment #3 due	11	12
		15		Paper #3 Eval in class	18	19
		22		Assignment #4 due	25	26
	Midterm 2	29		Paper #4 Eval & Project Checkpoint	2 No Class	3

Questions?

Topics for Second Half

	Monday	Tuesday	Wednesday	Thursday	Friday
Week 7	Voting				
Week 8	Social Choice & Fair Division				
Week 9	Decentralized Markets				
Week 10	Incentives in P2P Systems				Project proposal due
Week 11	Midterm 2 and No class				Project checkpoint 1
Week 12	Incentives in Network Routing & Blockchains				Project checkpoint 2
Finals period	Complexity of Equilibrium				Project presentations
	Final Project Report Due				

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

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

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.

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.

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 & Imbalance

- In general markets, there is competition (imbalance): $n + k$ candidates and n jobs
- Preference lists are not complete: rank only top d choices
- **Open problem**: how does size of matching relate to d and k?
    - [SP'25]: For random matching markets if the preference lists are $\approx \log n \log \frac{n}{k}$ in size then matching is perfect w.h.p, and if shorter then 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-size stable 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

Stark Effect of Competition

- Which side of the market has an advantage in a random matching market?
- **[AKL '13]** Size of core is a knife edge, and short side enjoys significant advantage. Follow up **[KMQ '21]** extends to incomplete lists.

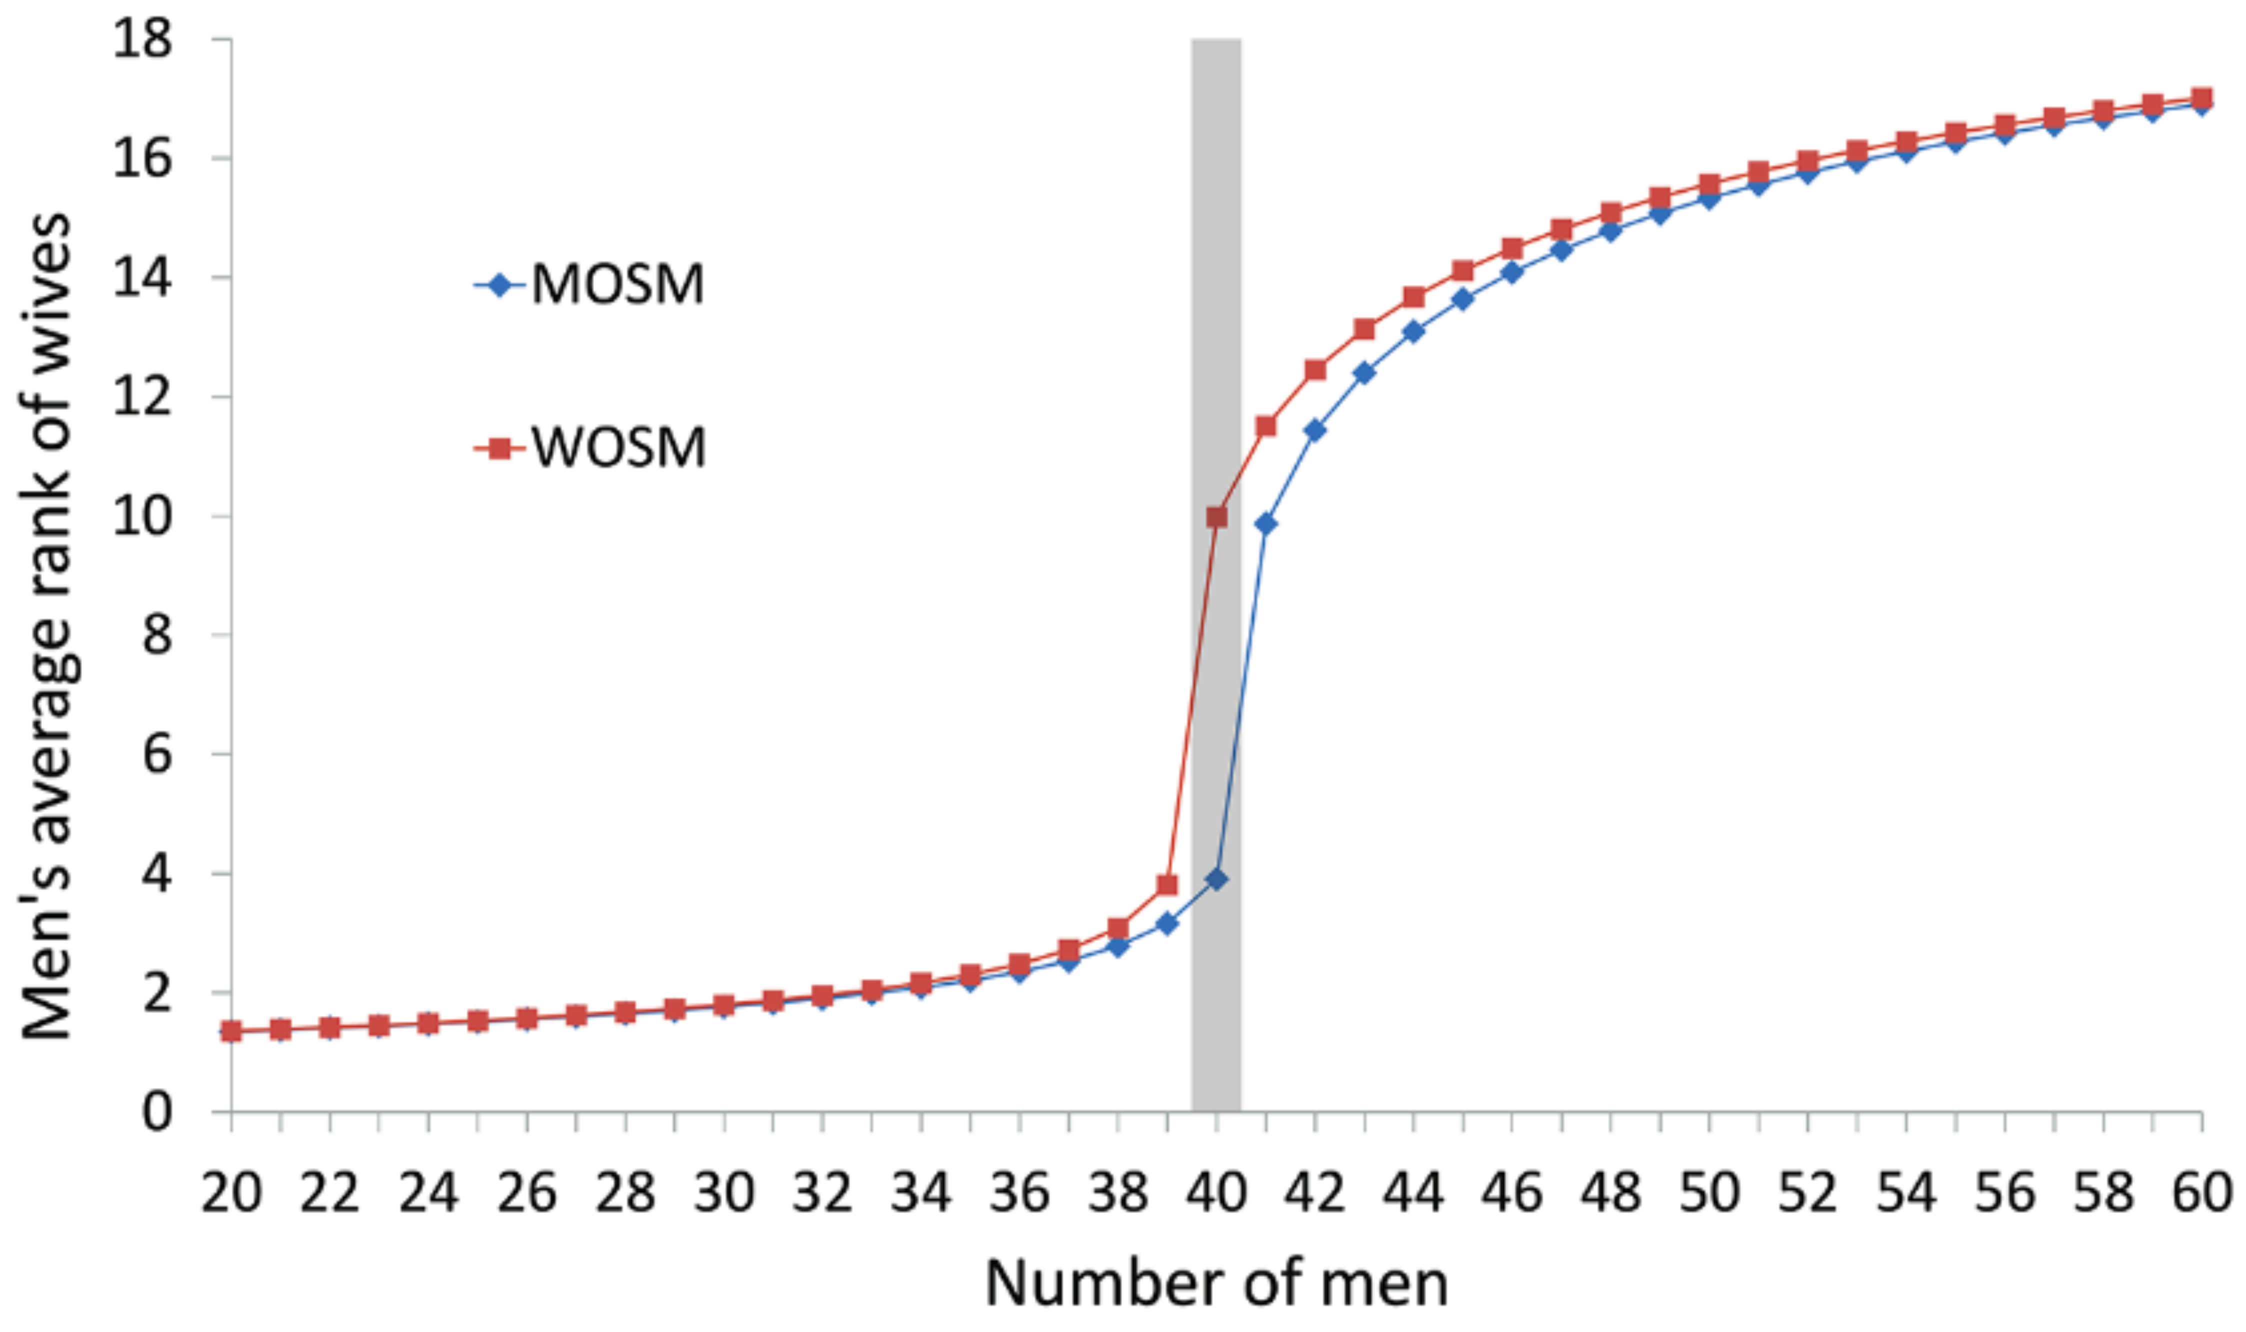

FIG. 2.—Men's average rank of wives under MOSM and WOSM in random markets with 40 women and a varying number of men. The lines indicate the average over 10,000 realizations.

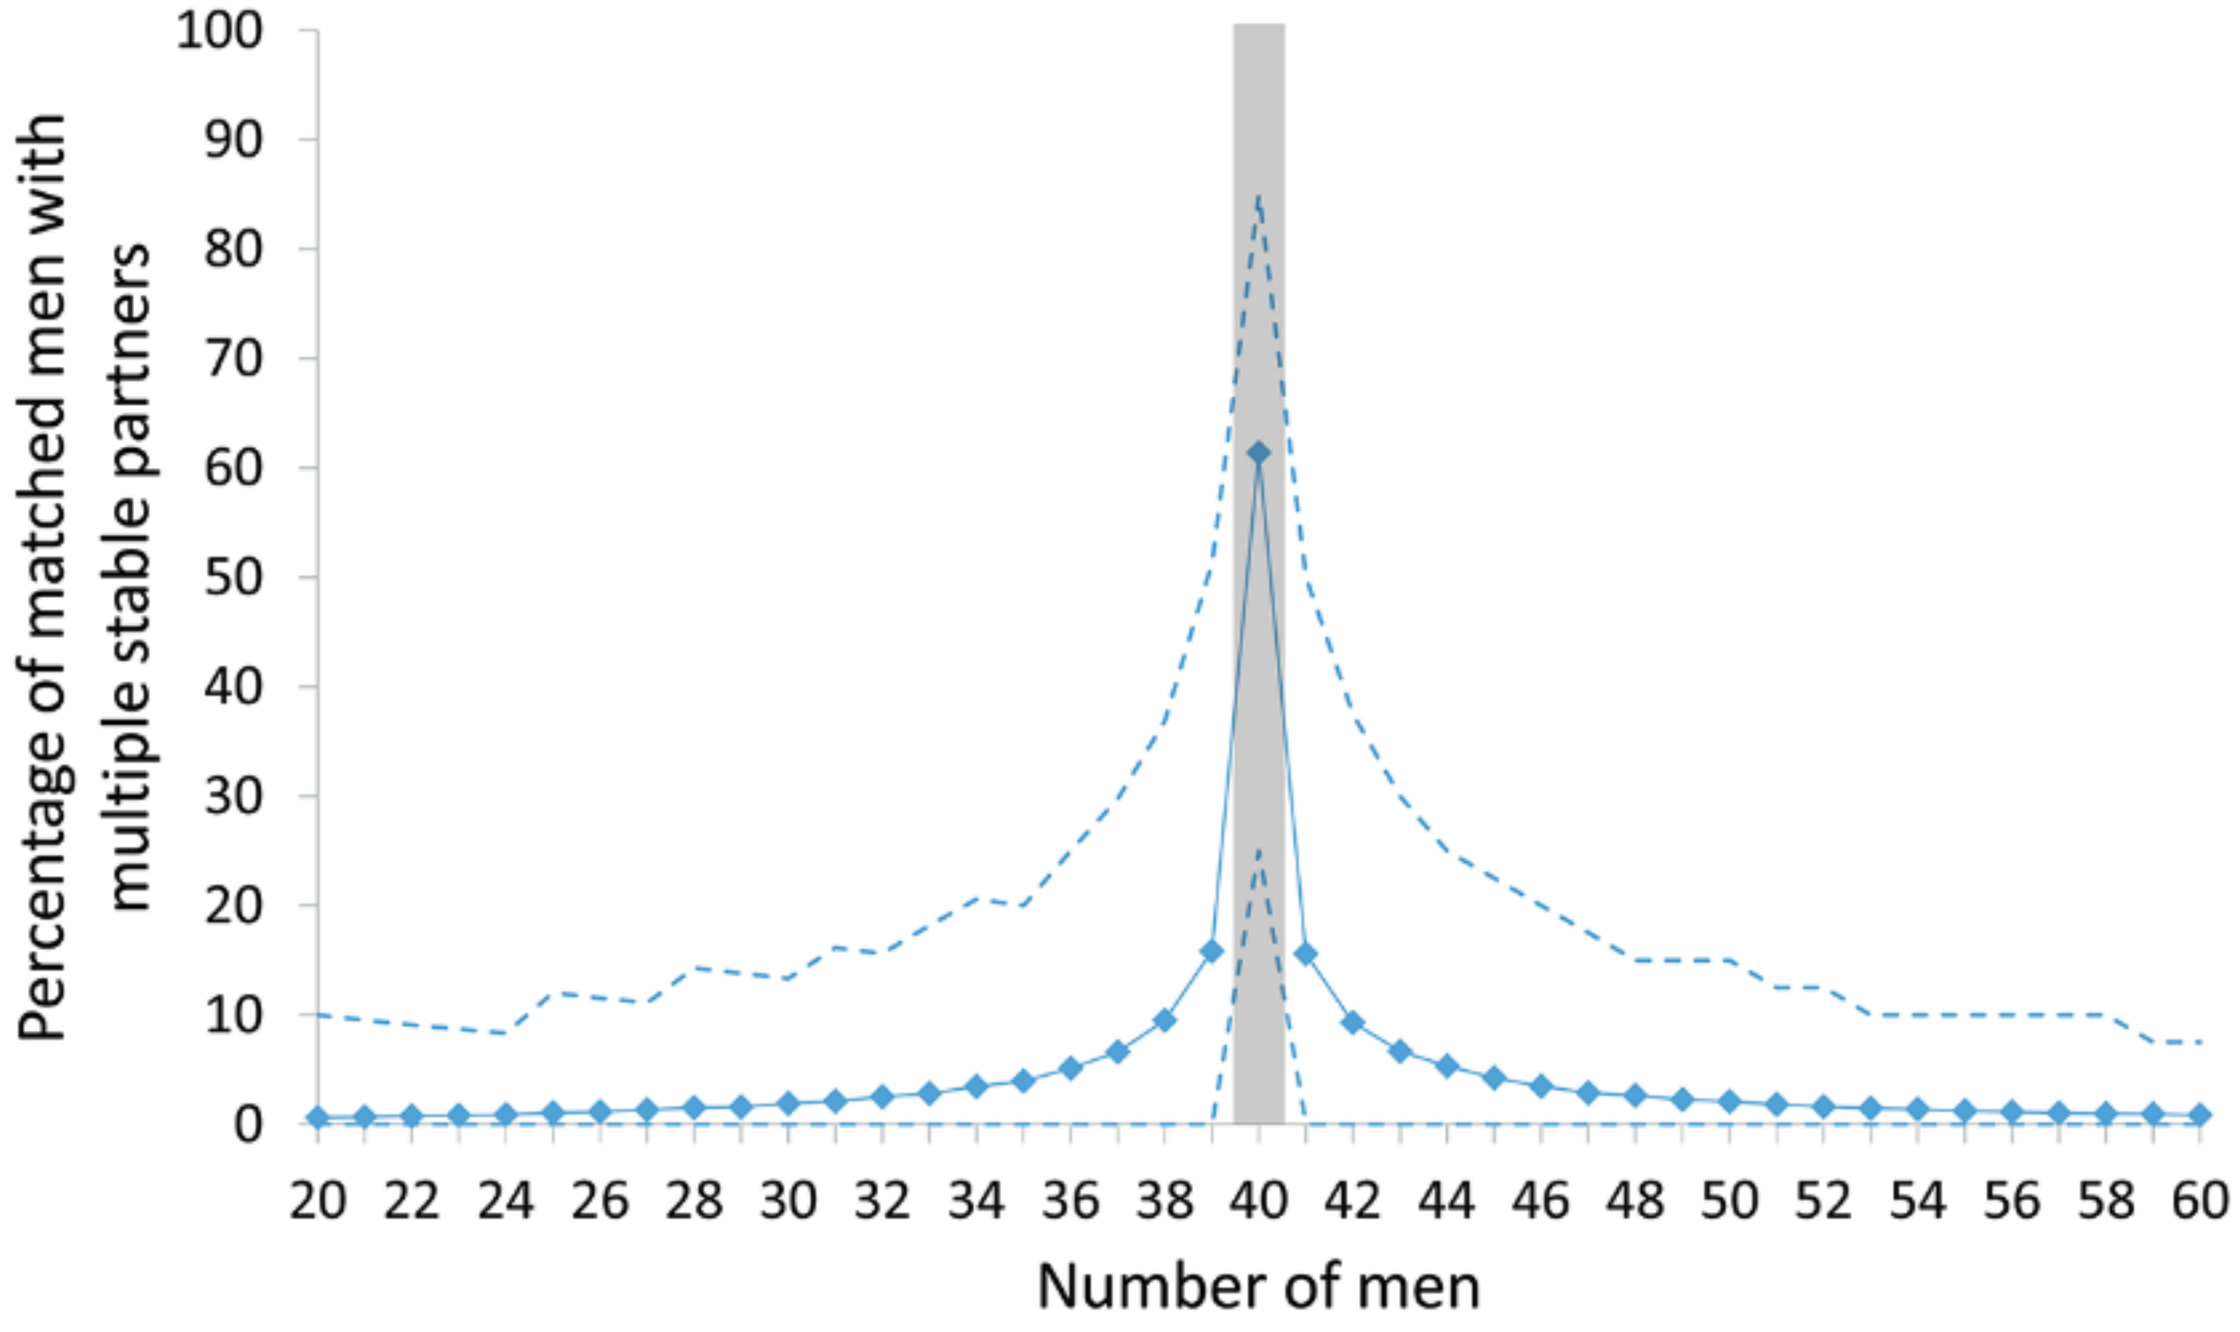

FIG. 1.—Percentage of men with multiple stable partners, in random markets with 40 women and a varying number of men. The main line indicates the average over 10,000 realizations. The dotted lines indicate the top and bottom 2.5th percentiles.

Many More Research Topics

2021

Deferred Acceptance with Compensation Chains

PIOTR DWORCZAK, Stanford University, Graduate School of Business

I introduce a class of algorithms called Deferred Acceptance with Compensation Chains (DACC). DACC algorithms generalize the DA algorithms by Gale and Shapley [1962] by allowing both sides of the market to make offers. The main result is a characterization of the set of stable matchings: a matching is stable if and only if it is the outcome of a DACC algorithm.

Mar 2024

Structural Complexities of Matching Mechanisms*

Yannai A. Gonczarowski[†] Clayton Thomas[‡]

March 30, 2024

Abstract

We study various novel complexity measures for two-sided matching mechanisms, applied to the two canonical strategyproof matching mechanisms, Deferred Acceptance (DA) and Top Trading Cycles (TTC). Our metrics are designed to capture the complexity of various structural (rather than computational) concerns, in particular ones of recent interest within economics. We consider a unified, flexible approach to formalizing our questions: Define a protocol or data structure performing some task, and bound the number of bits that it requires. Our main results apply this approach to four questions of general interest; for mechanisms matching applicants to institutions, our questions are:

(1) How can one applicant affect the outcome matching?
(2) How can one applicant affect another applicant's set of options?
(3) How can the outcome matching be represented / communicated?
(4) How can the outcome matching be verified?

Nov 2024

Stable Matching with Ties: Approximation Ratios and Learning

Shiyun Lin * Simon Mauras † Nadav Merlis ‡ Vianney Perchet §

November, 2024

Abstract

We study the problem of matching markets with ties, where one side of the market does not necessarily have strict preferences over members at its other side. For example, workers do not always have strict preferences over jobs, students can give the same ranking for different schools and more. In particular, assume w.l.o.g. that workers' preferences are determined by their utility from being matched to each job, which might admit ties. Notably, in contrast to classical two-sided markets with strict preferences, there is no longer a single stable matching that simultaneously maximizes the utility for all workers.

We aim to guarantee each worker the largest possible share from the utility in her best possible stable matching. We call the ratio between the worker's best possible stable utility and its assigned utility the *Optimal Stable Share* (OSS)-ratio. We first prove that distributions over stable matchings cannot guarantee an OSS-ratio that is sublinear in the number of workers. Instead, randomizing over possibly non-stable matchings, we show how to achieve a tight logarithmic OSS-ratio. Then, we analyze the case where the real utility is not necessarily known and can only be approximated. In particular, we provide an algorithm that guarantees a similar fraction of the utility compared to the best possible utility. Finally, we move to a bandit setting, where we select a matching at each round and only observe the utilities for matches we perform. We show how to utilize our results for approximate utilities to gracefully interpolate between problems without ties and problems with statistical ties (small suboptimality gaps).

Jan 2025

UNBALANCED RANDOM MATCHING MARKETS WITH PARTIAL PREFERENCES

ADITYA POTUKUCHI AND SHIKHA SINGH

ABSTRACT. Properties of stable matchings in the popular random-matching-market model have been studied for over 50 years. In a random matching market, each agent has complete preferences drawn uniformly and independently at random. Wilson (1972), Knuth (1976) and Pittel (1989) proved that in balanced random matching markets, the proposers are matched to their $\ln n$th choice on average. In this paper, we consider markets where agents have partial (truncated) preferences, that is, the proposers only rank their top d partners. Despite the long history of the problem, the following fundamental question remained unanswered: *what is the smallest value of d that results in a perfect stable matching with high probability?* In this paper, we answer this question exactly—we prove that a degree of $\ln^2 n$ is necessary and sufficient. That is, we show that if $d < (1-\varepsilon)\ln^2 n$ then no stable matching is perfect and if $d > (1+\varepsilon)\ln^2 n$, then every stable matching is perfect with high probability. This settles a recent conjecture by Kanoria, Min and Qian (2021).

We generalize this threshold for unbalanced markets: we consider a matching market with n agents on the shorter side and $n(\alpha+1)$ agents on the longer side. We show that for markets with $\alpha = o(1)$, the sharp threshold characterizing the existence of perfect stable matching occurs when d is $\ln n \cdot \ln\left(\frac{1+\alpha}{\alpha+(1/n(\alpha+1))}\right)$.

Finally, we extend the line of work studying the effect of imbalance on the expected rank of the proposers (termed the "stark effect of competition"). We establish the regime in unbalanced markets that forces this stark effect to take shape in markets with partial preferences.

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

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

Matching Application: Kidney Exchange

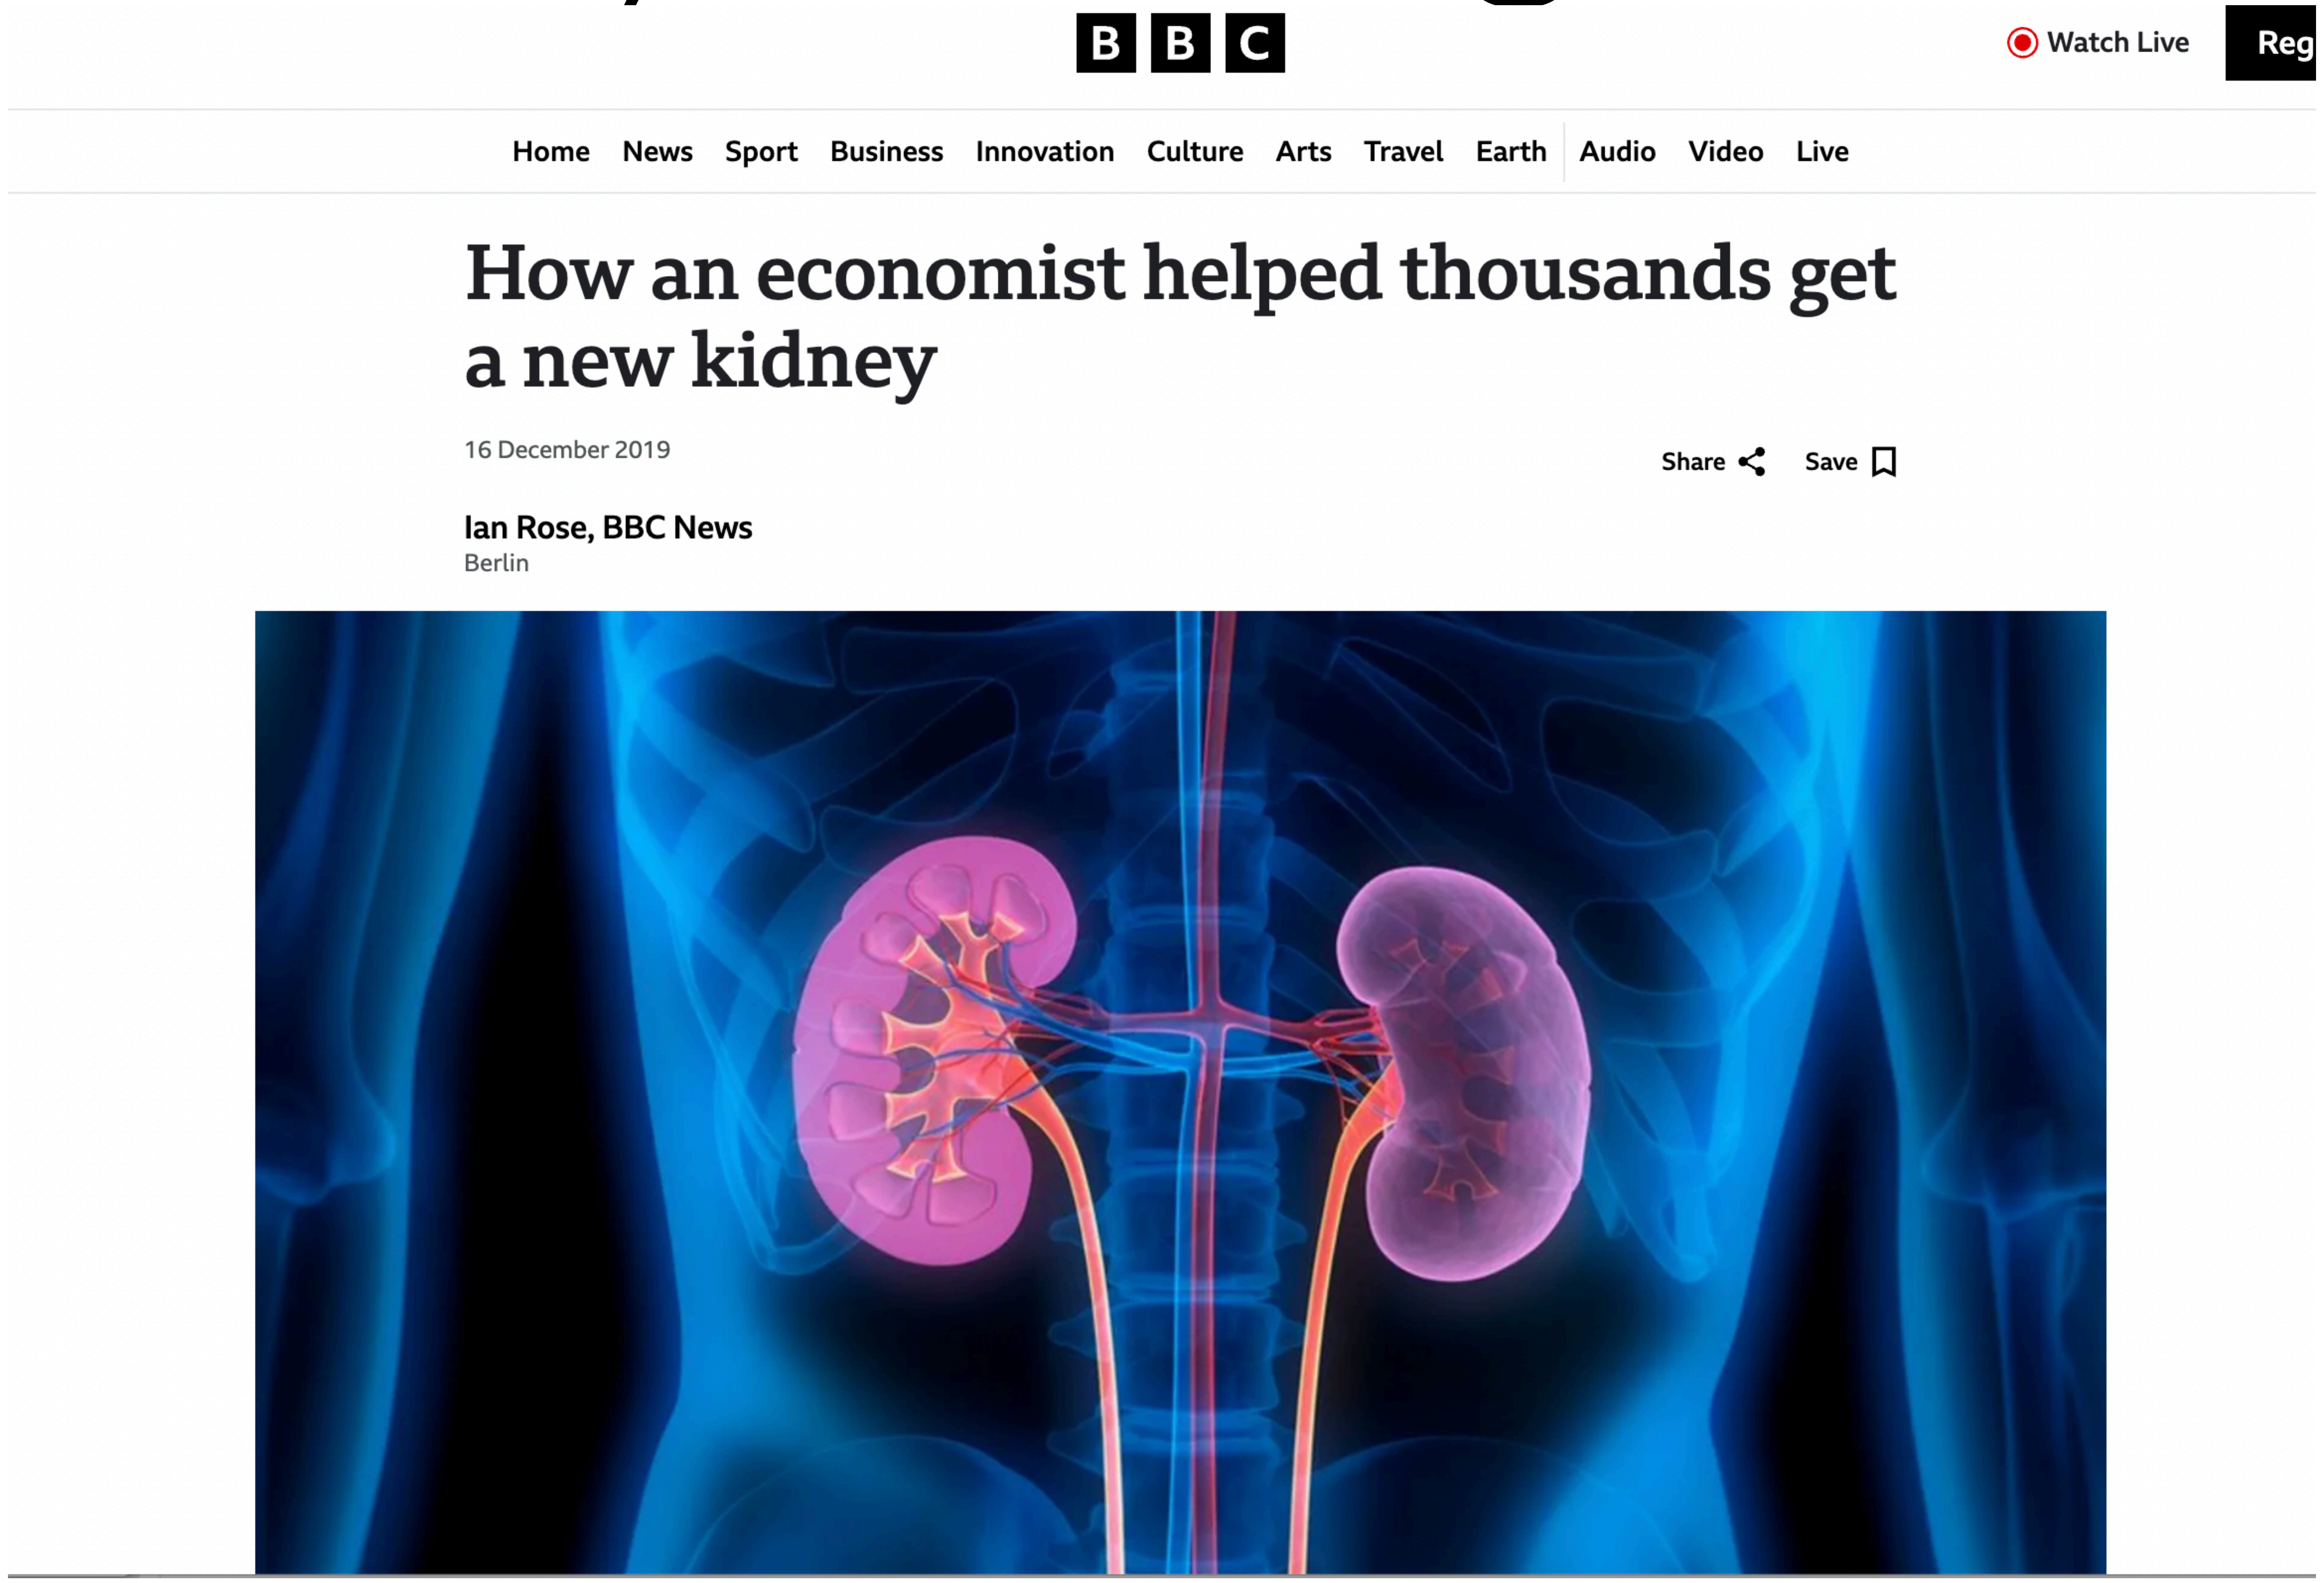
BBC
Watch Live
Reg

Home News Sport Business Innovation Culture Arts Travel Earth Audio Video Live

How an economist helped thousands get a new kidney

16 December 2019

Share Save

Ian Rose, BBC News
Berlin

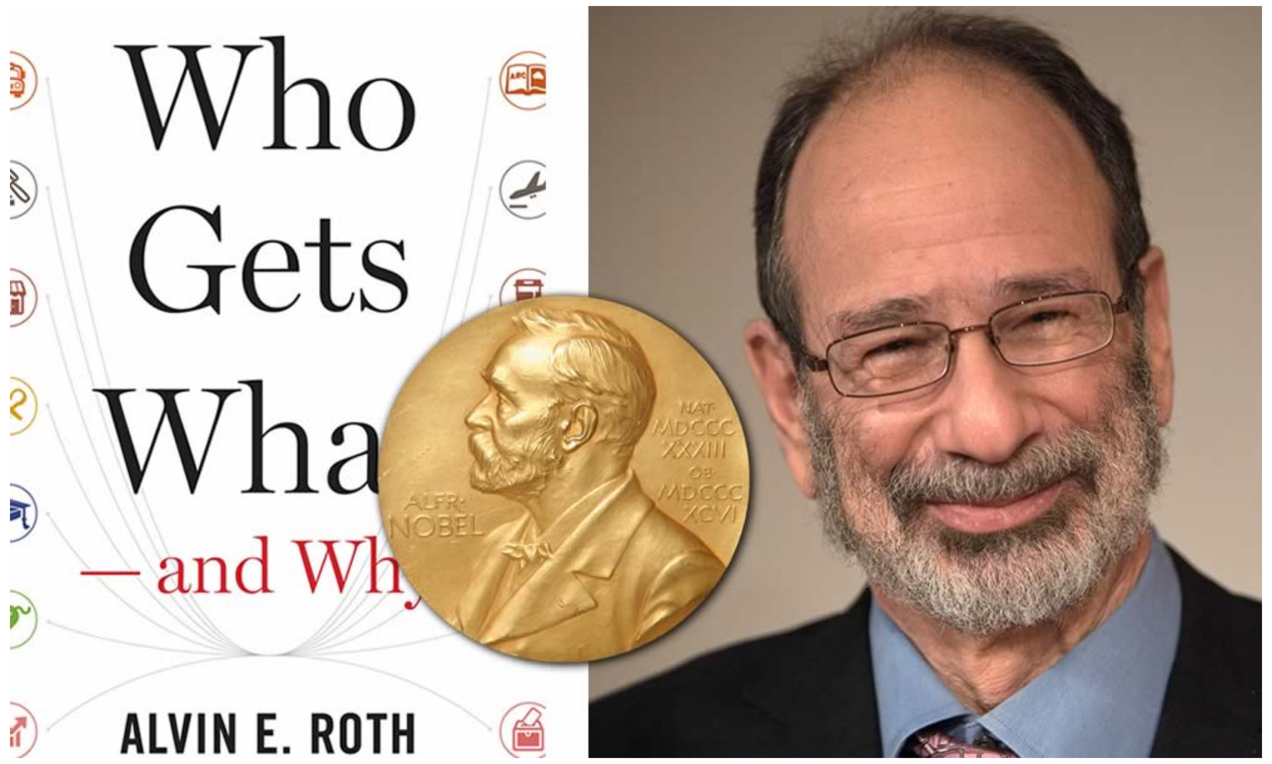

Kidney Exchange

- Many people suffer from kidney failure and need a transplat
- In the US, around **100,000 people** are on a waiting list to receive kidneys each year
- A third of kidney transplants come from living organ donors
- Unfortunately, having a kidney is not enough, sometimes a patient-donor pair is incompatible
- Two incompatible donor-patient pairs might be able to participate in an exchange
- National kidney exchanges have gain momentum
- Kidney exchange is legal but compensation for organ don is illegal in US (and every country except Iran)

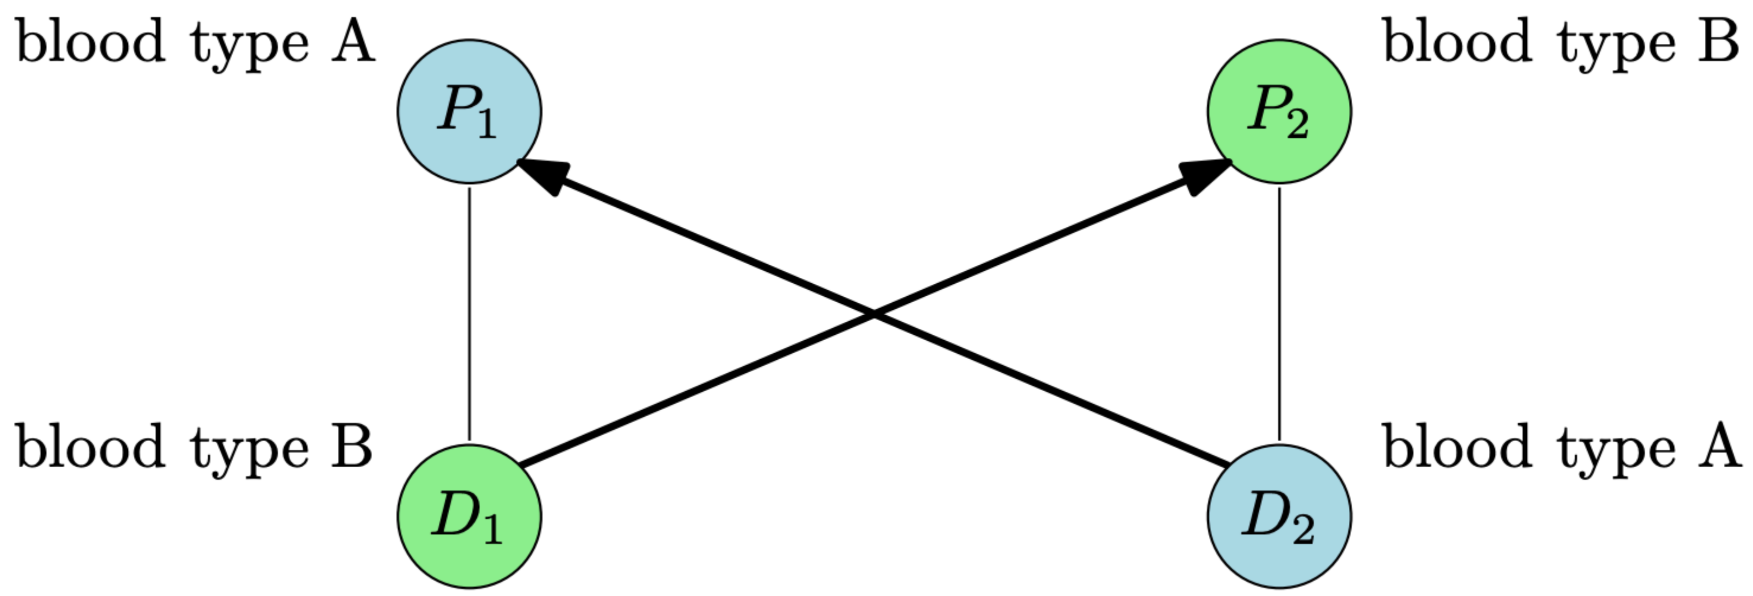

Using TTC: Challenges

- In an influential study in 2004, Roth Sonmez and Unver advocated for the TTC algorithm for kidney exchange
- Agent, house pairs are now patient, donor pairs
- A total ordering over kidneys can be determined by the likelihood of the transplant being successful
- The goal is to reallocate kidneys in way that everyone is collectively as better off as possible
- The actual problem is a bit more complicated and TTC extensions can handle some of them (e.g., accommodating patients without donors, and deceased donors)
- The biggest dealbreaker in TTC for kidney exchange is **long trading cycles**

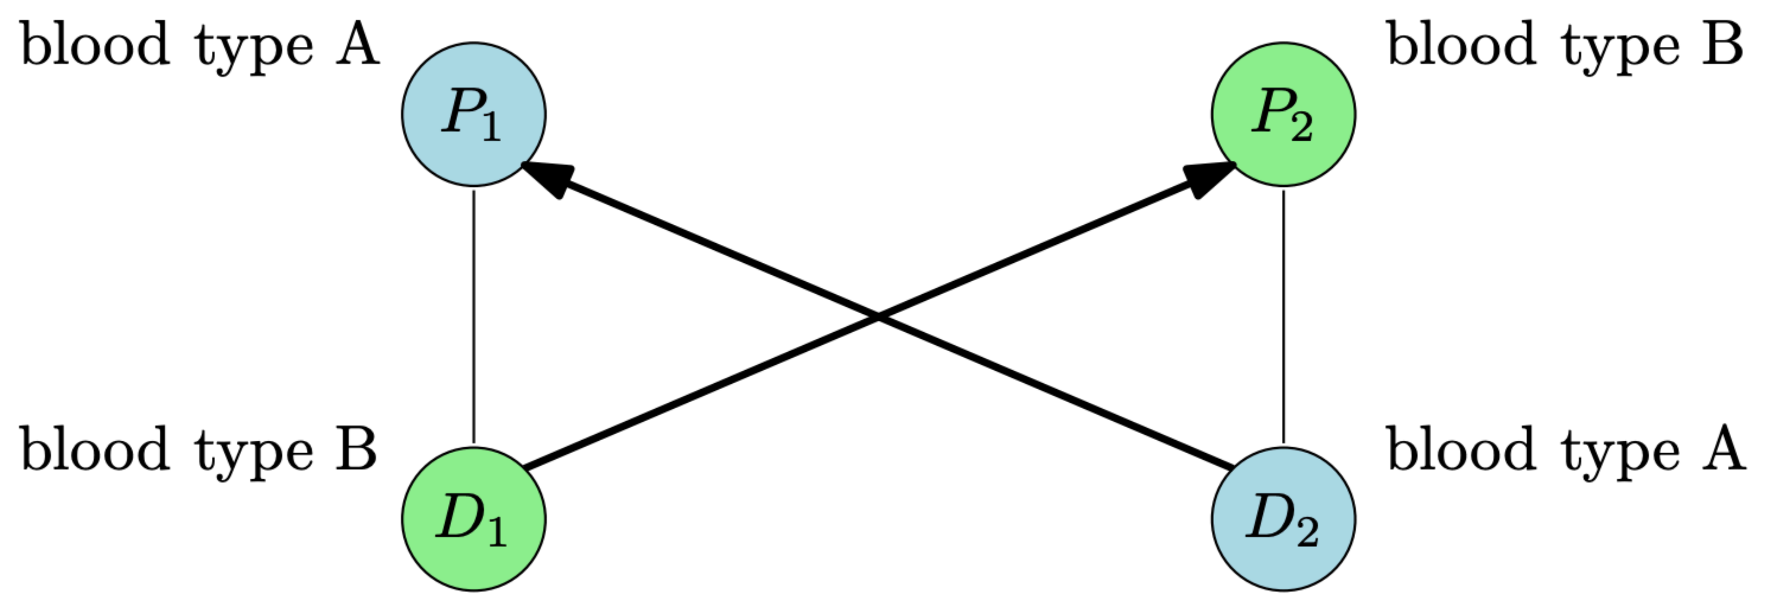

Using TTC: Challenges

- The biggest dealbreaker in TTC for kidney exchange is long trading cycles
    - Transplants must occur simultaneously due to incentive issues (if surgeries for P1 and D2 happen first, there is a risk that D1 will renege on its offer)
- TCC model requires a total ordering over kidneys
    - In reality patients don't care which kidney they get as long as it is compatible with them
    - Binary preferences are more appropriate

Archive: Guinness World Record Organization Distinguishes the National Kidney Registry for World's Longest Kidney Transplant Chain

December 3, 2020

Guinness World Records officially recognized the National Kidney Registry (NKR) for the longest kidney transplant chain in the world. This massive chain was supported by the combined efforts of 25 transplant centers, it included 70 surgeries, facilitated 35 transplants and was featured on ABC News's Nightline.

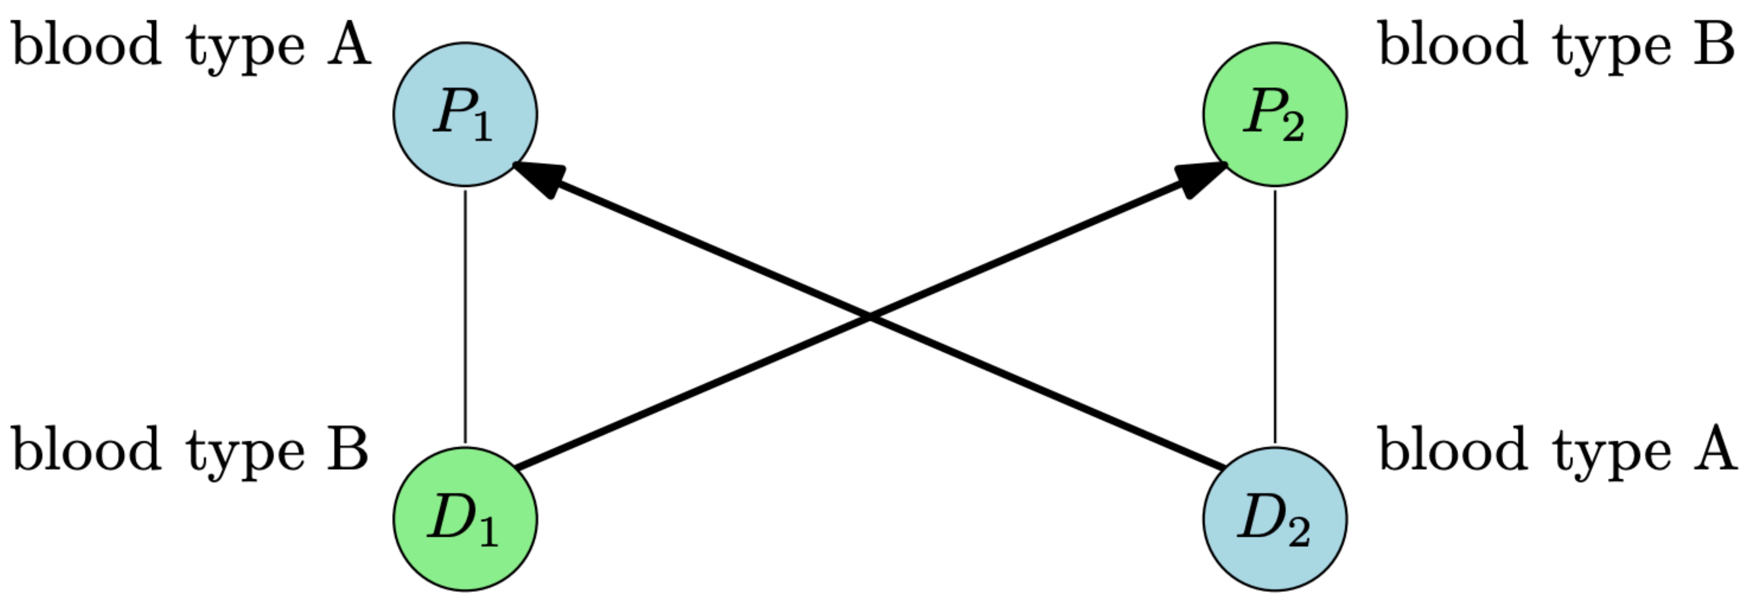

Max Cardinality Matching

- In a subsequent paper, Roth Sonmez and Unver propose using matchings
- The nodes of the graph are patient donor pairs and edges are between compatible pairs that can lead to an exchange
    - Matchings lead to 2-way swaps
- **Model**. Each agent i has a true edge set E_i and can report any subset $F_i \subseteq E_i$ to a mechanism (patients can refuse exchanges $E_i \backslash F_i$ for any reason)
- **Goal**. Compute a maximum-cardinality matching and to be DSIC (for each agent, truthfully reporting its full edge set is a dominant strategy.)

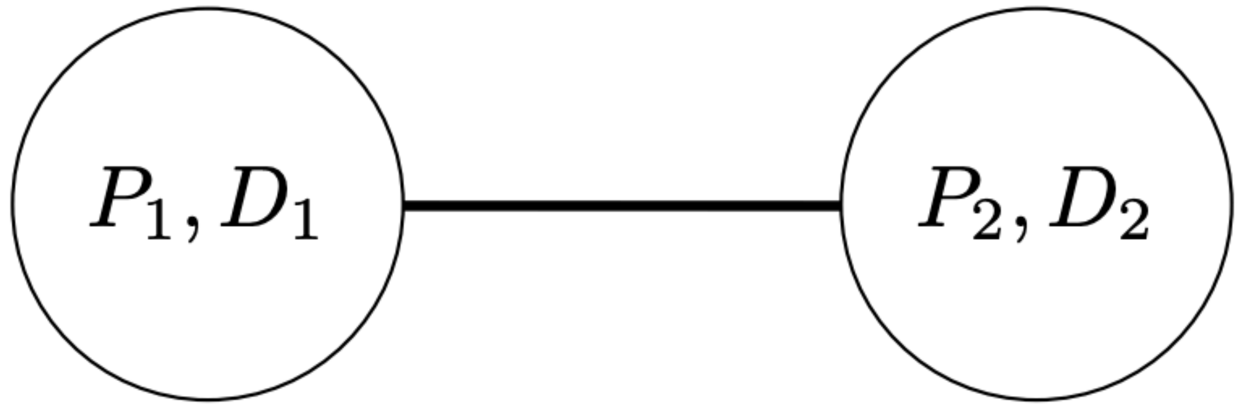

Multiple Matchings

- Even if we collect preferences, create a graph and find a maximum-cardinality matching, there is still a wrinkle
- A graph can have many matchings of the same cardinality
    - How do we handle tie breaks?

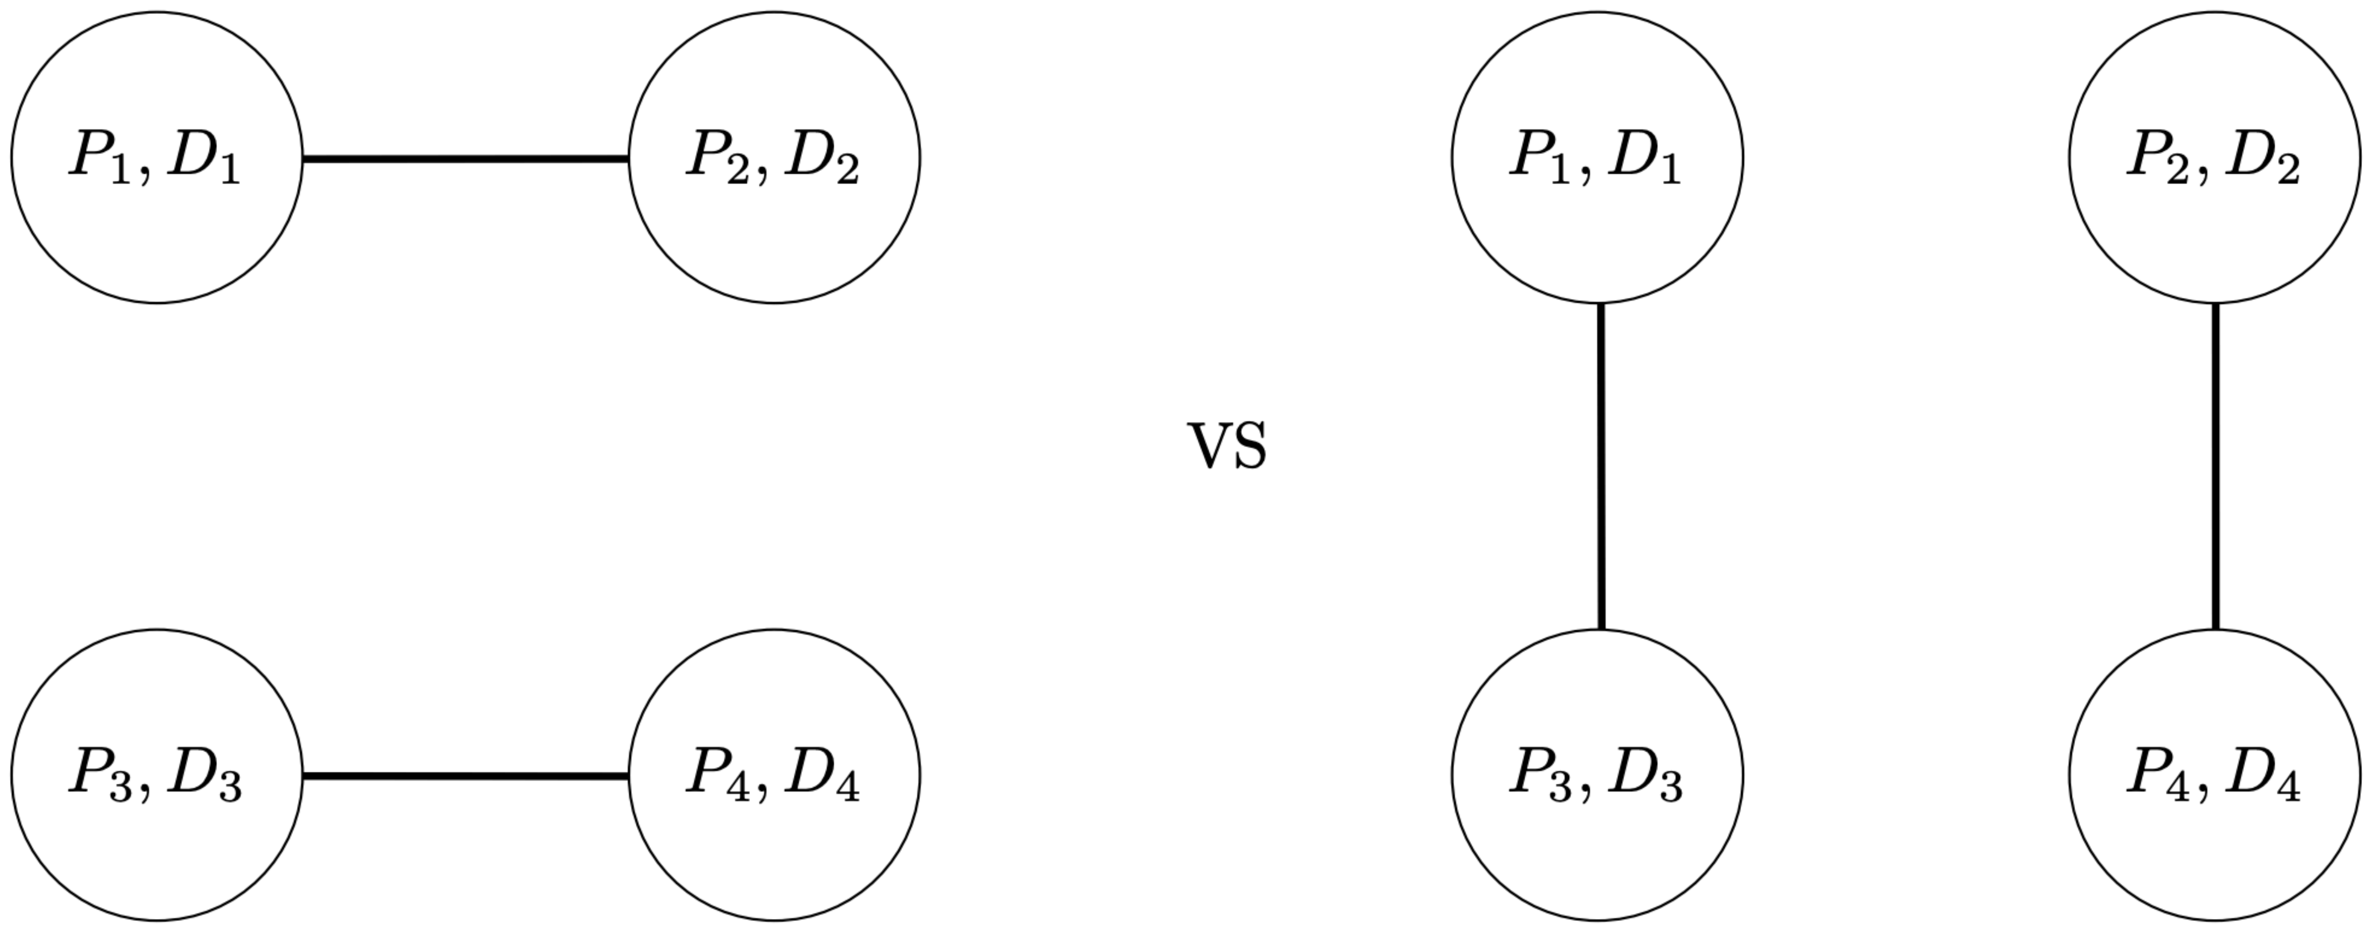

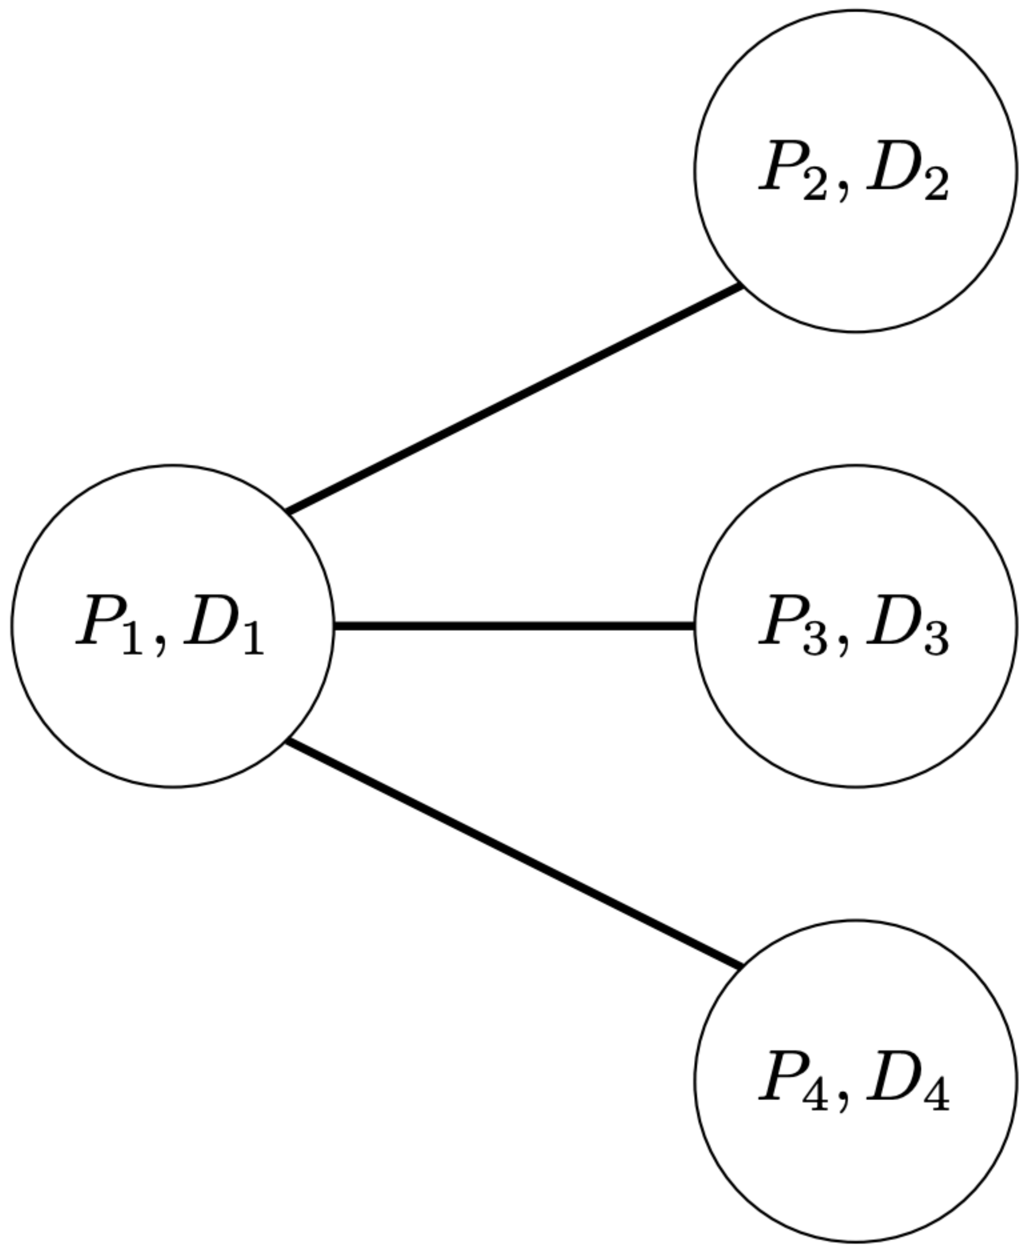

Priority Order Over Nodes

- One way this is resolved through a priority order over nodes
- A priority maximum matching mechanism turns out to be DSIC: no agent can go from unmatched to match by reporting a subset of its edges

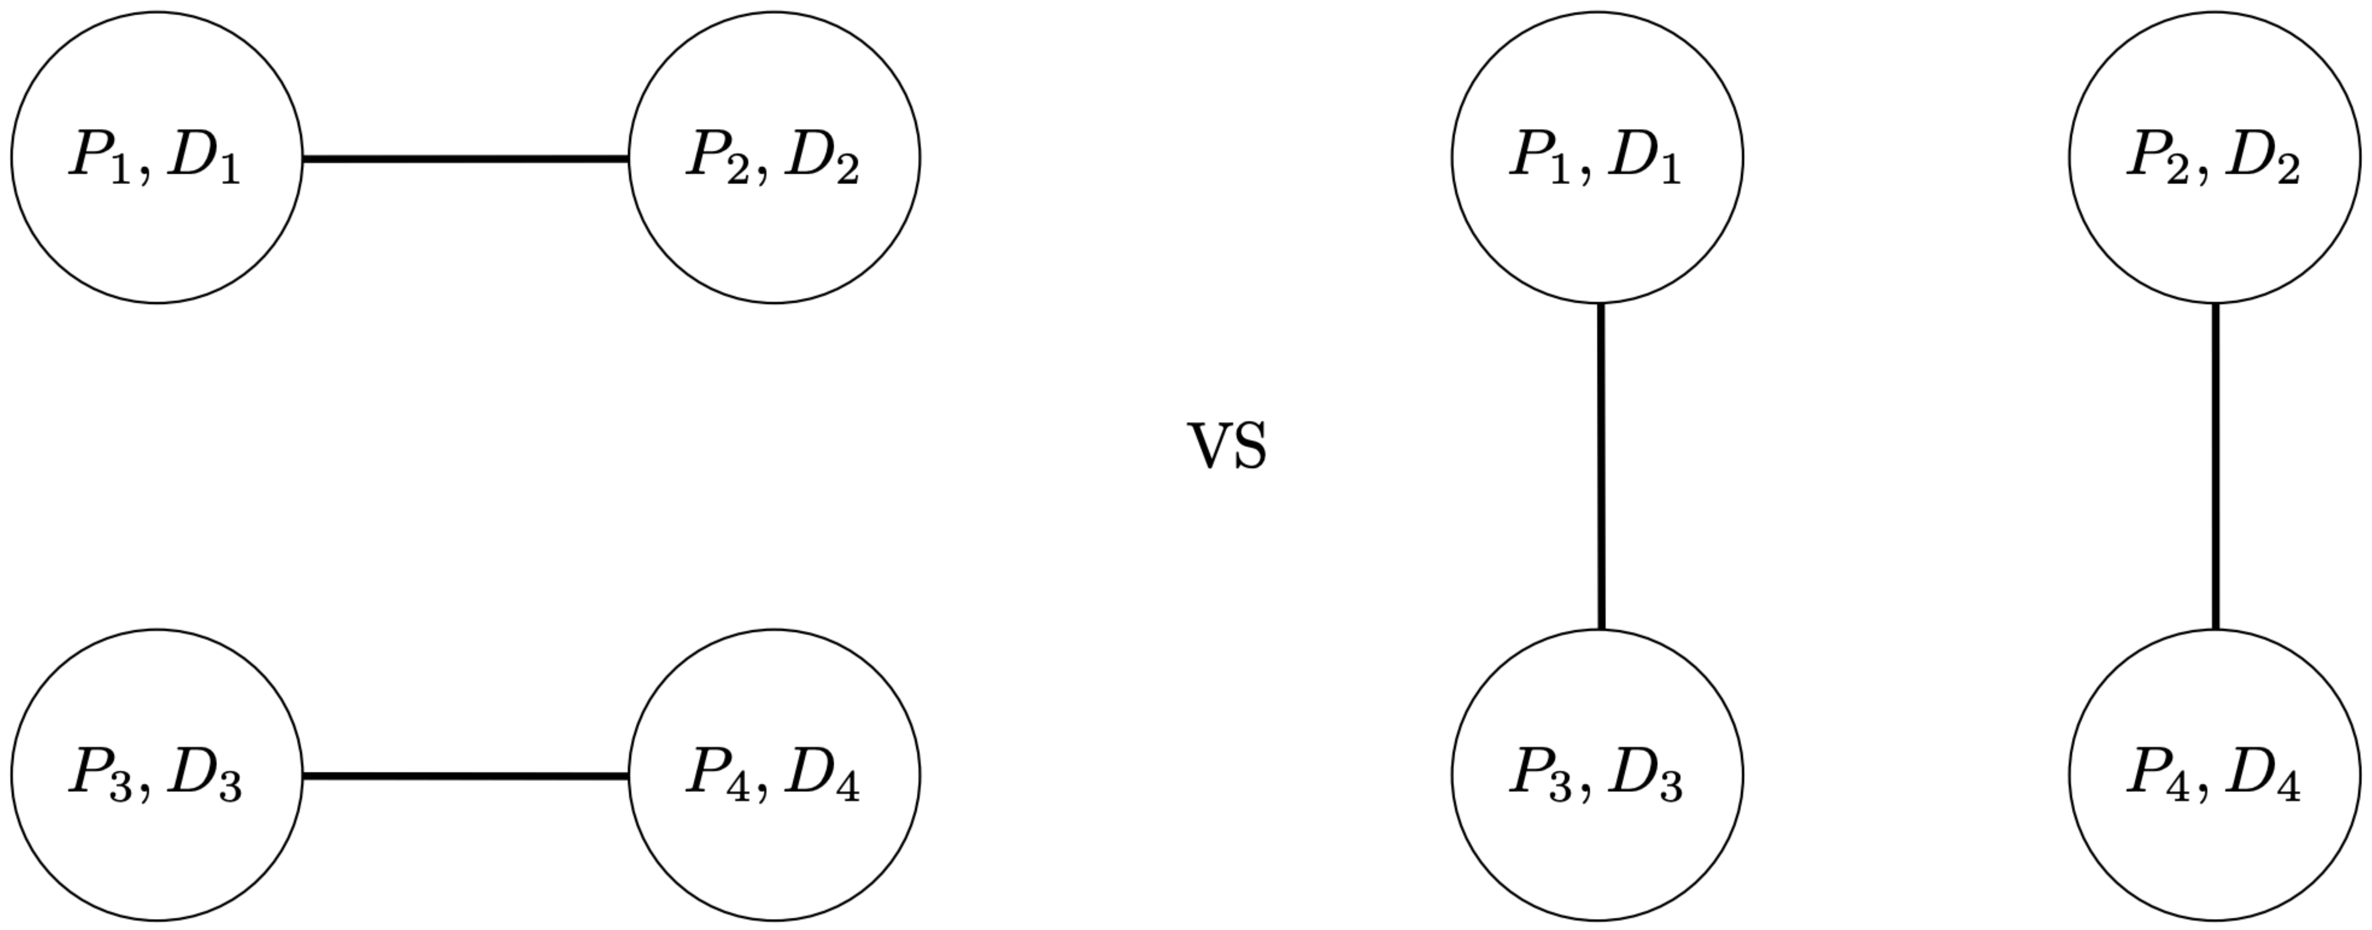

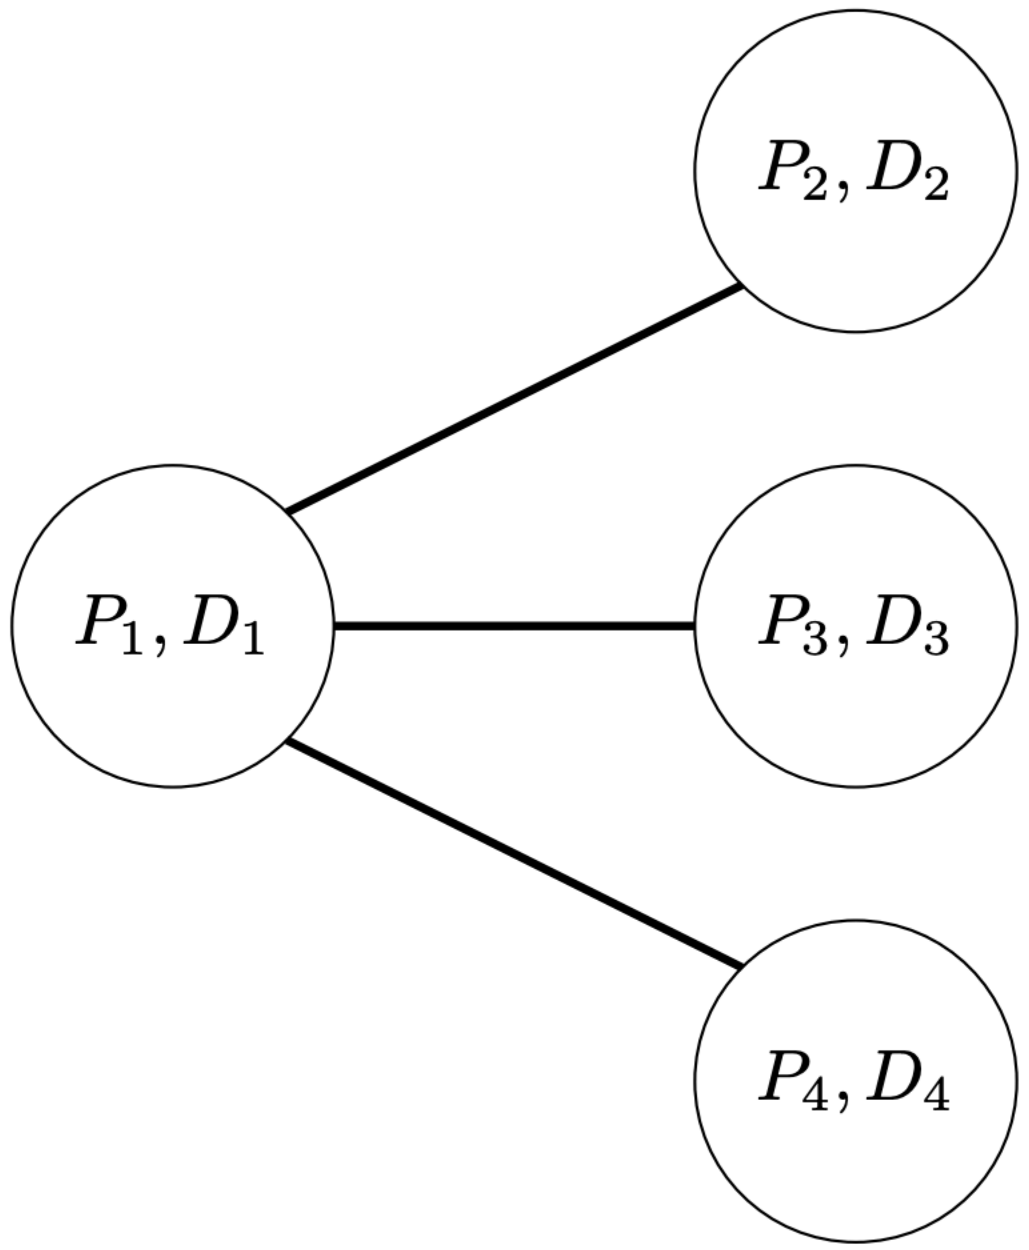

Challenges

- Need for full reporting at the hospital level and need hospitals to participate in global transplants
- Objective of individual hospitals: match as many of their patients as possible, perhaps internally to save time and ensure successful use of organs
- Objective of society: fairness as well as social welfare (match as many patients as possible overall)

Organ Transplant System ‘in Chaos’ as Waiting Lists Are Ignored

The sickest patients are supposed to get priority for lifesaving transplants. But more and more, they are being skipped over.

By Brian M. Rosenthal, Mark Hansen and Jeremy White Feb. 26, 2025

Challenges

- Example left: If H_1 internally matches 1 and 2, and only reports 3, then 3 cannot be matched
  - If both hospitals report all 6 patients, then all 3 exchanges can take place
- Example right: if H_1 hides patients 2 and 3 (while H_2 reports truthfully), what happens?
  - Similarly, if H_2 hides 5 and 6 (while H_1 reports truthfully), what happens?

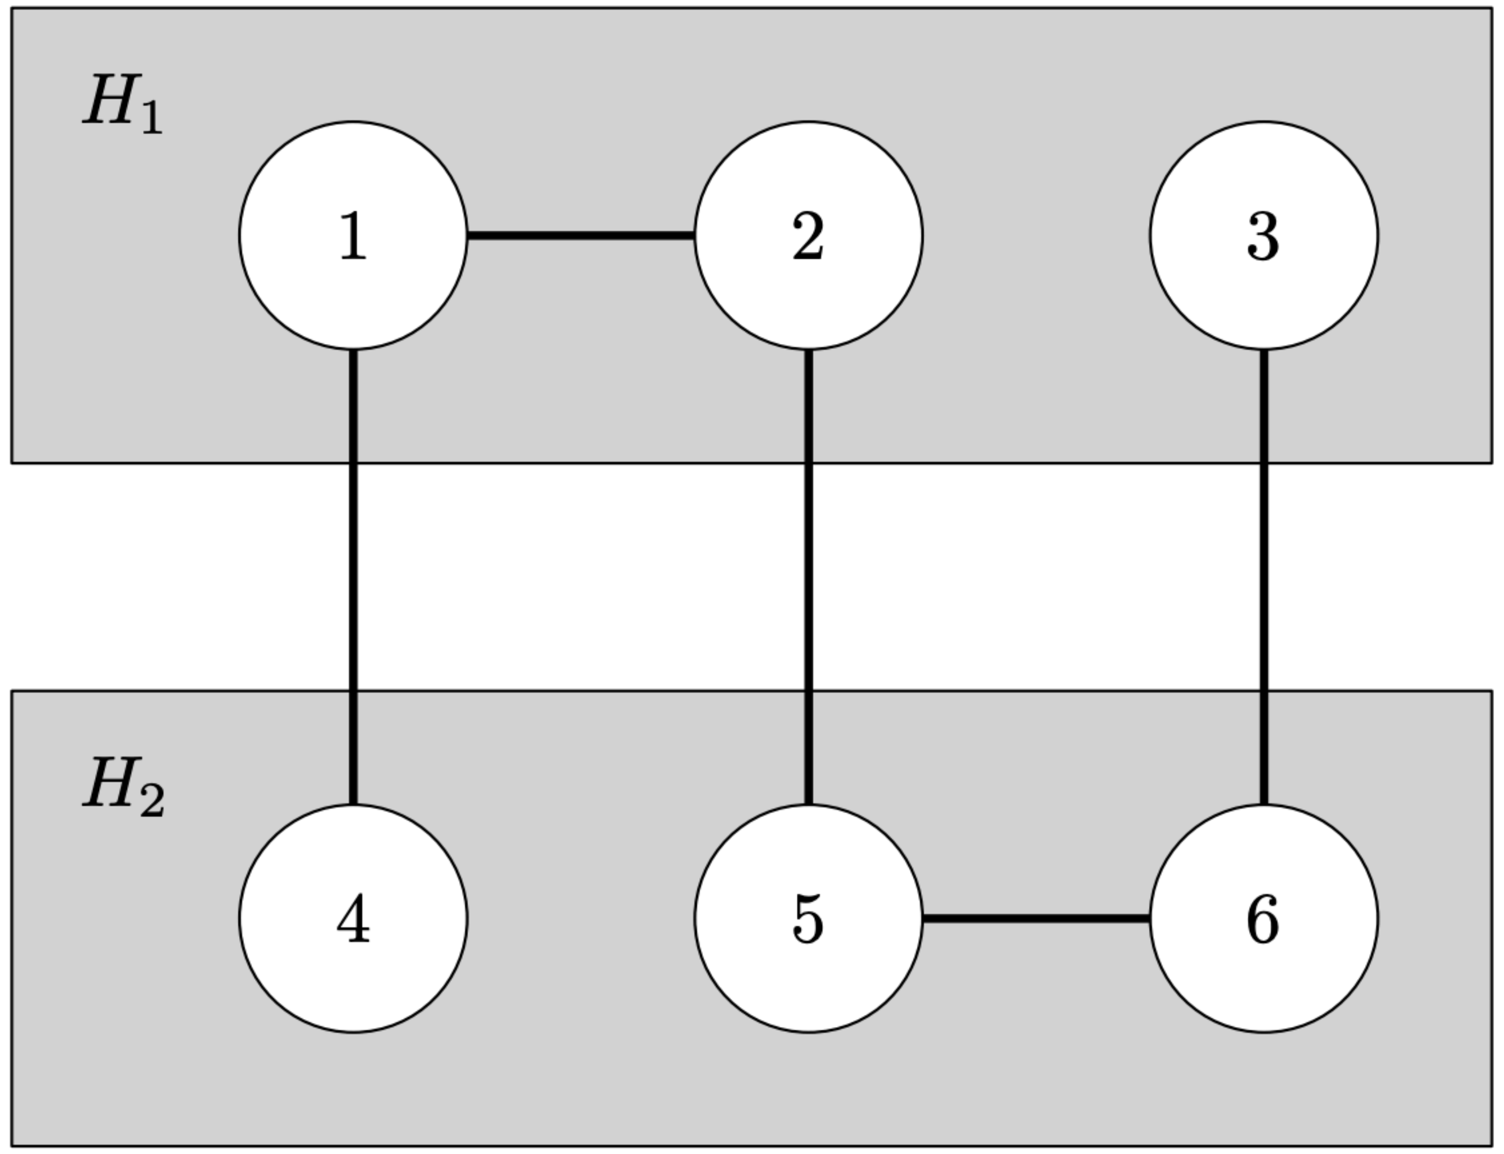

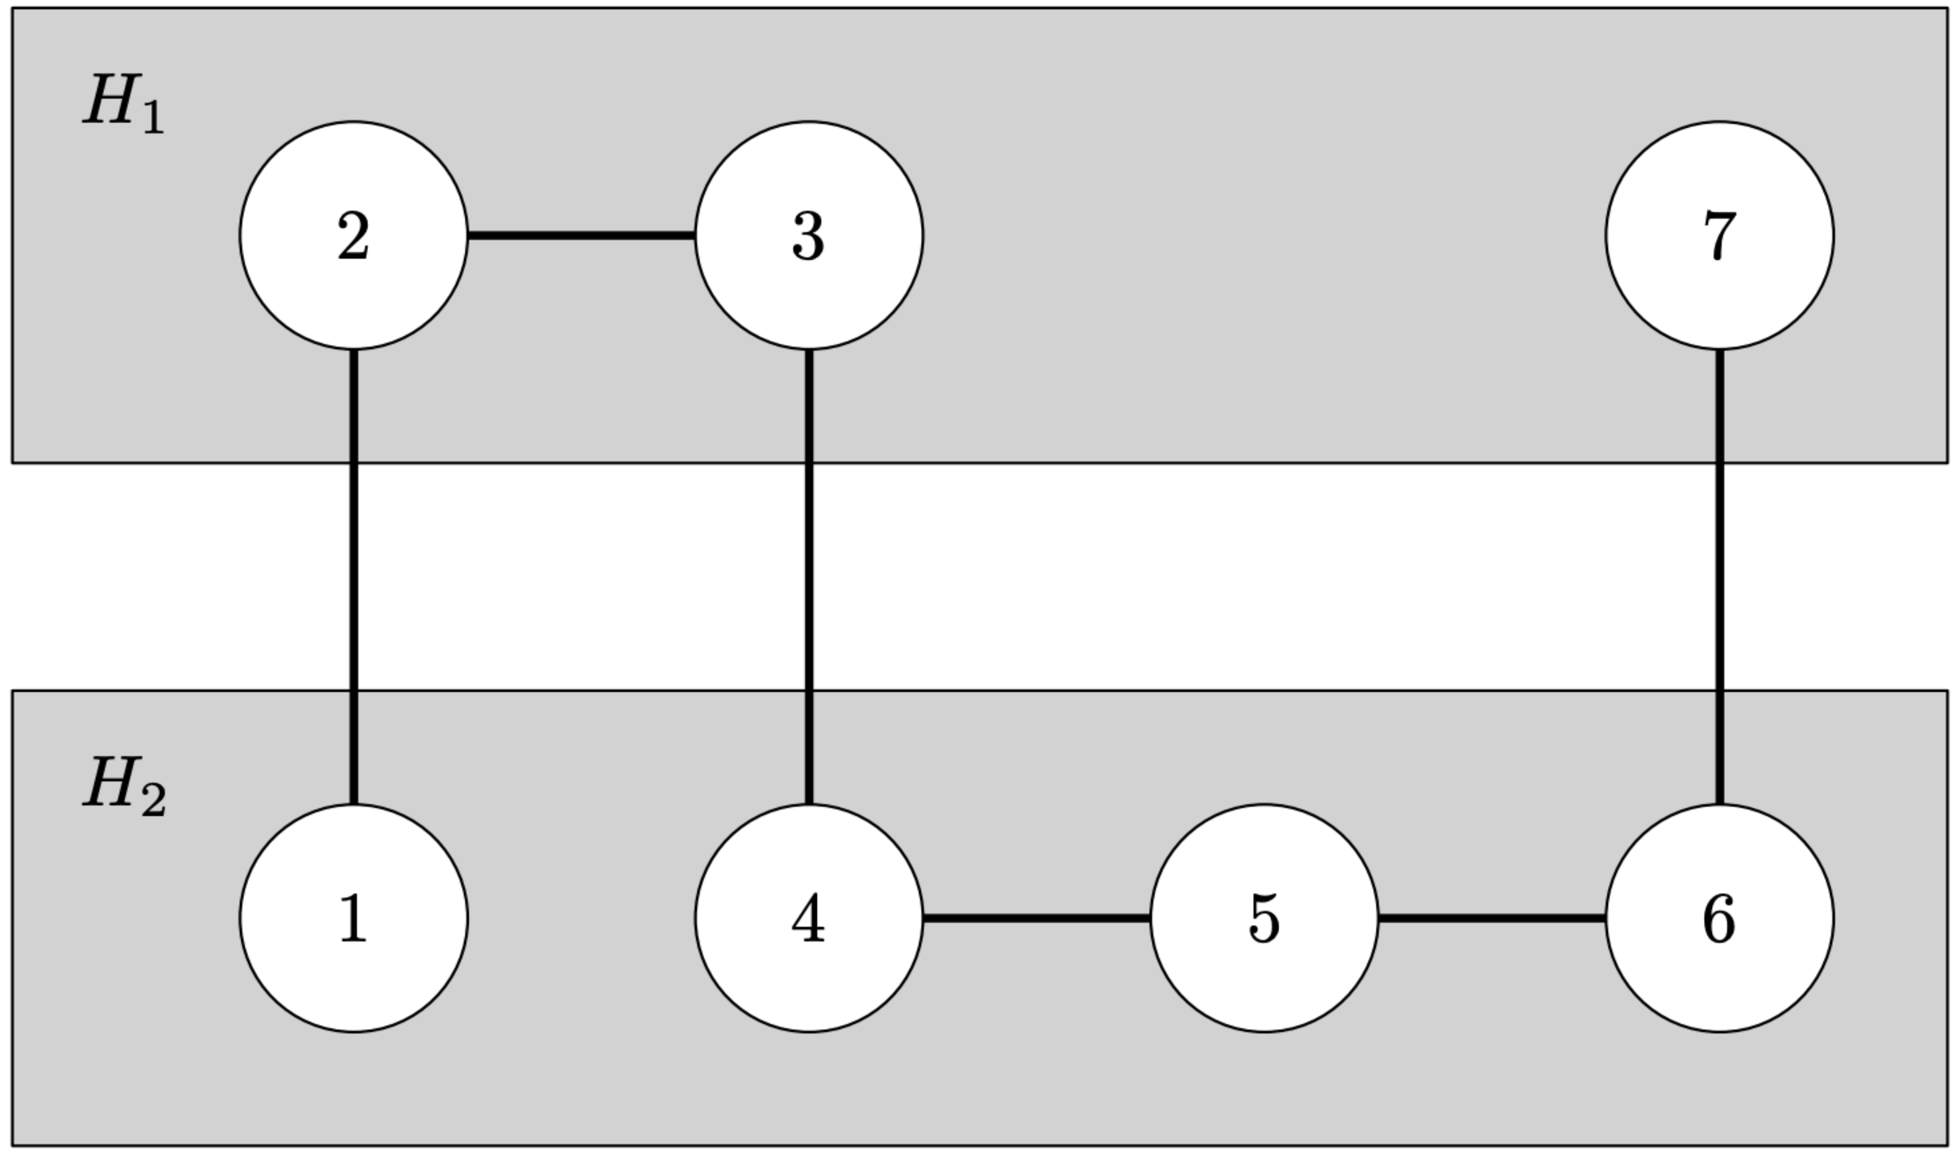

Research Directions

- `Cannot maximize matching size and ensure strategyproofness
- Research direction: how to approximately maximize matching overall that also maximizes the number of matched patients internally for hospitals

Mix and Match: A Strategyproof Mechanism for Multi-Hospital Kidney Exchange☆

Itai Ashlagi[1]

Department of Operations Management, Sloan School of Management, Massachusetts Institute of Technology

Felix Fischer[2]

Statistical Laboratory, University of Cambridge

Ian A. Kash

Microsoft Research Cambridge

Ariel D. Procaccia

Computer Science Department, Carnegie Mellon University

Abstract

As kidney exchange programs are growing, manipulation by hospitals becomes more of an issue. Assuming that hospitals wish to maximize the number of their own patients who receive a kidney, they may have an incentive to withhold some of their incompatible donor-patient pairs and match them internally, thus harming social welfare. We study mechanisms for two-way exchanges that are strategyproof, i.e., make it a dominant strategy for hospitals to report all their incompatible pairs. We establish lower bounds

A Random Graph Model of Kidney Exchanges: Efficiency, Individual-Rationality and Incentives

Panos Toulis
Harvard University, SEAS
Cambridge, MA 02138, USA
ptoulis@fas.harvard.edu

David C. Parkes
Harvard University, SEAS
Cambridge, MA 02138, USA
parkes@eecs.harvard.edu

ABSTRACT

In kidney exchanges, hospitals share patient lists and receive transplantations. A kidney-paired donation (KPD) mechanism needs to promote full sharing of information about donor-patient pairs, and identify a Pareto efficient outcome that also satisfies participation constraints of hospitals. We introduce a random graph model of the KPD exchange and then fully characterize the structure of the efficient outcome and the expected number of transplantations that can be performed. Random graph theory allows early experimental results to be explained analytically, and enables the study of participation incentives in a methodological way. We derive a square-root law between the welfare gains from sharing patient-donor pairs in a central pool and the individual sizes of hospitals, illustrating the urgent need for the nationwide expansion of such programs. Finally, we establish through theoretical and computational analysis that enforcing simple individual rationality constraints on the outcome can mitigate the negative impact of strategic behavior by hospitals.

1. INTRODUCTION

The scarcity of cadaver kidneys and significant medical benefits from live kidney donation has promoted the expansion of kidney-paired donation (KPD) in recent years. The idea is that kidney patients with one or more incompatible donors, might be able to receive compatible transplants through barter exchanges. It is typical for this to be performed as a two-way exchange, which involves four (typically simultaneous) operations.[1] Currently, there is a handful of such kidney-exchange programs in the USA and several others around the world [14]. Their expansion in large-scale has been hitherto hindered by the ethical, logistical and even incentive issues they entail.

Nevertheless, there are numerous reports favoring the benefit of kidney exchanges. From the medical literature, survival rates up to 100% are reported in a sample including 10 two-way, paired donations [8]. Further benefits in terms of total saved lives can be found where larger than two-way exchanges are considered [14]. Recent history has seen an increase in (KPD) through multi-regional KPD programs [11].

Naturally enough, if we would like to represent the patient-

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: committee selection, budgeting decisions
  - Fair division: how to divide indivisible goods fairly (cake-cutting problems)
- No money or transfers involved
  - Mechanism design without money

Voting 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 L_i alternatives A
- Voting rules can have two forms:
    - **Social-choice function** selects a single alternative for a given preferences profile, that is, $L_1, L_2, \ldots, L_n \mapsto a^*$ where $a^* \in A$
    - **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

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

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.

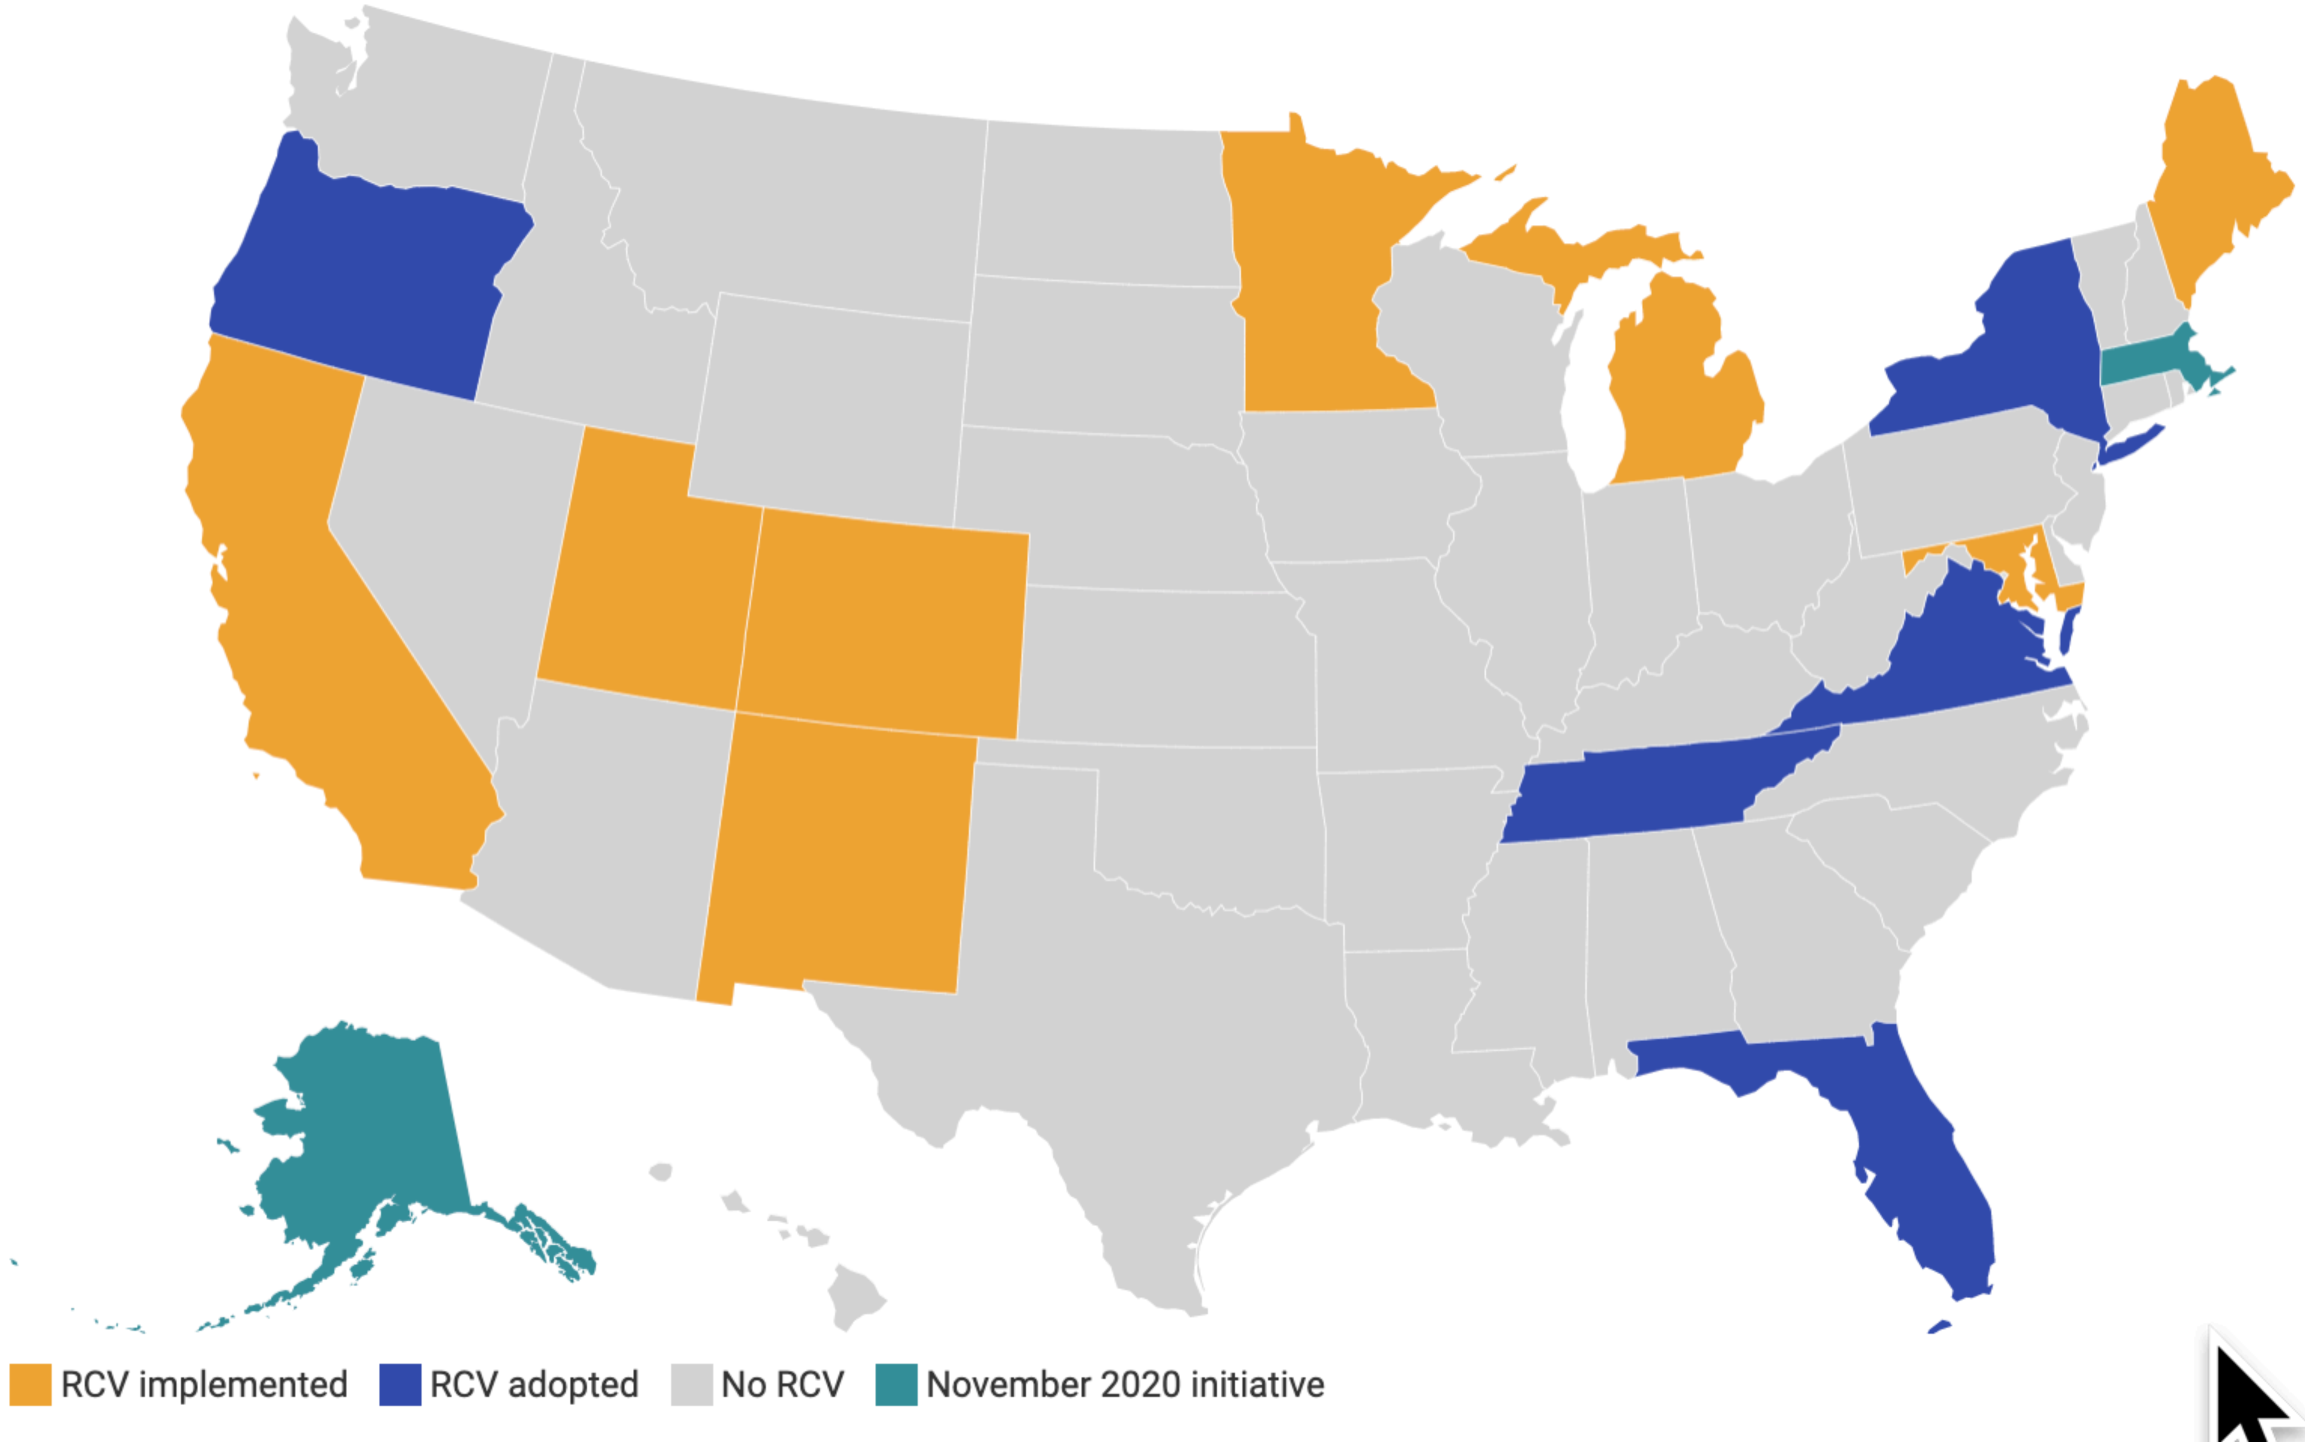

Source: https://www.southcoasttoday.com/news/20201003/in-massachusetts-yes-on-2-pitches-ranked-choice-voting

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
- Ranked-choice voting is preferred as it is less susceptible to extreme candidates
- How good is this voting rule?

	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

- 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

- **Question 1.** Is ranked-choice voting rule "fair"?
  - Let's wait on a criterion for this
- **Question 2.** Is ranked-choice voting rule strategyproof?
  - Can you see a useful misreport in this 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

Strategy Proof

- 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.

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
- **Question.** Do ranked-choice voting and plurality satisfy Condorcet criterion?

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

- **Question.** Is Borda count strategyproof?

Winner
b

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

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

Borda Count

- **Question.** Does Borda count satisfy Condorcet criterion?
  - Question in next homework

Many Rules, Many Applications

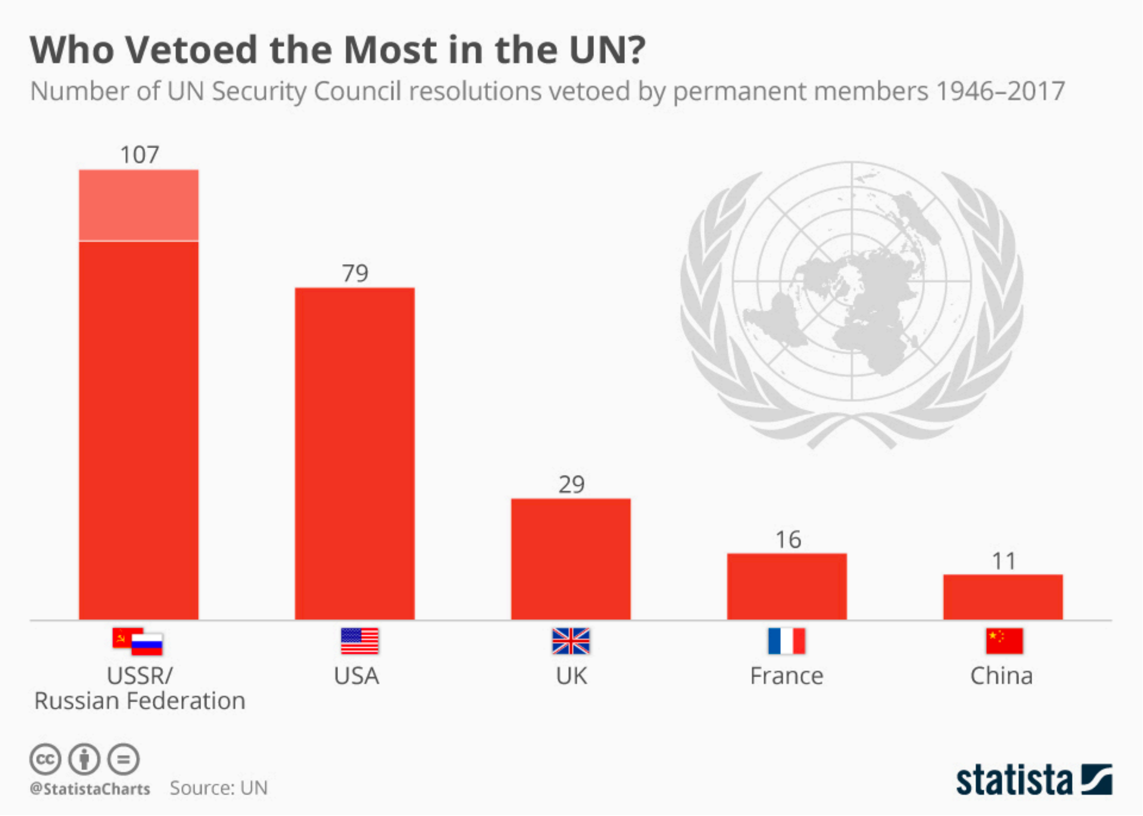

https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.pdf

One to Rule them All?

- For the same input profile, plurality, Borda and ranked-choice can all output a different winner!
    - Can you construct such an example?
- Changing the voting rule changes the outcome of the mechanism
- Leads to contention on which voting rule is the “best”
- Voting theorists have an "axiomatic" approach to study voting rules
- Identify "desirable" properties that one would like
- Compare rules based on that
- **Question**: Is there any voting rule that is strategyproof and reasonable?

Properties of Voting Rules

- **Onto**: For any candidate a, there exists an input profile where a wins

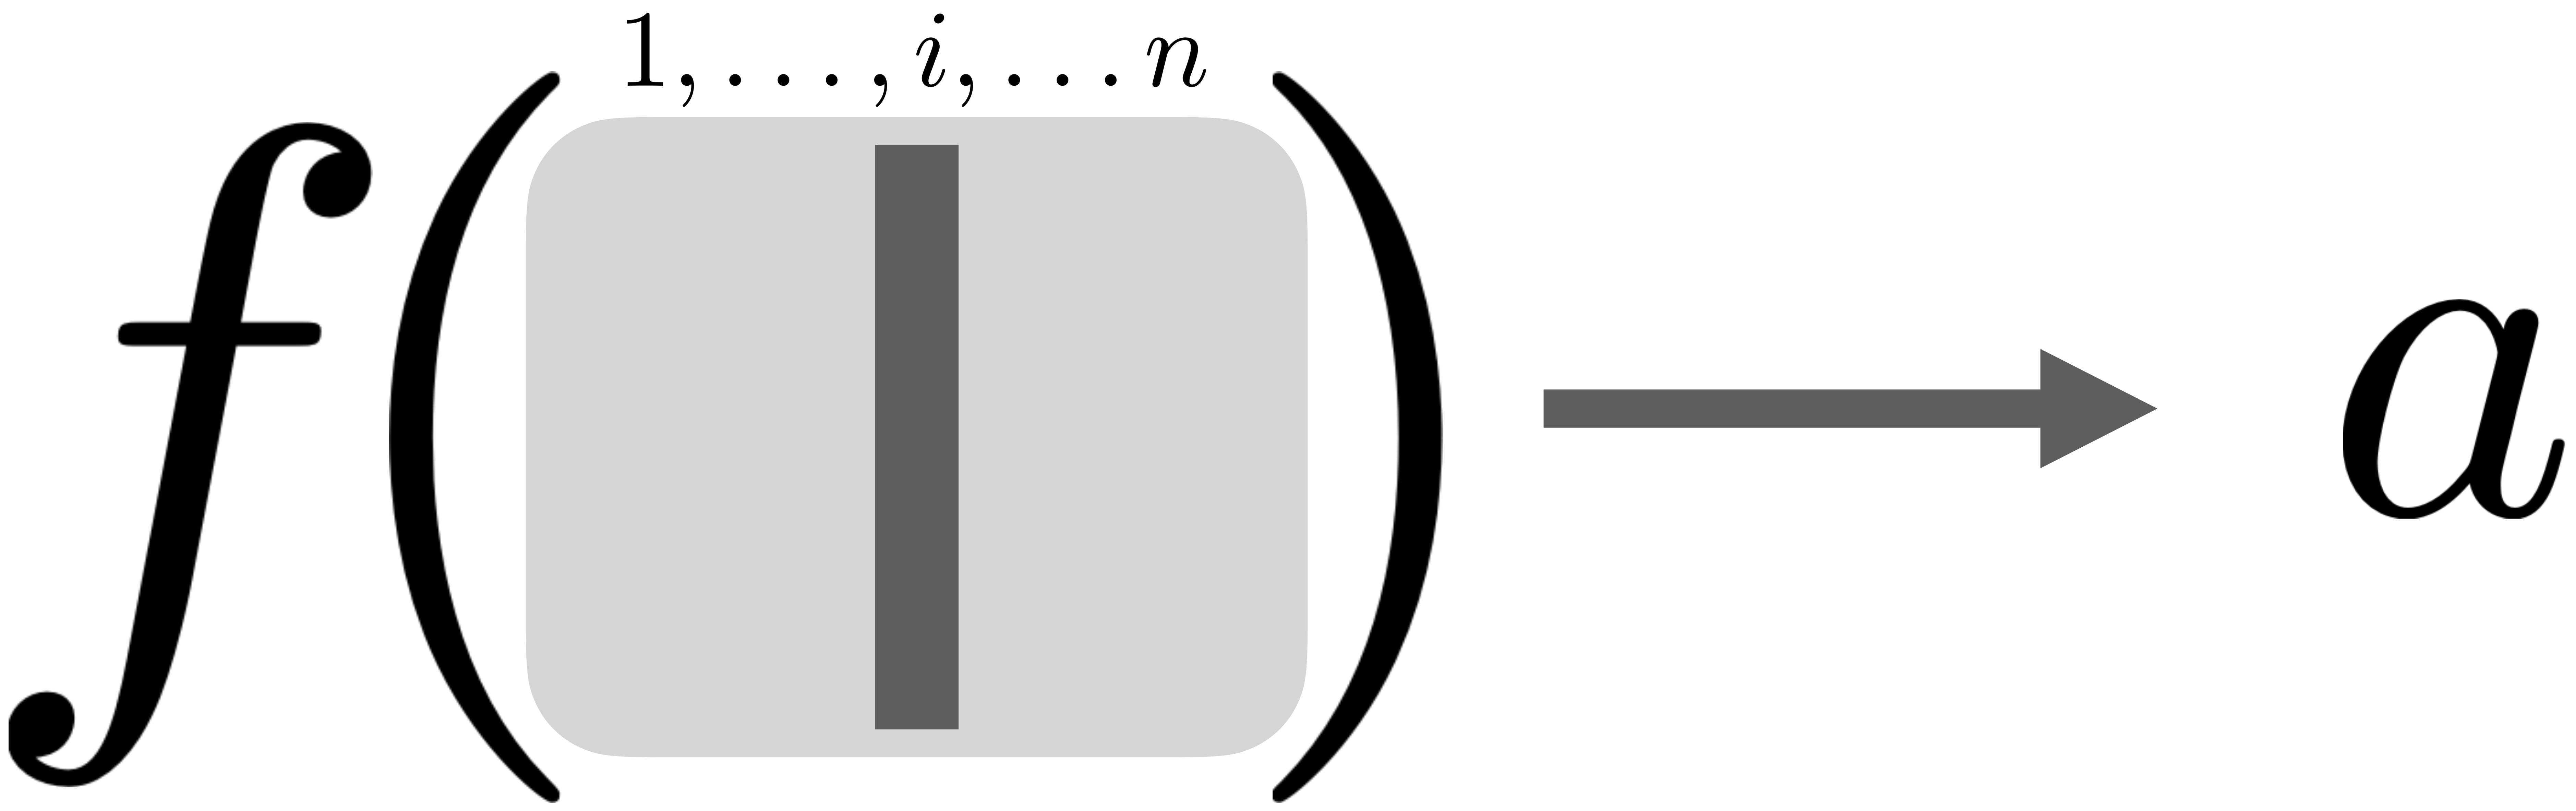

- Are Borda, plurality, ranked-choice etc onto?
  - Yes, can always construct a profile to make any candidate win

Properties of Voting Rules

- **Strategyproof**: No voter can improve by misreporting preferences

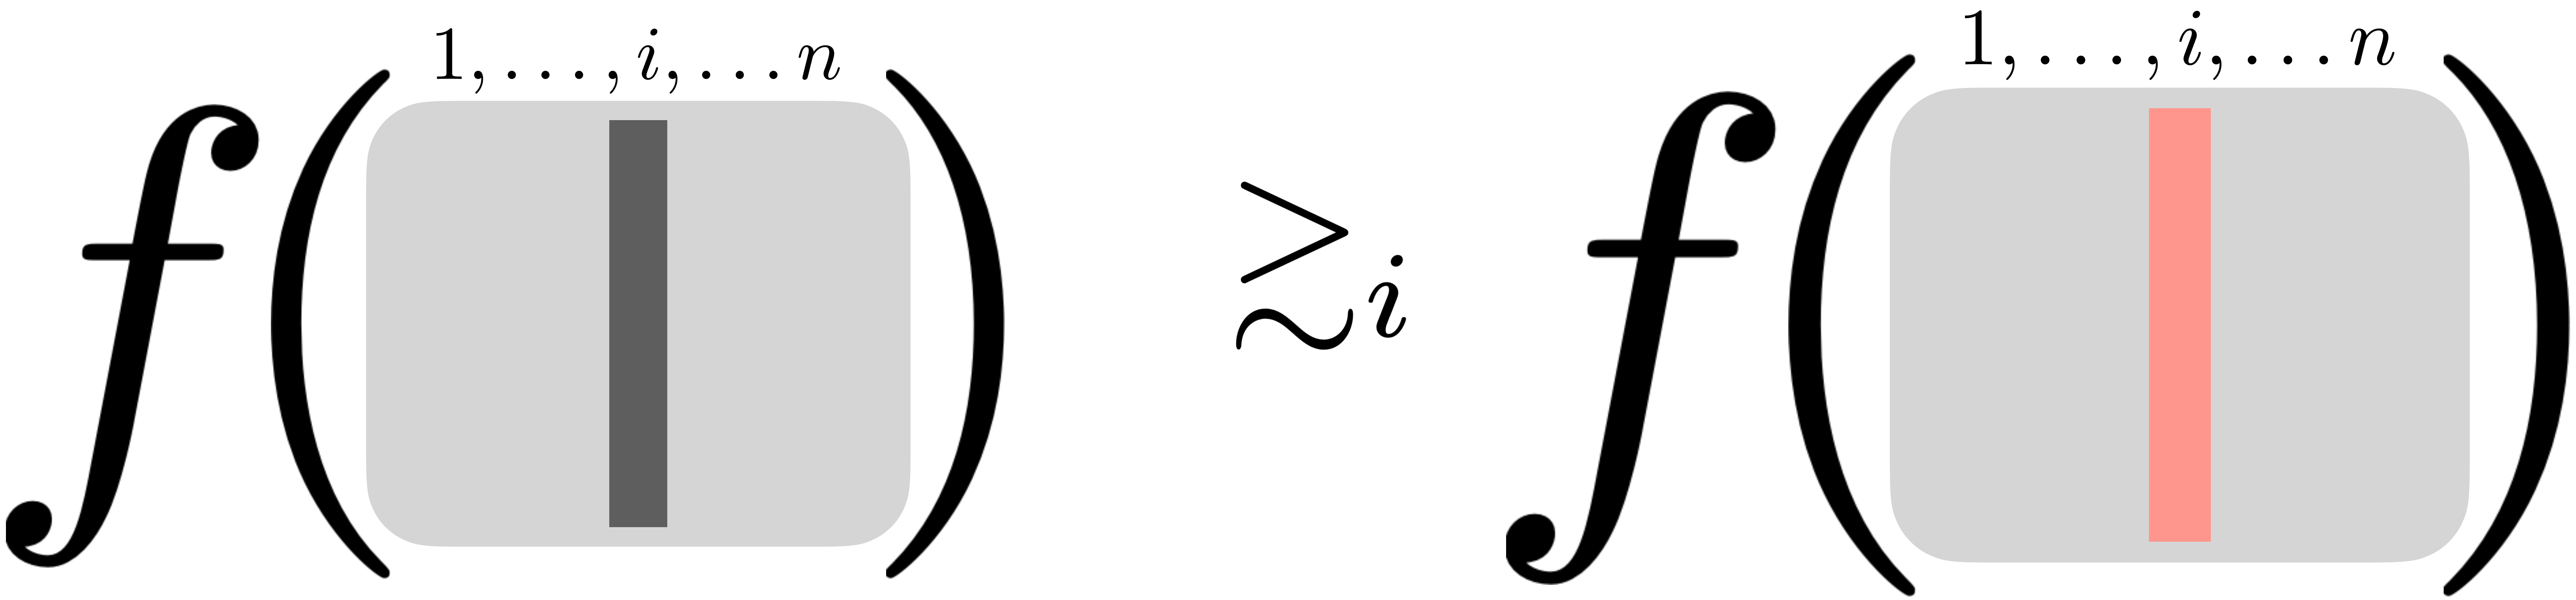

- Are Borda, plurality, ranked-choice etc strategyproof?
  - No

Onto and Strategyproof

- (3 or more alternatives) onto but not strategyproof? Borda, Plurality, Ranked-choice
- (3 or more alternatives) strategyproof but not onto? Constant or restricted majority

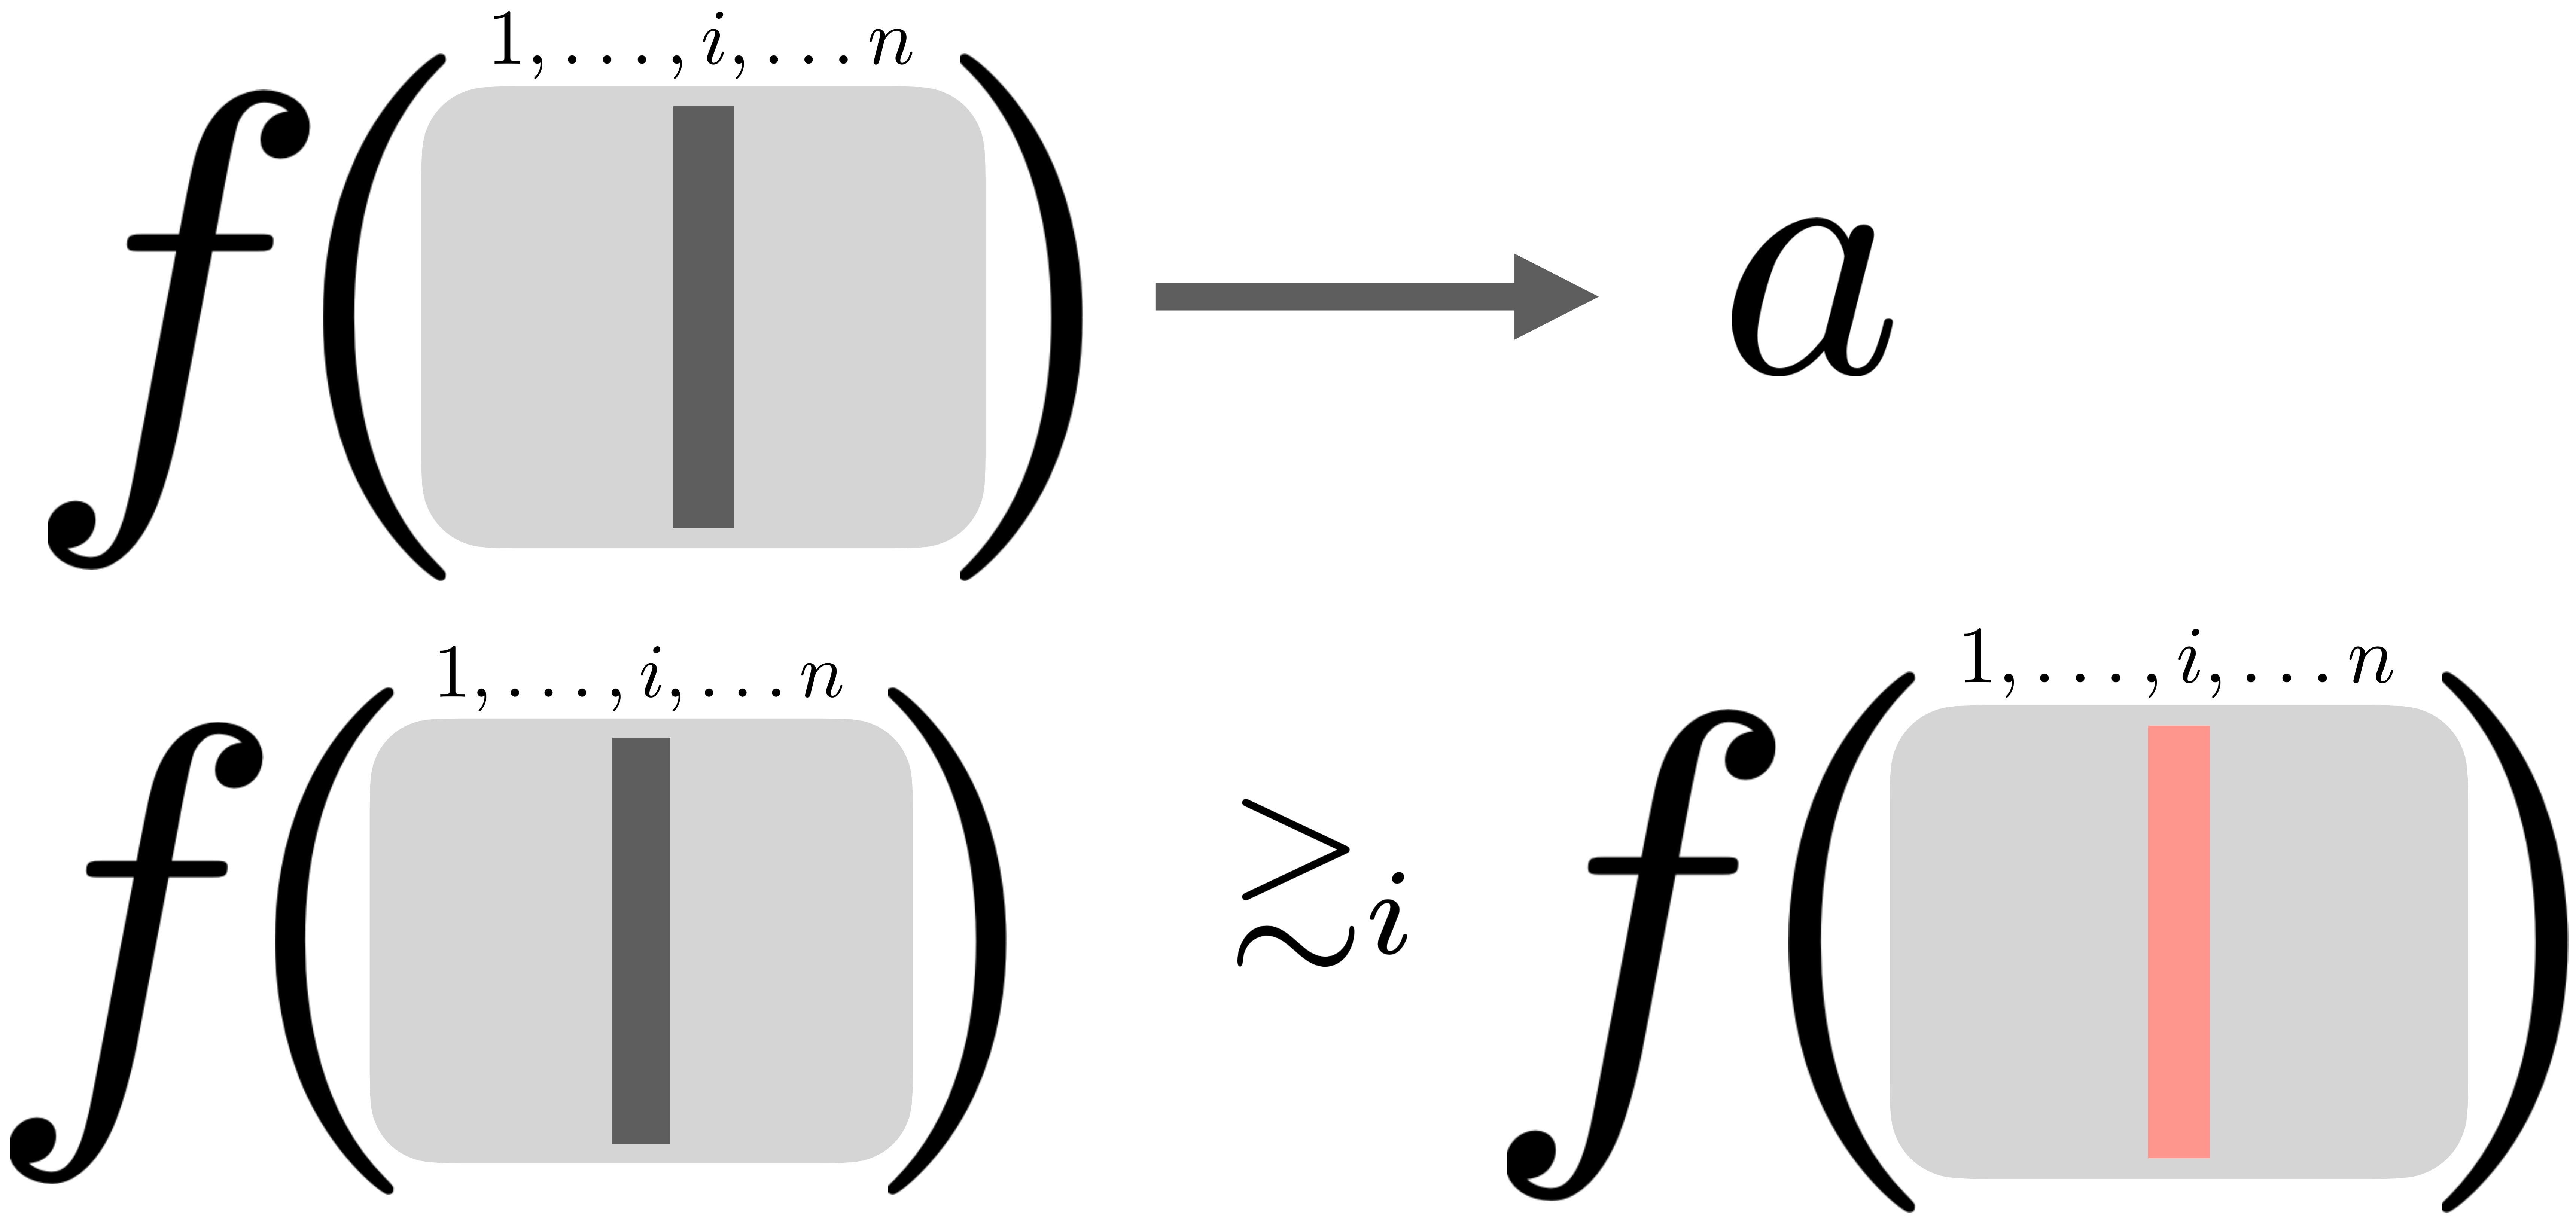

A Bad Voting Rule

- **Dictatorship** : A voting rule is **dictatorial** if there is a voter i such that the rule always elects i's first choice (regardless of others' preferences)

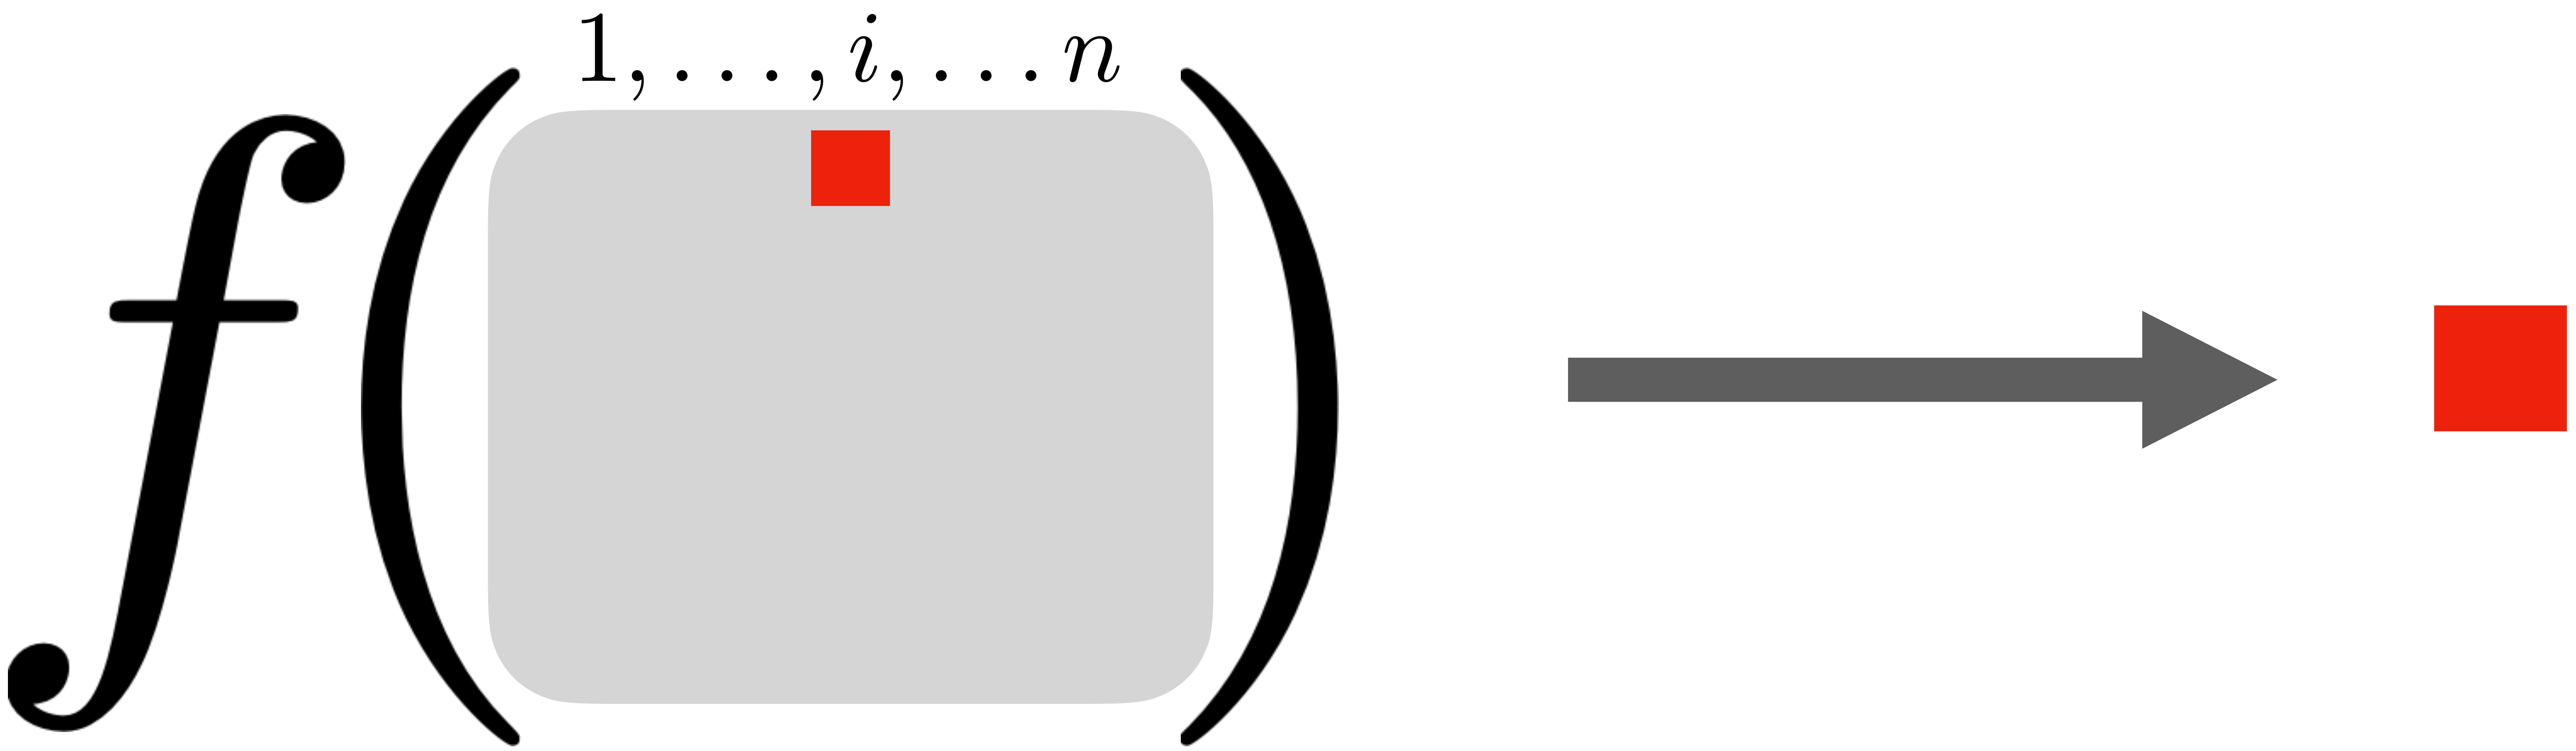

- Is a dictatorship straregyproof?
- Is a dictatorship onto?

[Gibbard '73, Satterthwaite '75]

When there are 3 or more alternatives, a voting rule is strategyproof and onto if and only if it is dictatorial.

Monotonicity

- **Definition**. Suppose a is the current winner (on profile L). For all input profiles L', in which for all voters, any candidate who was ranked below a in L is still ranked below a in L', then a should continue to win in L'.
    - Support of a either increases or stays the same: a's outcome cannot get worse
- **Theorem**. Strategyproof $\Longleftrightarrow$ monotone

Image credit: https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.p

[GS Theorem] With three or more candidates, a voting rule is **strategyproof** and **onto** if and only if it is a **dictatorship**.

Goal. Strategyproof + Onto $\Longrightarrow$ Dictatorship

[Proof Outline]

Part 1. Strategyproof $\Longleftrightarrow$ Monotonicity

Part 2. Monotone + Onto $\Longrightarrow$ Unanimous

Part 3. Monotone + Unanimous $\Longrightarrow$ Dictatorship

Strategyproof $\implies$ Monotone

- Suppose a rule is strategyproof but not monotone

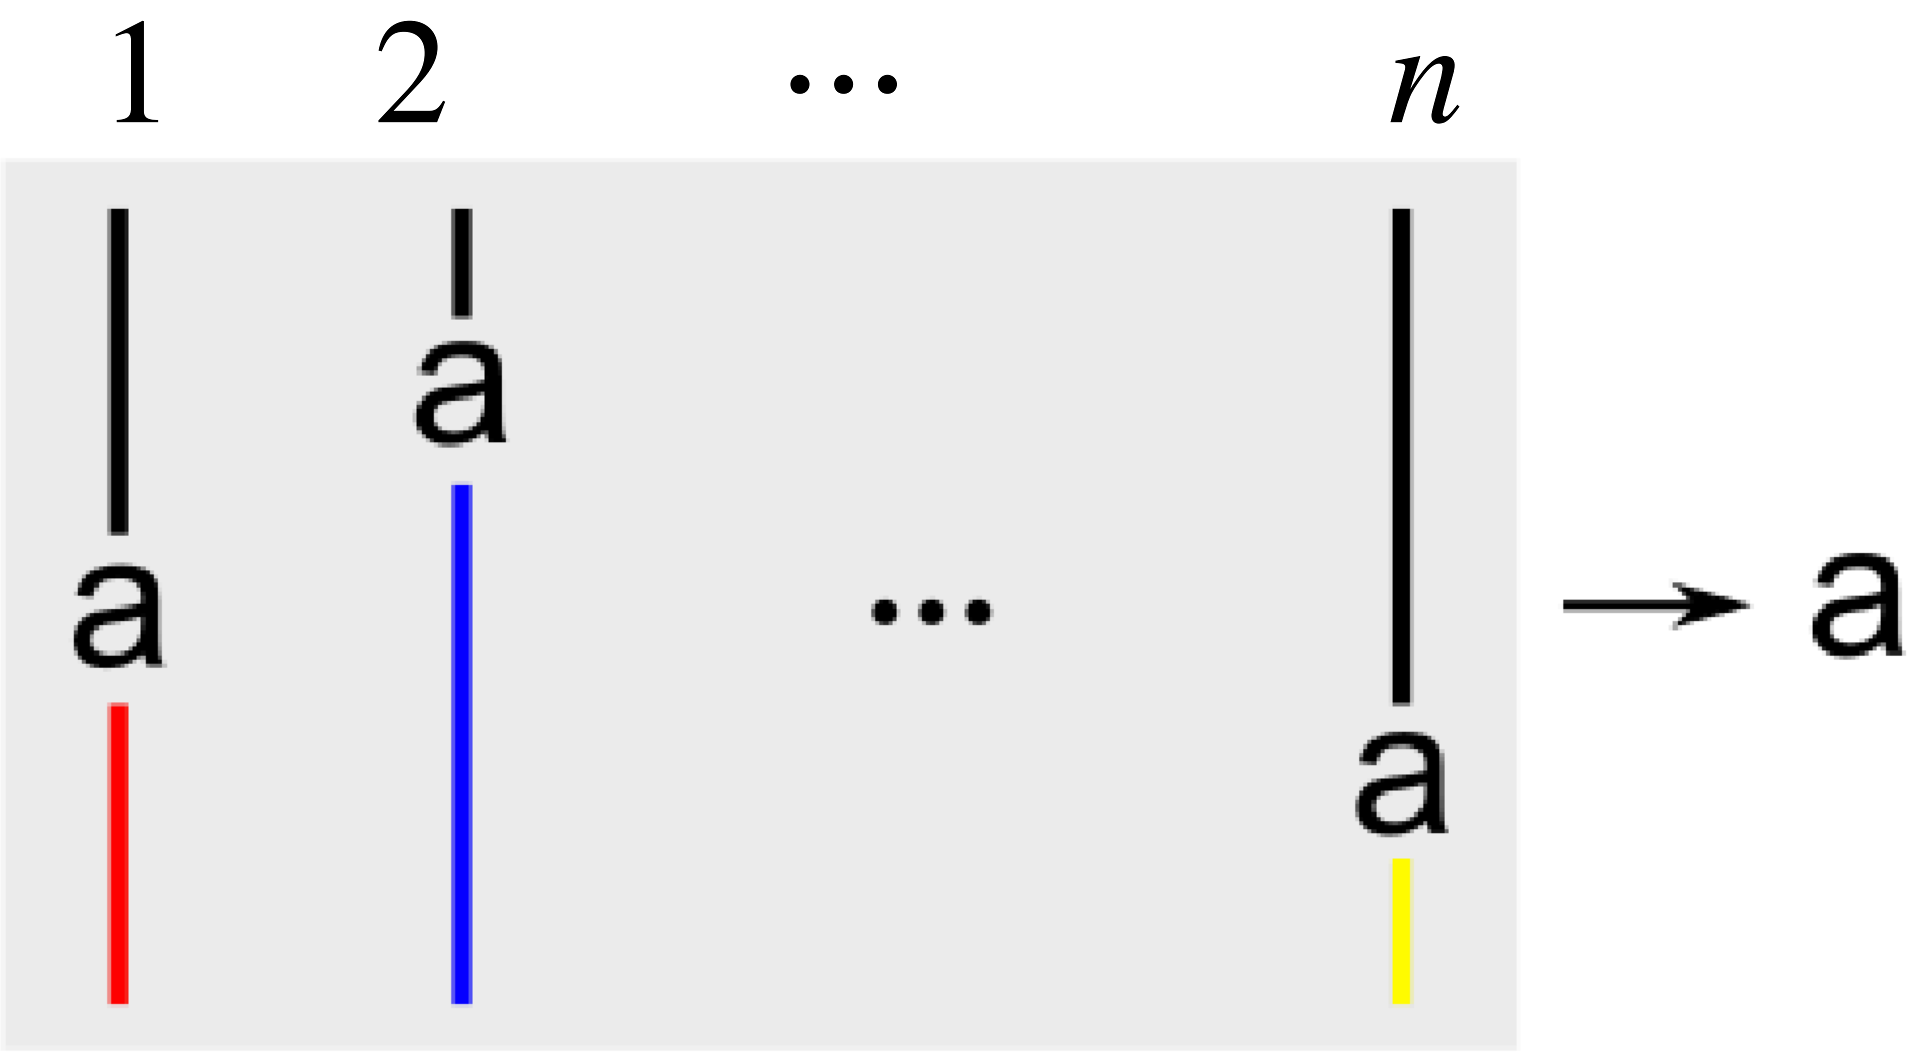

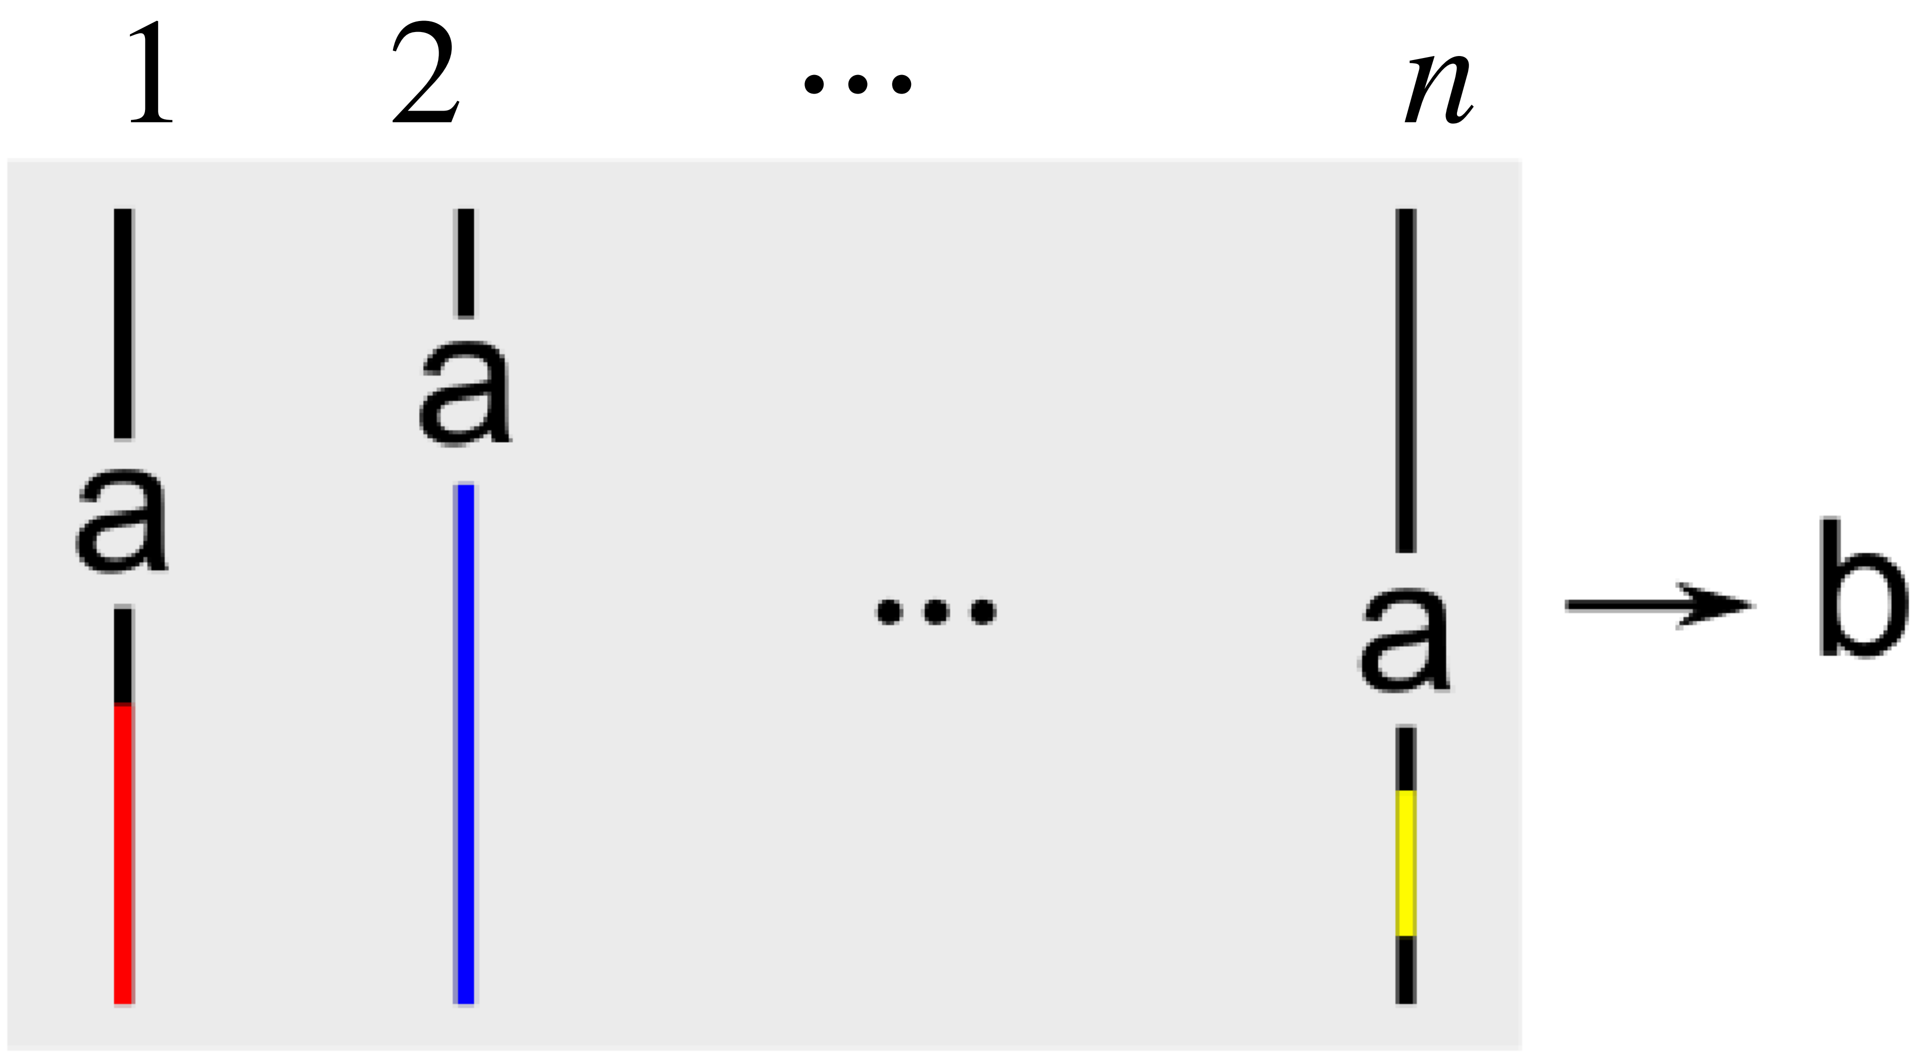

Image credit: https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.p

Strategyproof $\implies$ Monotone

- Suppose a rule is strategyproof but not monotone
- Strategyproof means:
  - No voter can change their individual ranking to make a more preferred candidate win
- **Not monotone** means:
  - Suppose a is the current winner (on profile L). For all input profiles L', in which for all voters, any candidate who was ranked below a in L is still ranked below a in L', then it is still possible for another candidate b to win in L'.

Image credit: https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.p

Strategyproof $\implies$ Monotone

- Suppose a rule is strategyproof but not monotone

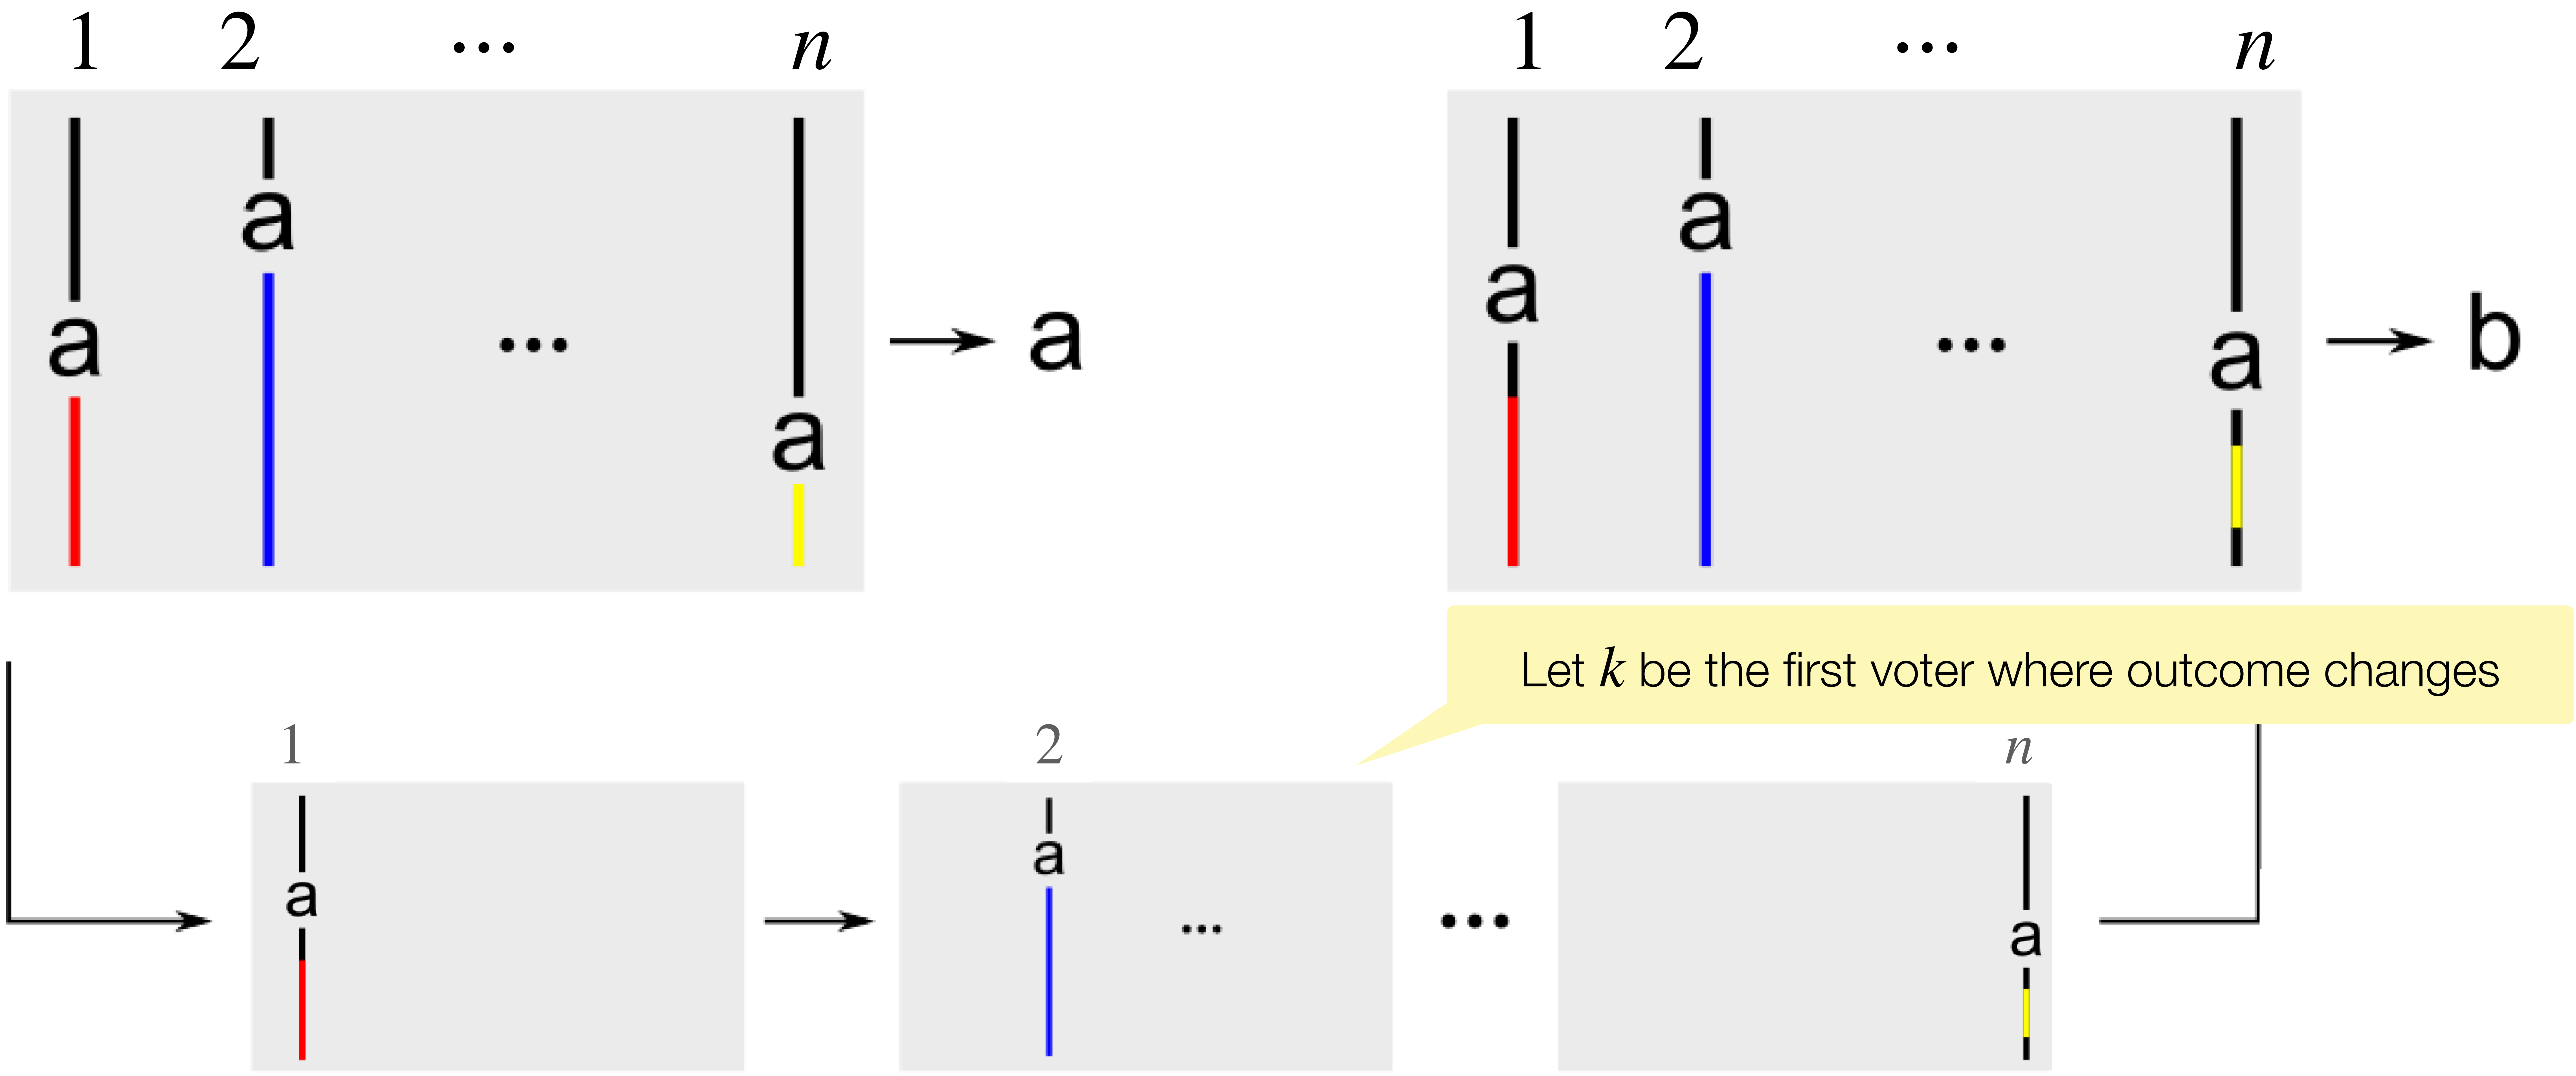

Image credit: https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.p

Strategyproof $\implies$ Monotone

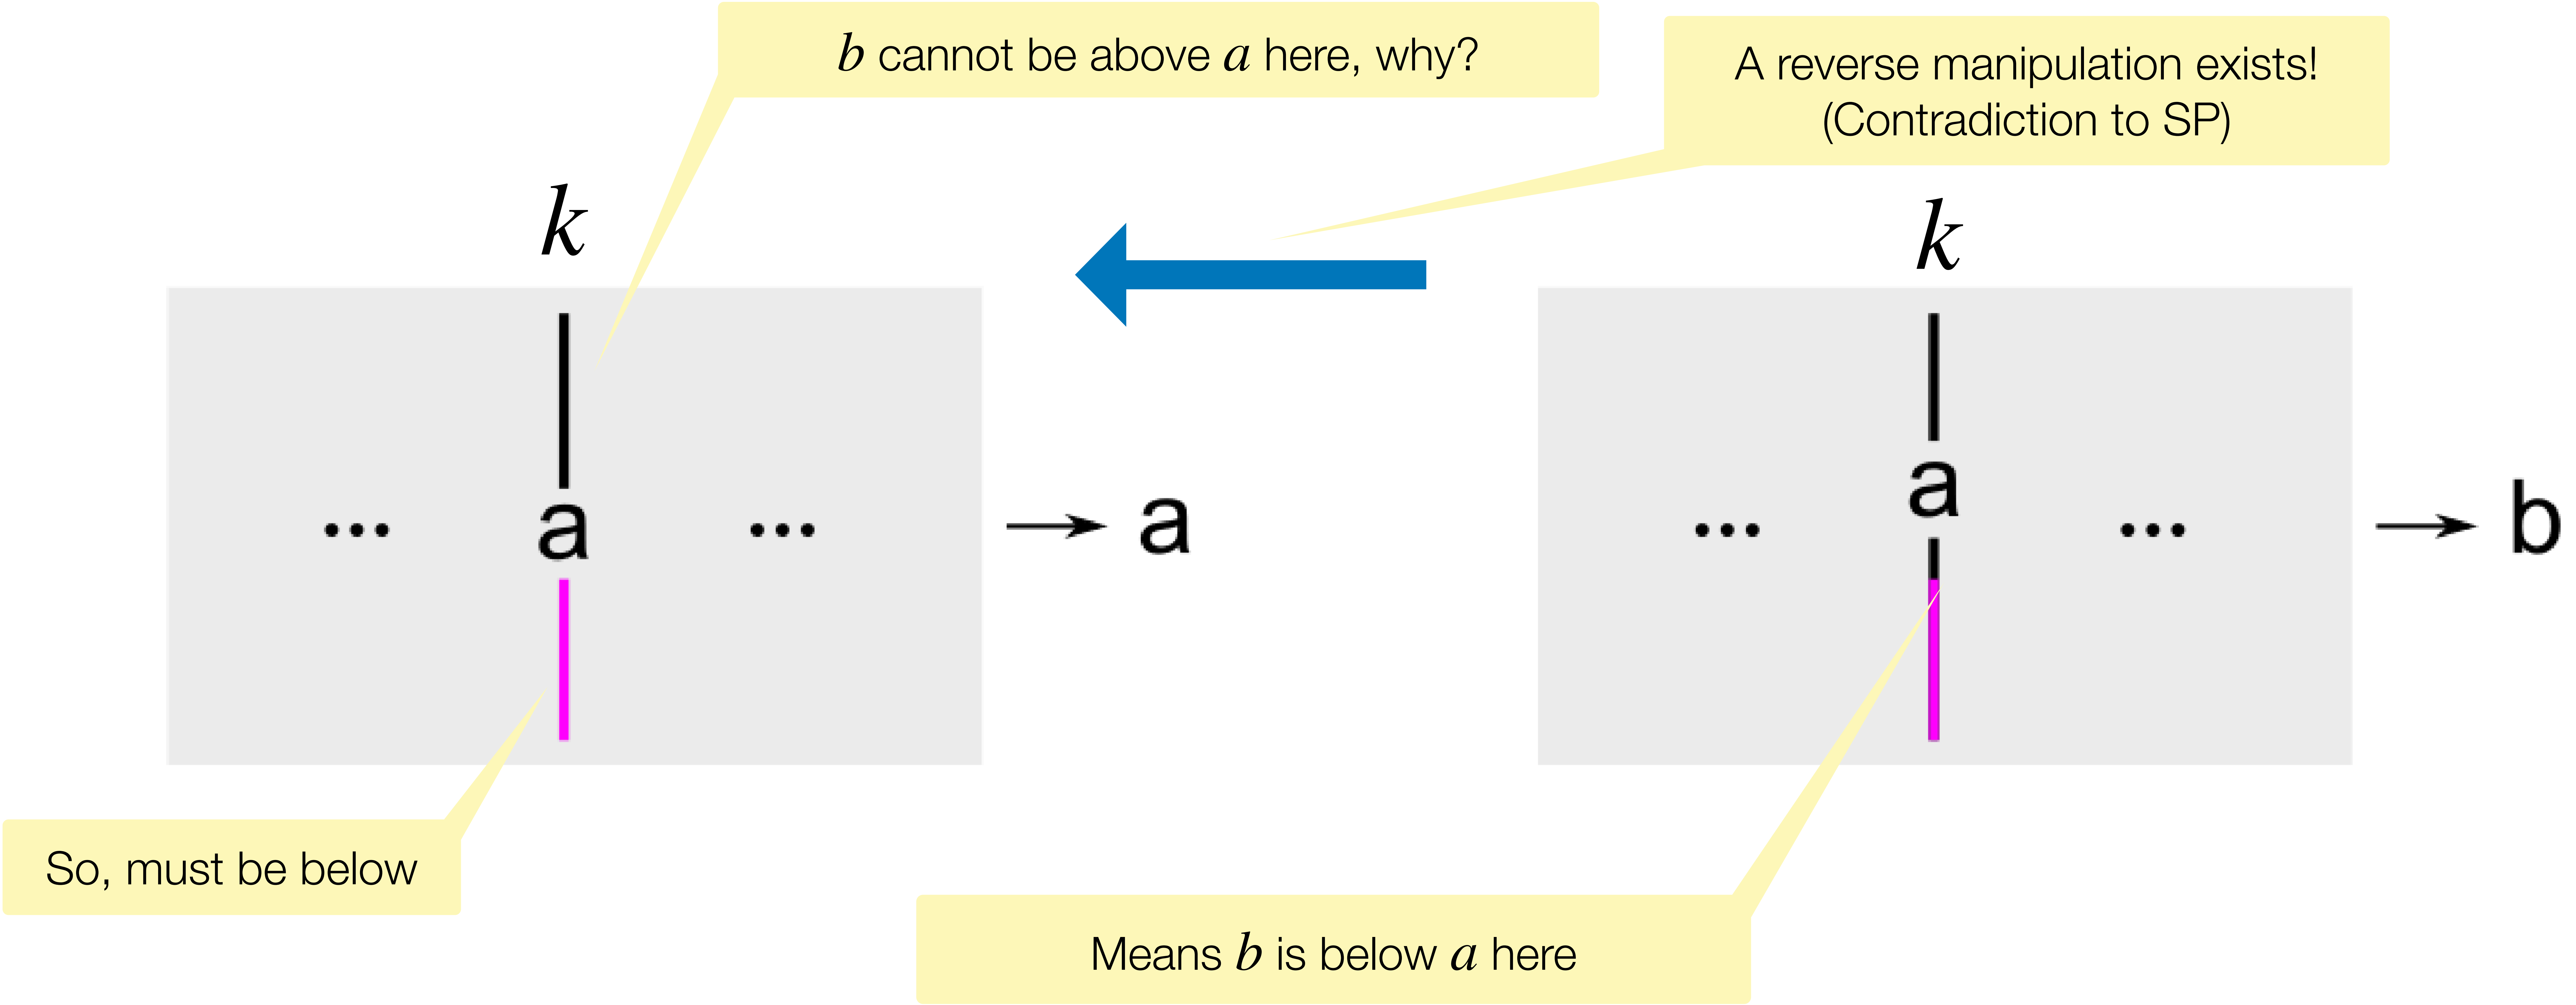

Image credit: https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.p

Monotone $\implies$ Strategyproof

- Suppose there is a voter v_k that prefers b to a
- Consider truthful instance on left where a wins

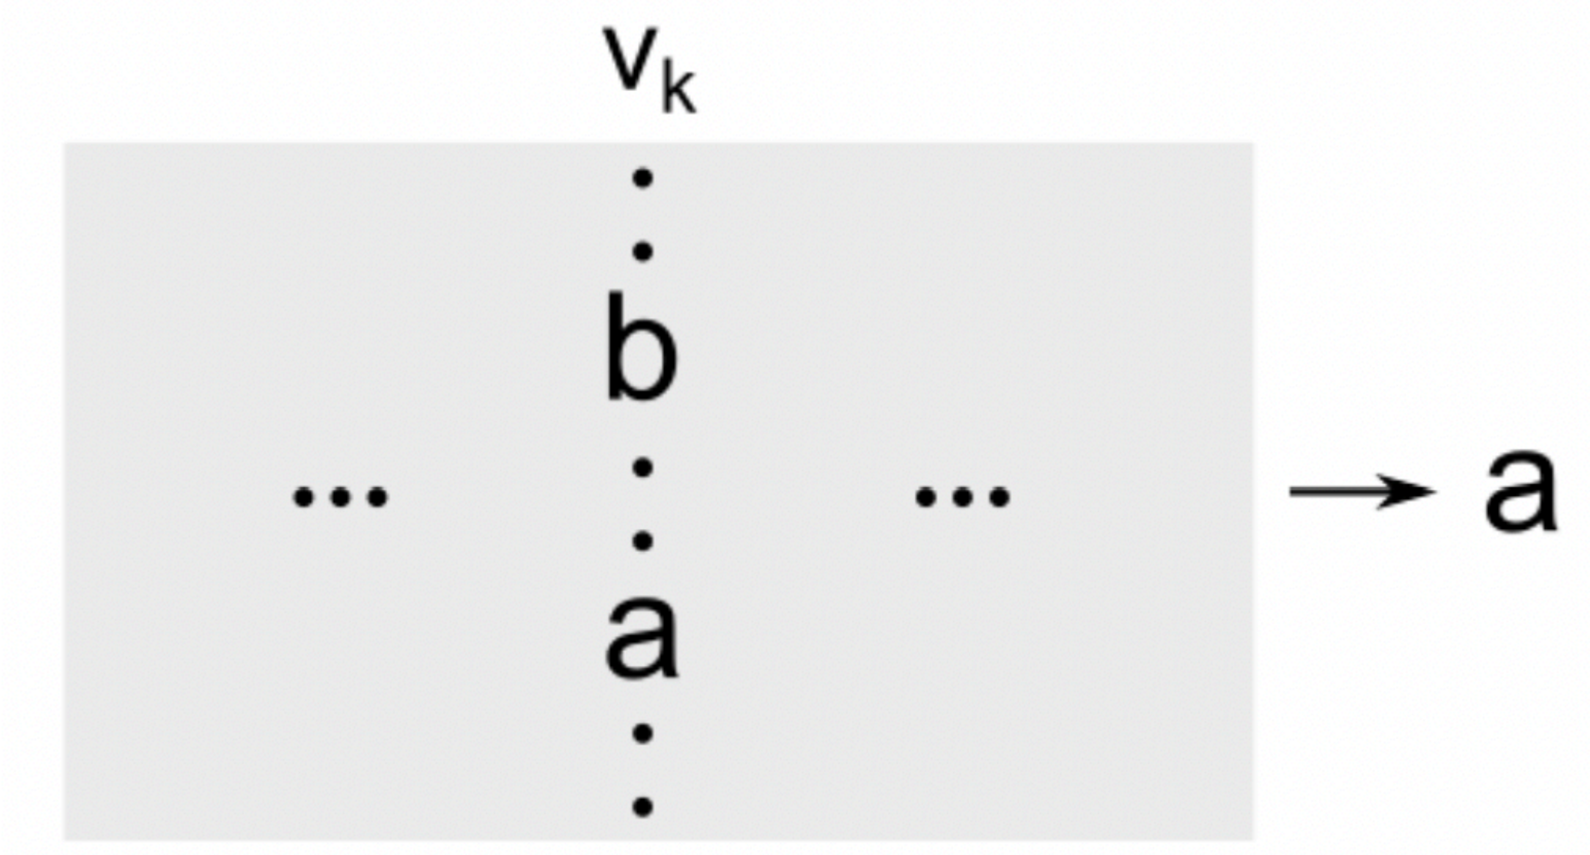

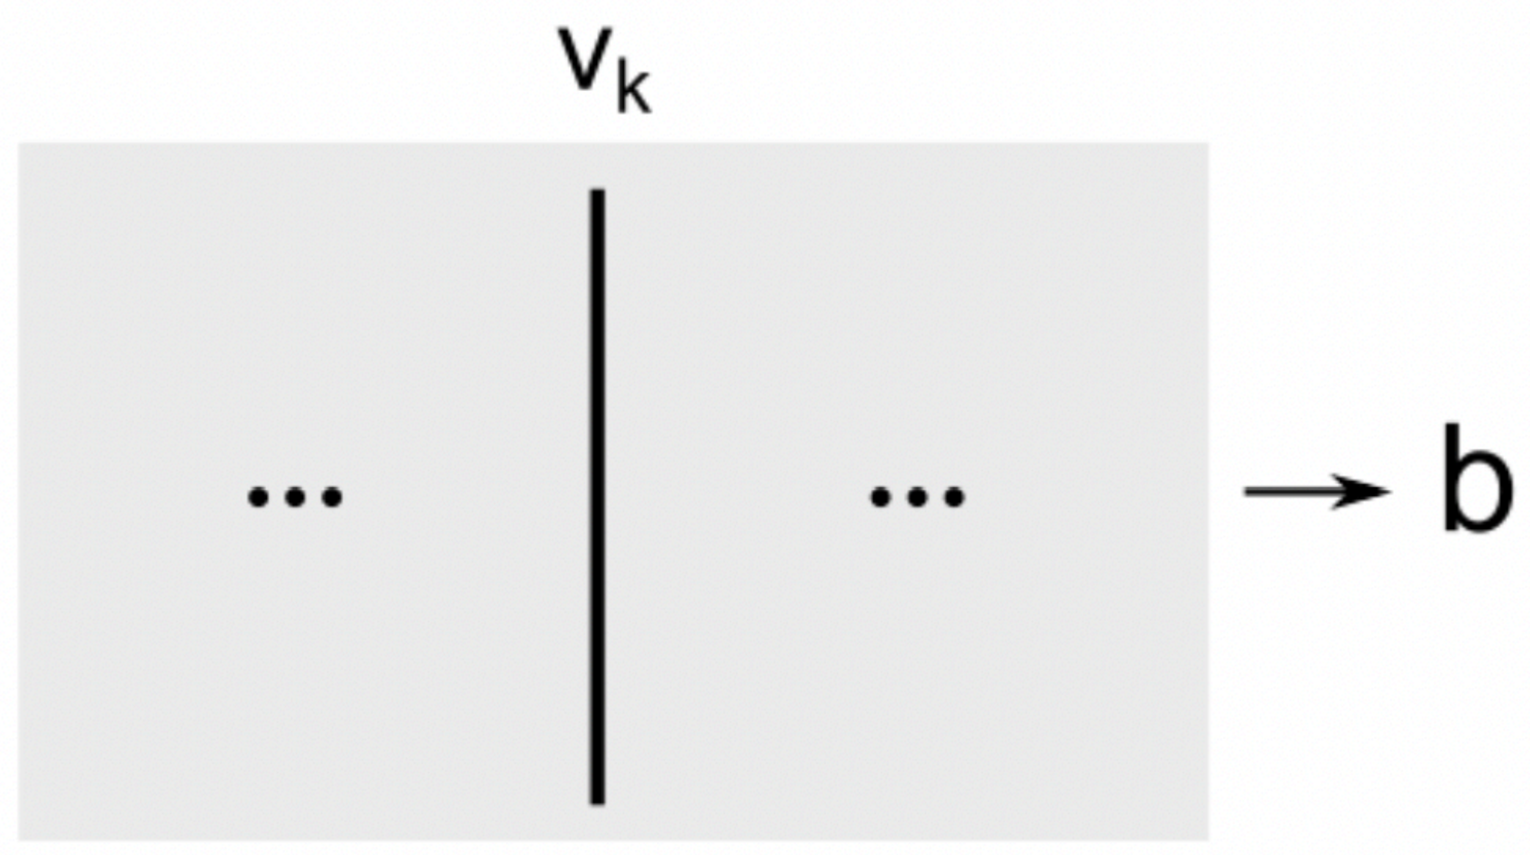

Image credit: https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.p

Monotone $\Longrightarrow$ Strategyproof

- Suppose there is a voter v_k that prefers b to a
- Consider truthful instance on left where a wins

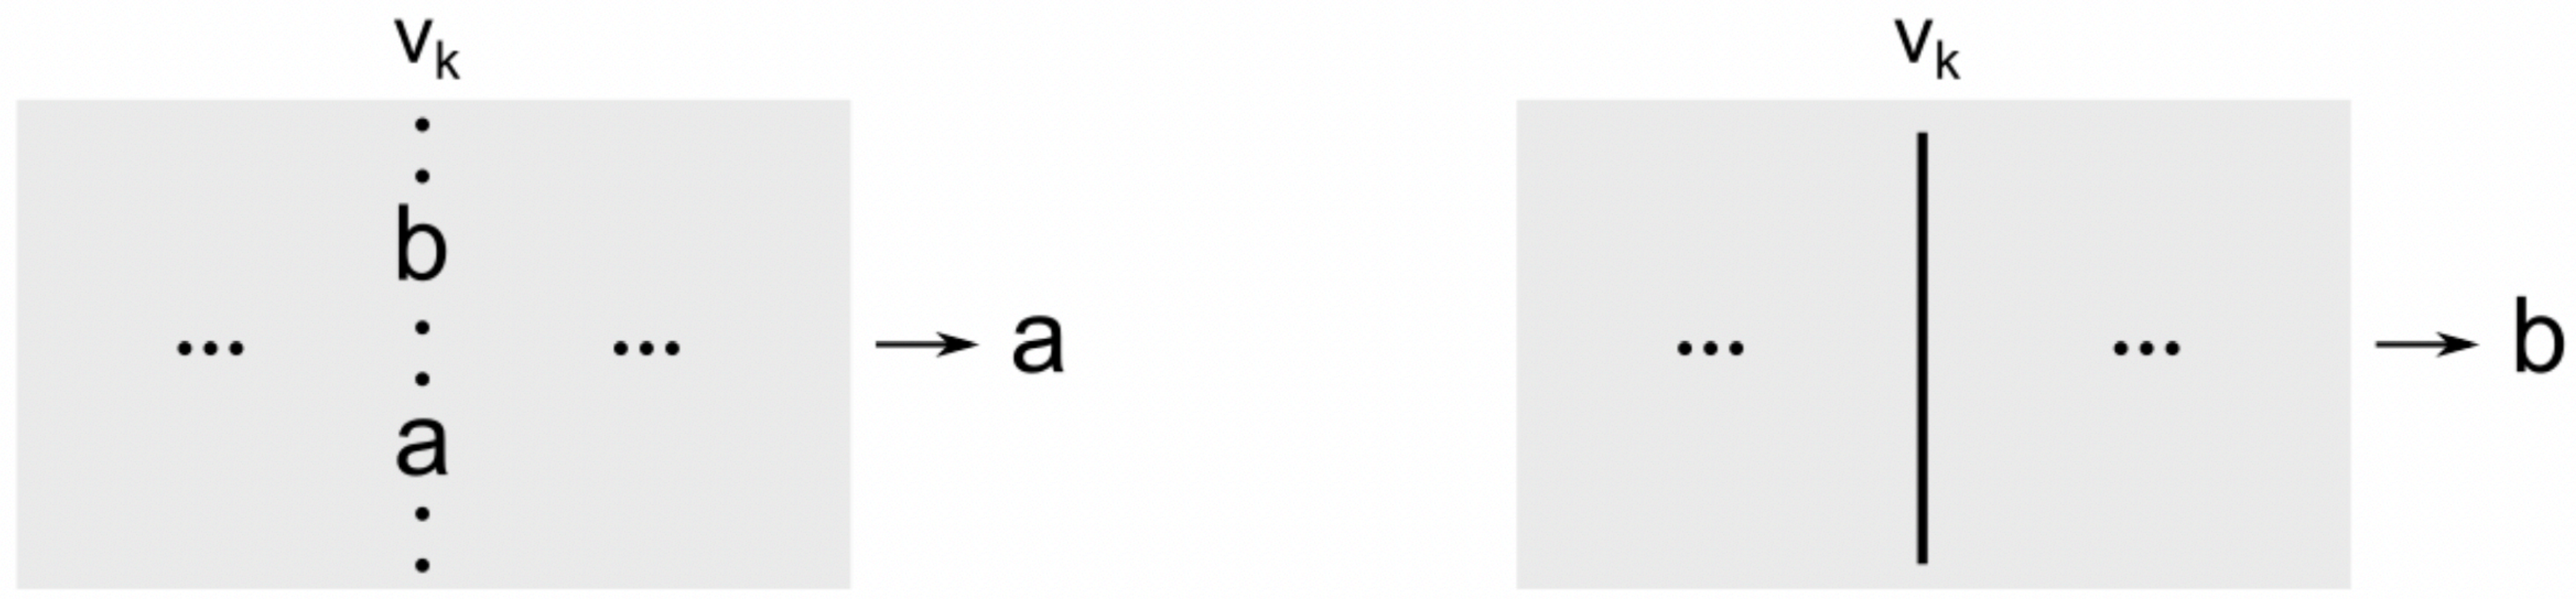

Image credit: https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.p

Monotone $\Longrightarrow$ Strategyproof

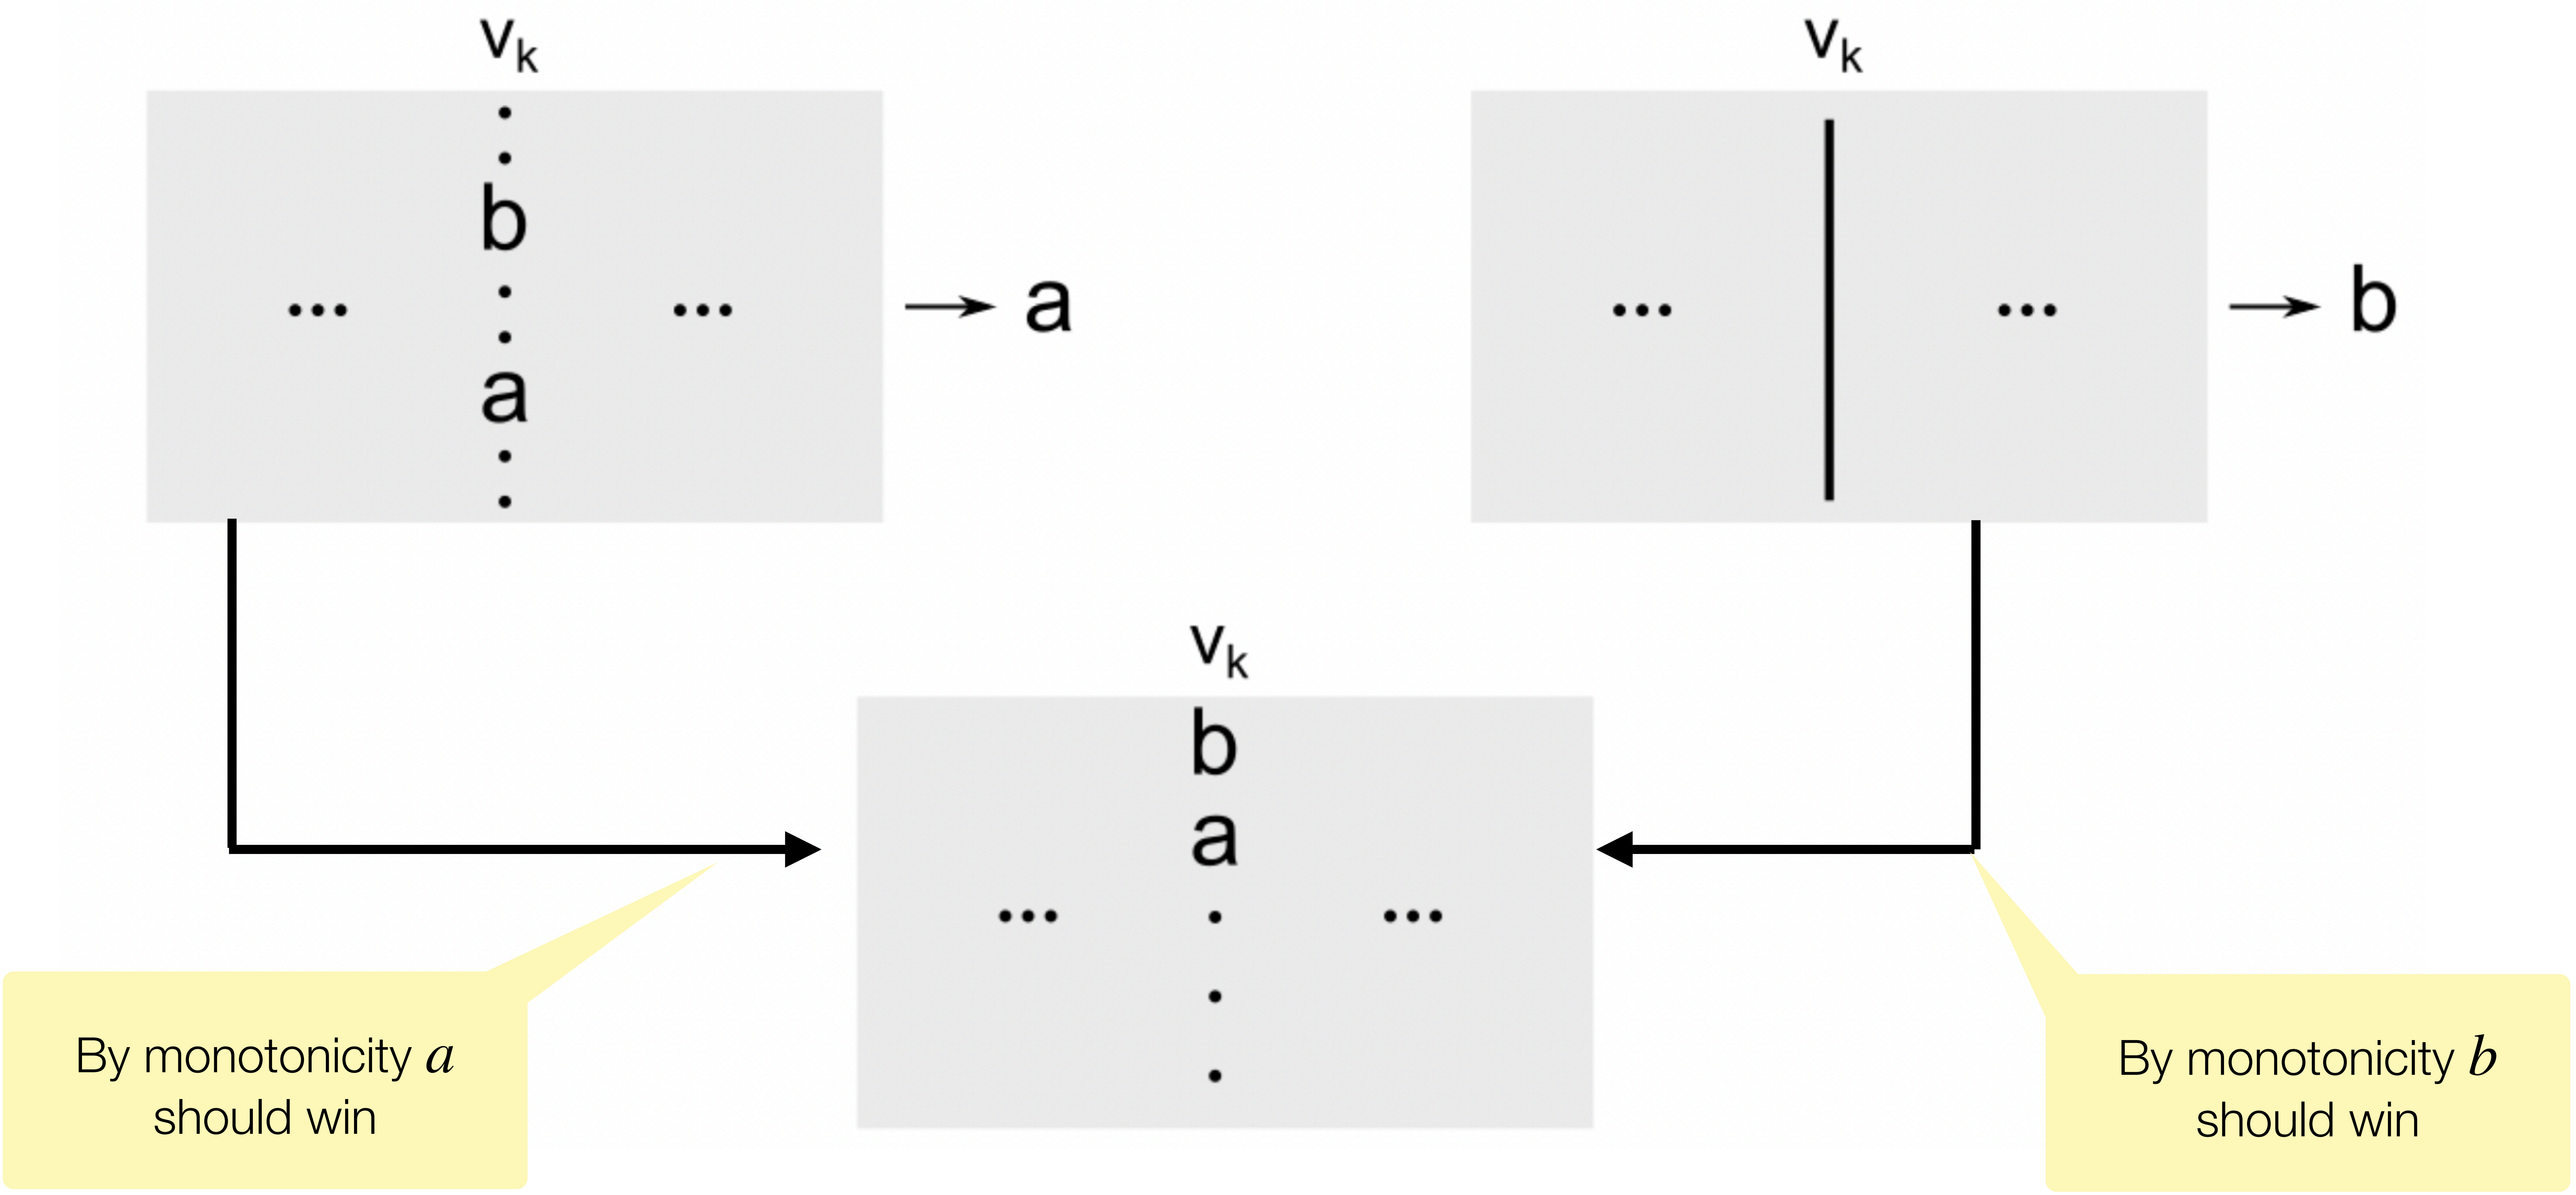

Image credit: https://rohitvaish.in/Teaching/2022-Spring/Slides/Lec%202.p

$\boxed{\text{Strategyproof} \iff \text{Monotonicity}}$

[GS Theorem] With three or more candidates, a voting rule is **strategyproof** and **onto** if and only if it is a dictatorship.

$\boxed{\text{Part 1. Strategyproof} \iff \text{Monotonicity}}$

[Alternate Statement] With three or more candidates, a voting rule is **monotone** and **unanimous** if and only if it is dictatorship.

$\boxed{\text{Part 2. Monotone + Onto} \implies \text{Unanimous}}$

SP + Onto $\Longrightarrow$ Unanimous

- **Definition (Unanimity)**. Given preference profile L, if there is an alternative a that every voter prefers to b, then $f(L) \neq b$.
- **Lemma**. SP + Onto $\Longrightarrow$ Unanimous

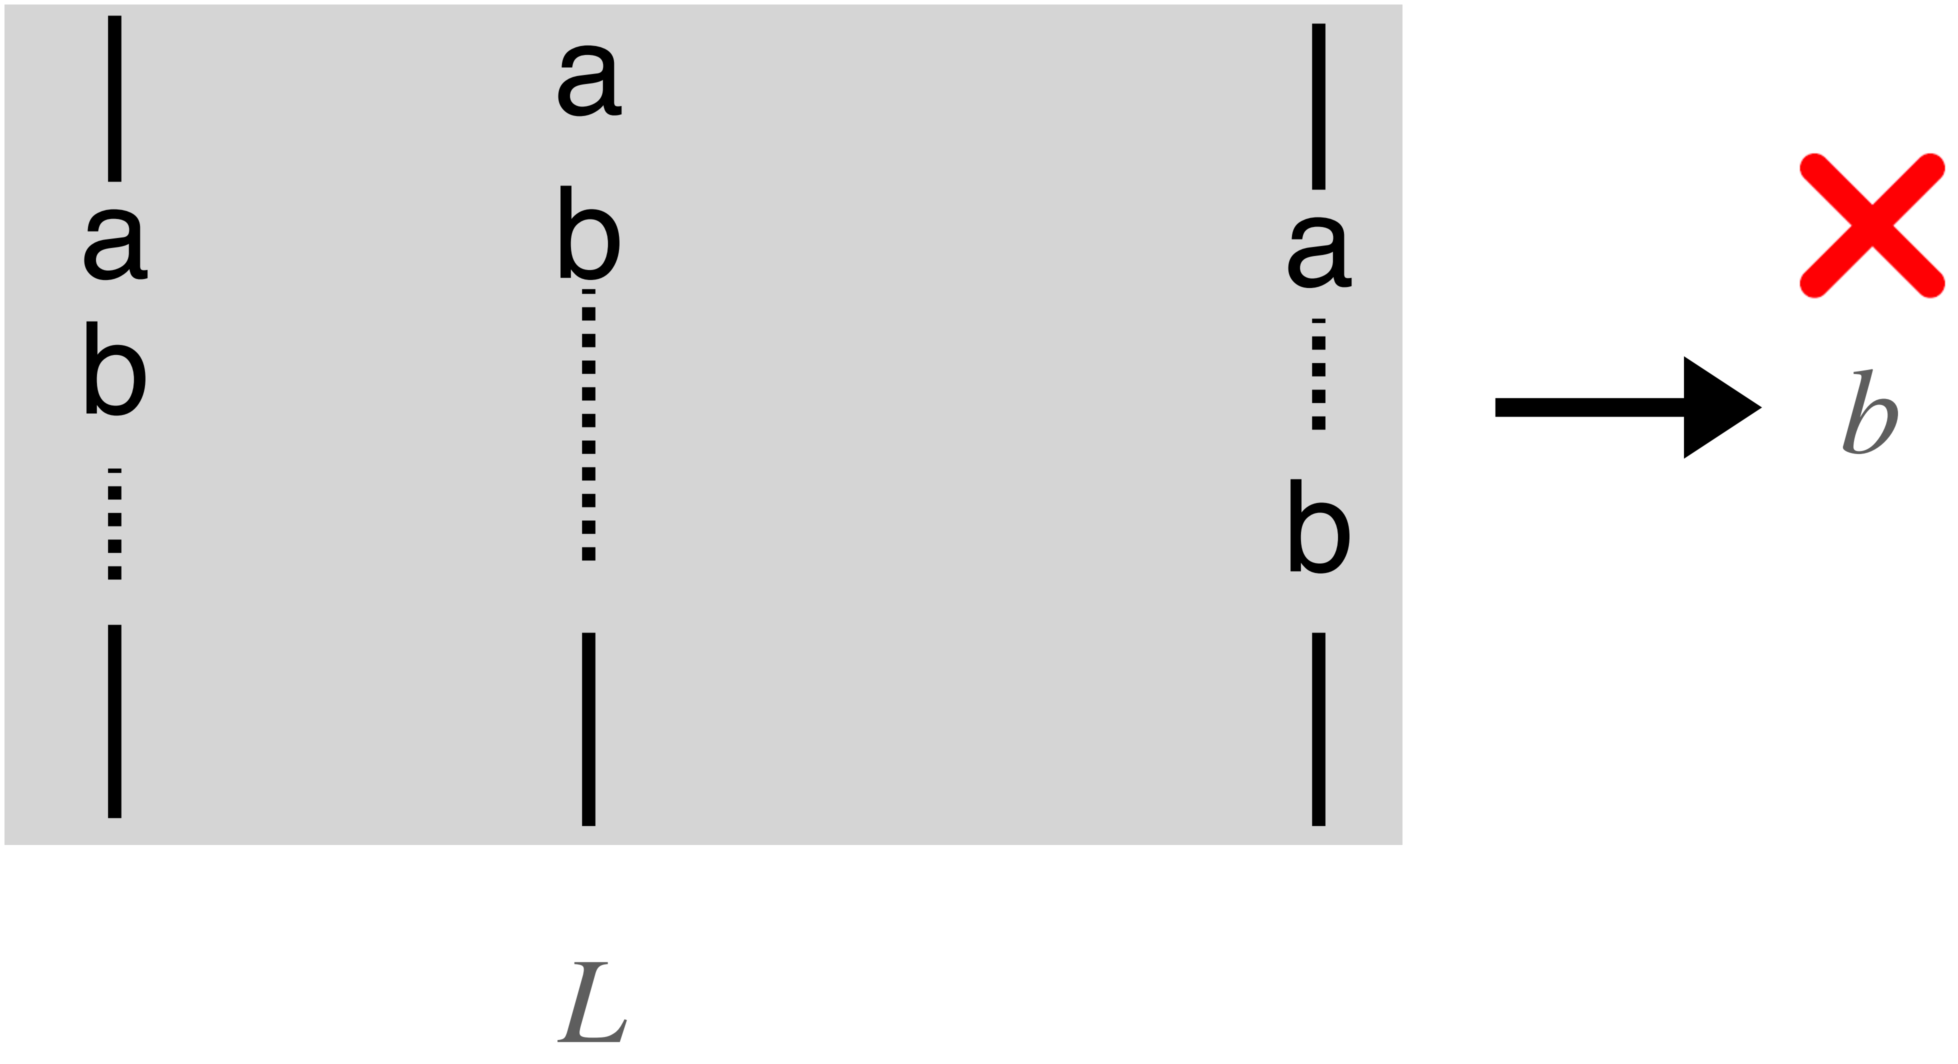

SP + Onto $\Longrightarrow$ Unanimous

- **Definition (Unanimity)**. Given preference profile L, if there is an alternative a that every voter prefers to b, then $f(L) \neq b$.
- **Lemma**. SP + Onto $\Longrightarrow$ Unanimous
- **Proof**. Suppose $f(L) = b$. Consider L' below. $f(L') = ?$

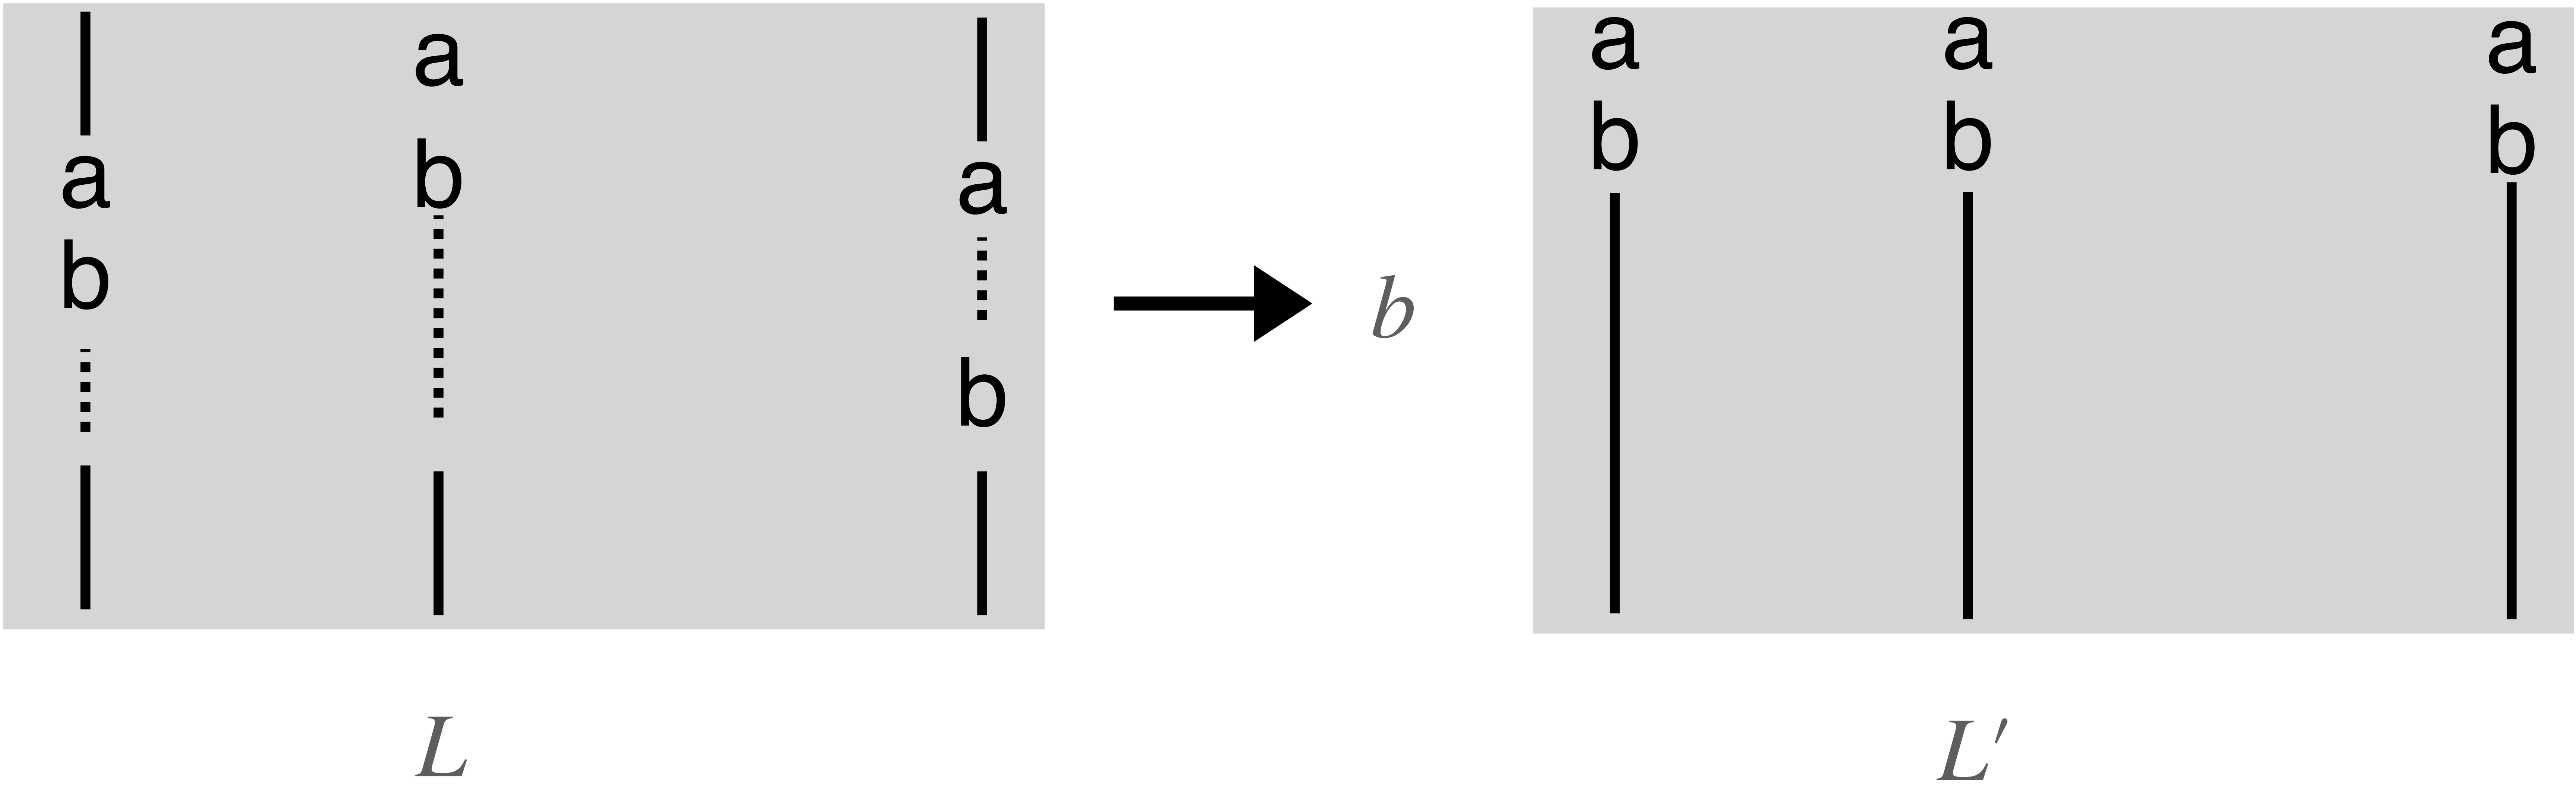

SP + Onto $\Longrightarrow$ Unanimous

- **Definition (Unanimity)**. Given preference profile L, if there is an alternative a that every voter prefers to b, then $f(L) \neq b$.
- **Lemma**. SP + Onto $\Longrightarrow$ Unanimous
- **Proof**. Suppose $f(L) = b$. Consider L' below. $f(L') = ?$

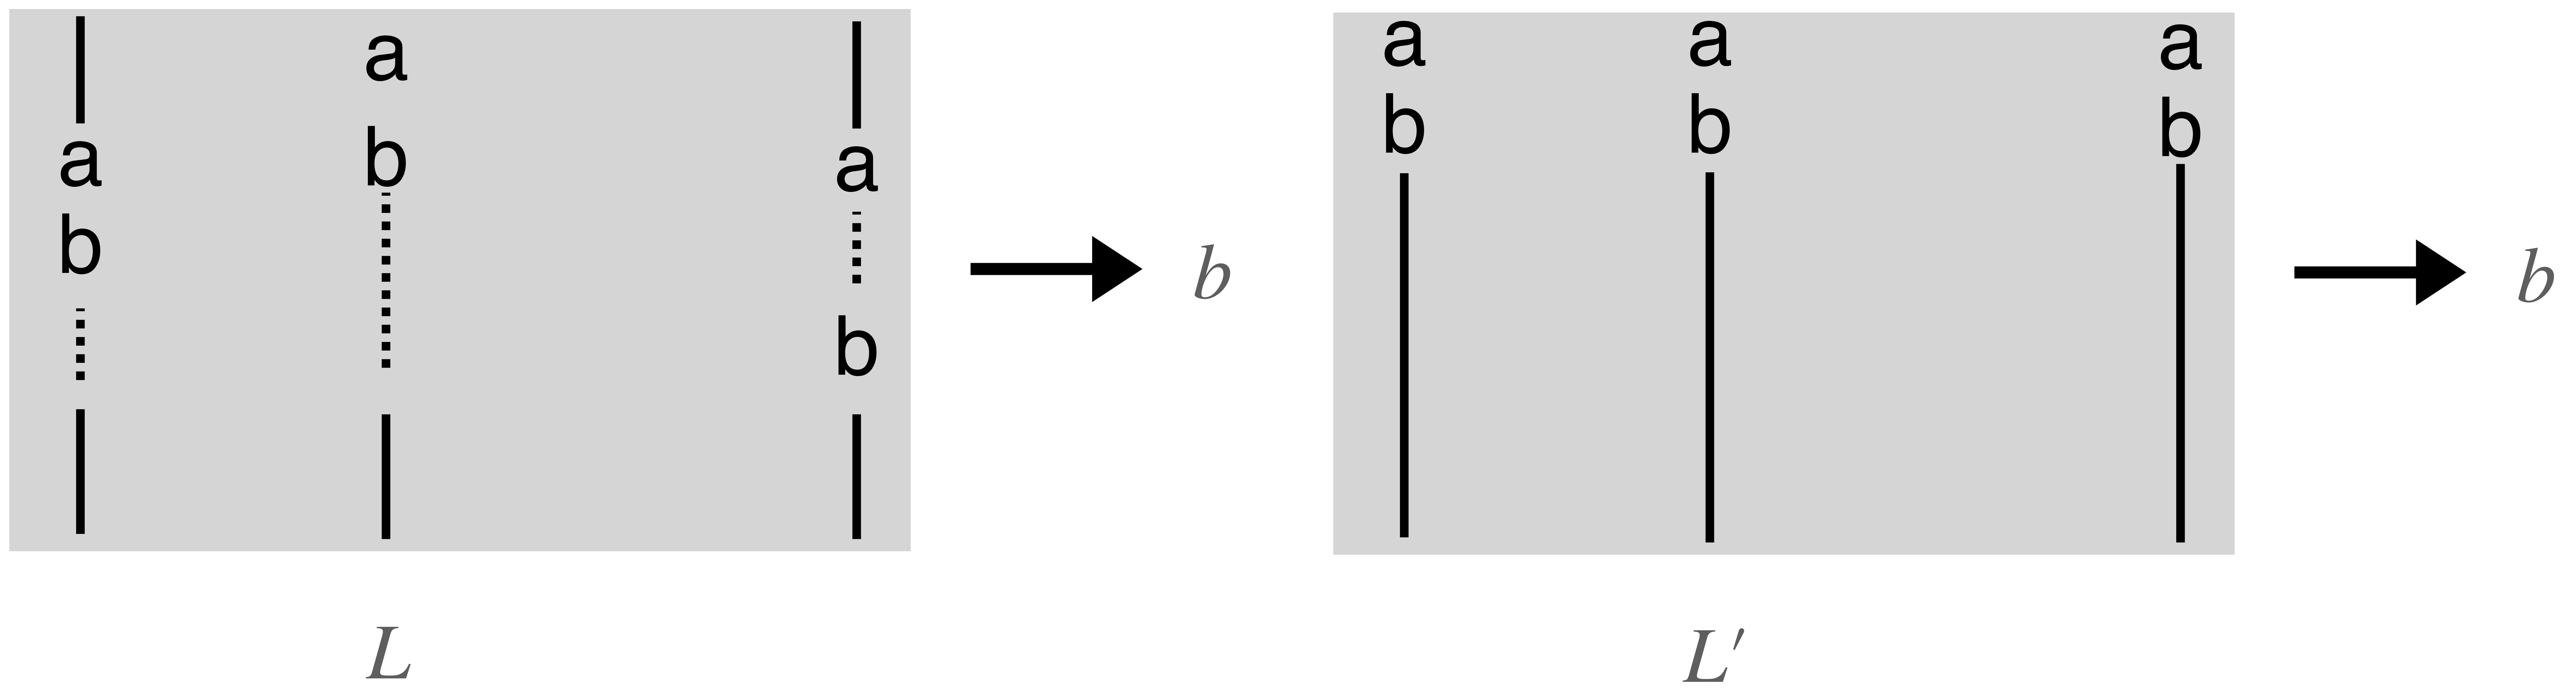

SP + Onto $\Longrightarrow$ Unanimous

- **Definition (Unanimity)**. Given preference profile L, if there is an alternative a that every voter prefers to b, then $f(L) \neq b$.
- **Lemma**. SP + Onto $\Longrightarrow$ Unanimous
- **Proof**. We know $f(L') = b$ by monotonicity. By onto, there exists a profile L'' where a wins.

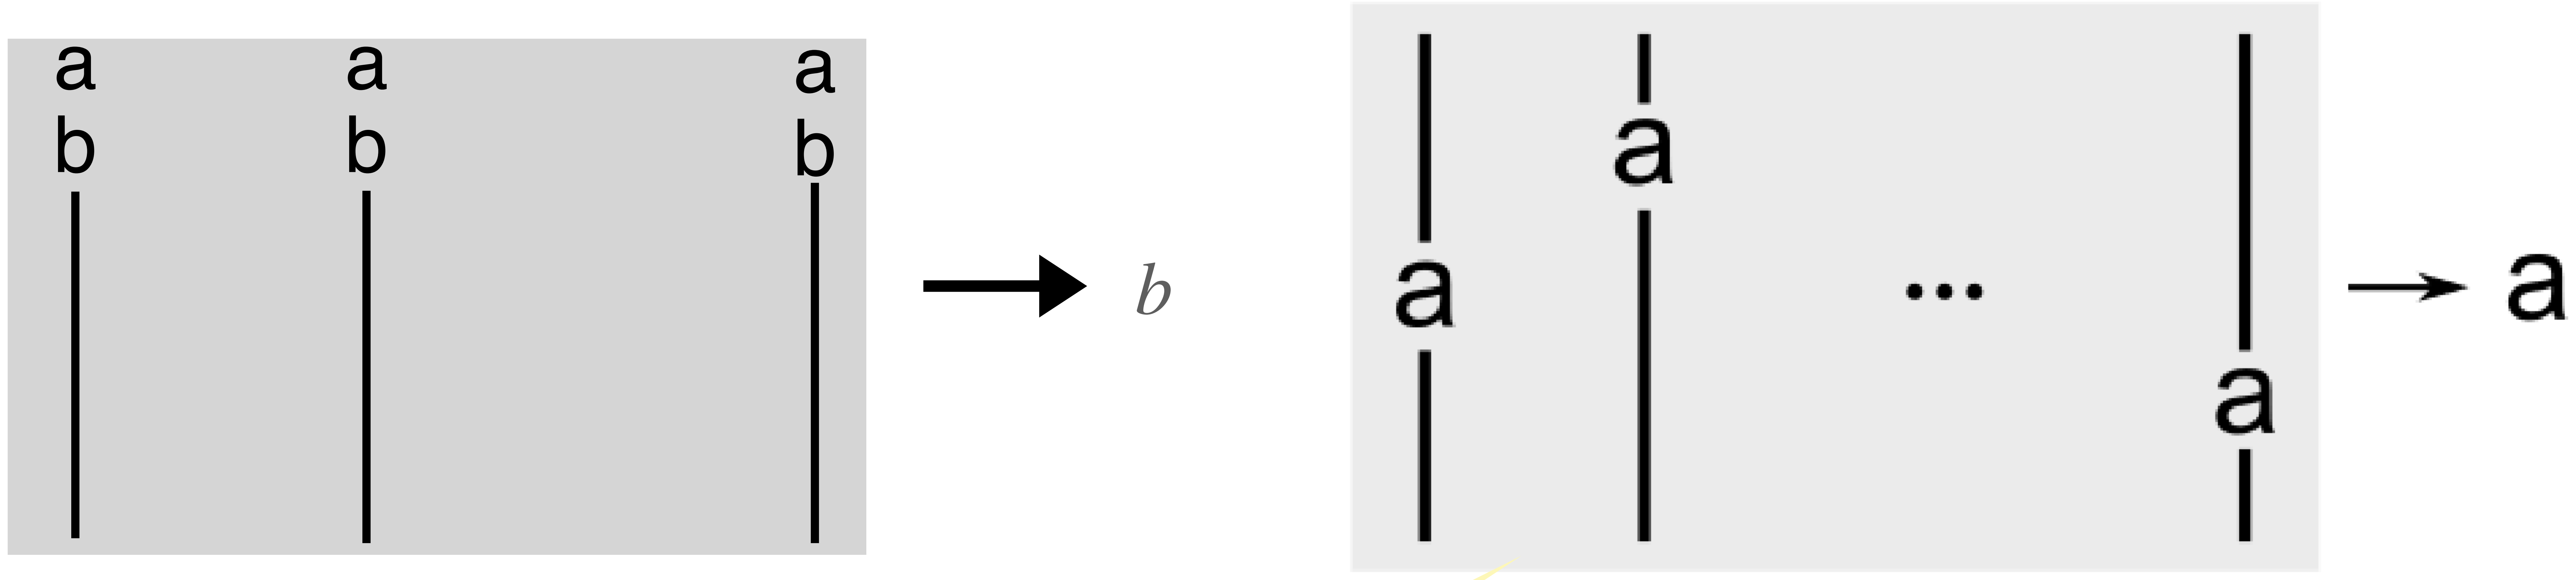

L'

L''

L'' to L', a's support only goes up,
by monotonicity b cannot win.

[GS Theorem] With three or more candidates, a voting rule is **strategyproof** and **onto** if and only if it is a **dictatorship**.

$\boxed{\textbf{Goal.}\ \text{Strategyproof} + \text{Onto} \Longrightarrow \text{Dictatorship}}$

[Proof Outline]

$\boxed{\text{Part 1.}\ \text{Strategyproof} \Longleftrightarrow \text{Monotonicity}}$

$\boxed{\text{Part 2.}\ \text{Monotone} + \text{Onto} \Longrightarrow \text{Unanimous}}$

$\boxed{\textbf{Part 3.}\ \text{Monotone} + \text{Unanimous} \Longrightarrow \text{Dictatorship}}$

Next time