Social Choice

CSC304 Lecture 21

November 28, 2016

Allan Borodin

Adapted from Craig Boutilier’s slides

Todays agenda and announcements

- Today: Review of popular voting rules. Axioms, Manipulation, Impossibility Results. Single-peaked preferences
  - Reading : Ch.13 (plus some ideas not discussed in the text)
  - Next week: Ch. 11   Fair Division
- Announcements
  - Last assignment is due noon November 30. A3.pdf to submit
  - Office hours by appointment for next couple weeks.
  - There may be a question on final exam regarding fair division.
  - Exam is 3 hours. Eight questions, 130 points. Mainly covering second half of course but there will be one basic game theory question. Usual 20% for saying you do not know how to answer a question. We allow one sheet, both sides, of *hand-written* notes No other aids allowed.

Voting and Preference Aggregation

- Last time
  - Introduced *social choice*: preference aggregation to make a single "consensus" decision for a group
  - The concept of a *voting rule*:
    - Given: a set N of n voters and a set A of m alternatives
    - Input: a preference profile (a ranking of alternatives by each voter)
    - Output: winning alternative from A
    - Also discussed the idea of deriving a *consensus ranking* over A
  - Different voting rules (Plurality, Borda, approval, STV, etc.) and properties
    - Different rules give different results on same profiles!

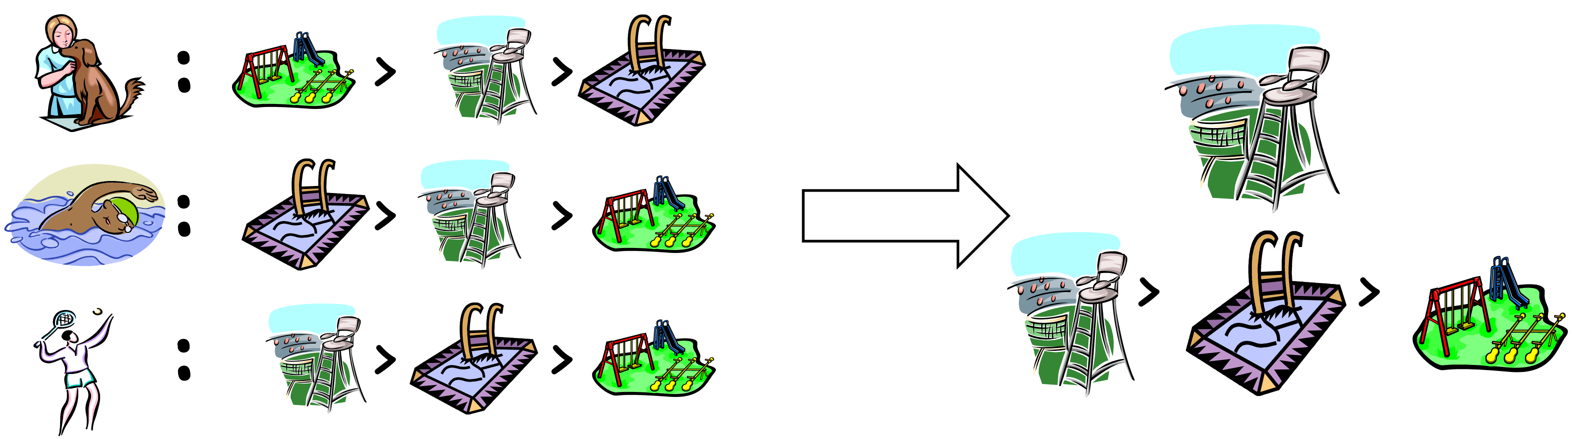

Plurality Voting

- *Plurality voting:*
  - **Input:** rankings of each voter
  - **Winner:** alternative ranked 1st by greatest number of voters
    - number of 1st-place rankings is *a*'s *plurality score*
    - *complete* rankings not needed, just votes for most preferred alternatives
    - we'll ignore ties for simplicity
  - This is a most familiar scheme, used widely:
    - locally, provincially, nationally for electing political representatives
  - With only 2 alternatives, often called *majority voting*

- Example preference profile (three alternatives):
  - $A \succ B \succ C$:   5 voters
  - $C \succ B \succ A$:   4 voters
  - $B \succ C \succ A$:   2 voters
- Winner: A wins (plurality scores are   A: 5;   C: 4;   B:2)

The Borda Rule

- *Borda voting rule:*
  - **Input:** rankings of each voter
  - *Borda score* for each alternative *a*: *a* gets m-1 points for every 1st-place rank, *m-2* points for every 2nd-place, etc.
  - **Winner:** alternative with highest Borda score
  - Used in sports (Heismann, MLB awards), variety of other places
  - Proposed by Jean-Charles, chevalier de Borda in 1770 to elect members to the French Academy of Sciences (also Ramon Llull, 13th century)

- Example profile (three alternatives, positional scores of 2, 1, 0):
  - $A \succ B \succ C$: 5 voters
  - $C \succ B \succ A$: 4 voters
  - $B \succ C \succ A$: 2 voters
- Winner: B wins (Borda scores are: B: 13; A: 10; C: 10)
  - Notice: more sensitive to *the entire range of preferences* than plurality is (which ranked *B* last)

Approval Voting

- Approval Voting
  - **Input:** voters specify a *subset* of alternatives they "approve of"
  - Approval score: a point given to a for each approval
    - variant: *k*-approval, voter lists exactly *k* candidates
  - **Winner:** alternative with highest approval score
  - Used in many informal settings (at UN, Doge of Venice, …)
  - Steven Brams a major advocate (see Wikipedia article)

- Example profile (three alternatives, approvals in bold):
  - **A** $\succ$ B $\succ$ C: 5 voters (approve of only top alternative)
  - **C** $\succ$ B $\succ$ A: 4 voters (approve of only top alternative)
  - **B** $\succ$ **C** $\succ$ A: 2 voters (approve of top two alternatives)
- Winner: C wins (approval scores are: C: 6; A: 5; B: 2)
  - Notice: can't predict vote based on ranking alone!

Positional Scoring (Voting) Rules

- Observe that plurality, Borda, *k*-approval, *k*-veto are all each *positional scoring rules*
- Each assigns a *score* $\alpha(j)$ to each rank position j
  - almost always non-increasing in j
- The winner is the candidate a with max total score: $\sum_i \alpha(r_i(a))$

In general:	$\alpha(1)$	$\alpha(2)$	$\alpha(3)$	$\alpha(4)$
Plurality:	1	0	0	0
Borda:	3	2	1	0
2-Approval:	1	1	0	0
Veto:	1	1	1	0
and another:	10	2	0	0

Which of these is Better?

- Notice that on the same vote profile, plurality, Borda, and approval gave different winners!

- Which is best?
  - hard to say: depends on social objective one is trying to meet
  - common approach: identify *axioms/desirable properties* and try to show certain voting rules satisfy them
    - we will see it is not possible in general!
- Note: all these voting rules must have some tie breaking breaking rule or allow for a re-vote. In some cases, that rule is simply a flip of the coin. See the tie vote in a 2015 election in Mississippi. Even with a large number of voters ties can happen.
- Let's now look at a few more voting rules to get a better sense of things.

There are Hundreds of Voting Rules

- *Single-transferable vote (STV)* or *Hare* system
  - Round *1*: vote for favorite candidate; eliminate candidate with lowest plurality score;
  - Round *t*: if your favorite eliminated at round *t-1*, recast vote for favorite remaining candidate; eliminate candidate with lowest plurality score
  - Round *m-1*: winner is last remaining candidate
    - terminate at any round if plurality score of top candidate is at least *n/2* (i.e., there is a majority winner)
  - Used: Australia, New Zealand, Ireland, Some variant of this is used in political conventions.

    Needn't be online: voters can submit rankings once
  - When would this be a bad voting rule?
- *Nanson's rule*
  - Just like STV, but use Borda score to eliminate candidates

There are Hundreds of Voting Rules

- *Egalitarian (maxmin fairness)*
  - Winner maximizes minimum voter's rank: $argmax_a \; min_j \; (m\text{-}r_j\,(a))$

- *Copeland*
  - Let $W(a,b,\mathbf{r}) = 1$ if more voters rank $a \succ b$; 0 if more $b \succ a$; ½ if tied
  - Score $s_c(a,\mathbf{r}) = \sum_b W(a,b,r)$; winner is *a* with max score
    - *i.e., winner is candidate that wins most pairwise elections*

- *Tournament/Cup*
  - Arrange a (usually balanced) tournament tree of pairwise contests
  - Winner is last surviving candidate
  - We'll discuss this in more detail later

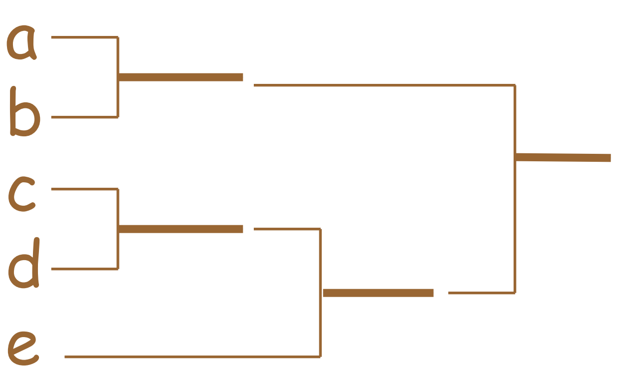

Condorcet Principle

- How would you determine "societal preference" between a pair of alternatives *a* and *b*?
- A natural approach: run a "pairwise" majority vote: if a *majority* of voters prefer *a* to *b*, then we say *the group prefers a to b*
- *Condorcet winner:* an alternative that beats every other in a pairwise majority vote
  - proposed by Marie Jean Antoine Nicolas de Caritat, marquis de Condorcet in 1785
  - if there is a Condorcet winner, it must be unique
  - a rule is *Condorcet-consistent* if it selects the Condorcet winner (if one exists)

- Condorcet winners need not exist (next slide)
  - Moreover, many natural voting rules are not Condorcet consistent (e.g., plurality, Borda, STV are not), but some are: Nanson, Copeland, Cup, etc.

Condorcet Paradox

- *Condorcet paradox:*
  - suppose we use the pairwise majority criterion to produce a societal preference ranking
  - pairwise majority preferences may induce *cycles* in societal ranking (i.e., the preference relation is not transitive)
- Simple example:
  - A $\succ$ B $\succ$ C: *m/3* voters
  - C $\succ$ A $\succ$ B: *m/3* voters
  - B $\succ$ C $\succ$ A: *m/3* voters
  - Societal ranking has $A \succ B$, $B \succ C$, and $C \succ A$ (!)
  - No clear way to produce a consensus ranking
  - Also evident that this preference profile has no Condorcet winner

Violations of Condorcet Principle

- Plurality violates Condorcet
  - 499 votes: $A \succ B \succ C$
  - 3 votes: $B \succ C \succ A$
  - 498 votes: $C \succ B \succ A$
  - plurality choses A; but B is a CW (B>A 501:499; B>C 502:498)

- Borda violates Condorcet
  - 3 votes: $A \succ B \succ C$
  - 2 votes: $B \succ C \succ A$
  - 1 vote: $B \succ A \succ C$
  - 1 vote: $C \succ A \succ B$
  - Borda choses B (9 pts) ; but A is a CW (A>B 4:3; A>C 4:3)
  - notice that for this preference profile, *any* scoring rule (not just Borda) will choose B if scores strictly decrease with rank

The Axiomatic Method

- Considerable work studies various *"axioms"* or principles that we might like voting rules to satisfy and asks whether we can devise rules that meet these criteria

- For example, the Condorcet principle is an axiom/property we might consider desirable. We've seen some standard voting rules satisfy it, and others do not.

- Let's consider a few more rather intuitive axioms…

Weak Mononoticity

- *Weak monotonicity:* Let *V* be a set of vote profiles and let *V'* be identical to *V* except that one alternative *a* is ranked higher in some of the votes. Then if *a* is the winner under voting rule r with profile *V*, it should also be the winner with profile *V'*.
  - That is, if *a* is the winner under some voting rule given some voter preferences, then a should remain the winner if a few voters raise their ranking of *a*, but everything else is unchanged.

- STV violates weak monotonicity
  - 22 votes: $A \succ B \succ C$
  - 21 votes: $B \succ C \succ A$
  - 20 votes: $C \succ A \succ B$
  - A wins (C, then B eliminated)…
  - … but if anywhere from 2 to 9 voters in the BCA group "promote" A to top of their rankings, C wins (B, then A eliminated)
- Lot of rules satisfy weak monotonicity (e.g. plurality, Borda, …)

Independence of Irrelevant Alternatives (IIA)

- *Independence of Irrelevant Alternatives (IIA):* Suppose *V'* is a vote profile that is different than *V*, but every vote in *V'* gives the same relative ordering to *a, b*, as it does in *V*. Then if *a* is the winner under a voting rule *r* given profile *V*, the *b* cannot be the winner under profile *V'*.
  - In other words, if the votes are changed, but the *relative* (pairwise) preference for *a* and *b* are identical for every voter, then we can't change the winner from *a* to *b.*

- Borda violates IIA (as do quite a few other voting systems):
  - 3 votes: $A \succ B \succ C \succ D \succ E$
  - 1 vote: $C \succ D \succ E \succ B \succ A$ (switch: $C \succ B \succ E \succ D \succ A$)
  - 1 vote: $E \succ C \succ D \succ B \succ A$ (switch: $E \succ C \succ B \succ D \succ A$)
  - C wins under red votes (Borda scores: C:13, A:12, B:11, D:8, E: 6)
  - … but with the blue switches, B wins (scores B:14, C:13, …).
  - Winner from C to B, despite all paired B,C prefs identical in both cases.

Independence of Irrelevant Alternatives (IIA)

- Another view of IIA: suppose *a* wins over *b* in an election. Then we add a new alternative. *Without changing anyone's relative preferences for a and b*, suddenly *b* can win.

- Consider the following preferences:
    - 100 votes: Bush $\succ$ Gore $\succ$ Nader
    - 12 votes: Nader $\succ$ Gore $\succ$ Bush
    - 95 votes: Gore $\succ$ Nader $\succ$ Bush
  - Run a plurality election with only two candidates, Bush and Gore
    - Gore wins over Bush (plurality score of 107 to 100)
  - At the least minute, Nader enters the race:
    - Bush wins the election now (plurality score of 100 to 95 to 12)

Other Principles

- *Unanimity:* if all $v \in V$ rank a first, then a wins
  - relatively uncontroversial
- *Weak Pareto:* if for all $v \in V$ rank $a \succ b$, then b cannot win
  - relatively uncontroversial
  - Implies unanimity
- *Non-dictatorial:* there is no voter k s.t. a is the winner whenever k ranks a first (no matter what other voters say)
- *Anonymity:* permuting votes within a profile doesn't change outcome
  - e.g., if all votes are identical, but provided by "different" voters, result does not change (can't depend on voter's identities)
  - implies non-dictatorship
- *Neutrality:* permuting alternatives in a profile doesn't change outcome
  - i.e., result depends on relative position of an alternative in the votes themselves, not on the identity of the alternative
  - implies non-imposition (i.e. every possible ranking is achievable)

Arrow's Theorem

- So can we satisfy all (or even some of these axioms)?
- Arrow's Theorem (1951): Assume at least three alternatives. No consensus ranking rule can satisfy IIA, unanimity, and non-dictatorship.
  - Most celebrated theorem in social choice
  - Broadly (perhaps too broadly) interpreted as stating there is no good way to aggregate preferences
  - Key point: Arrow's Theorem is phrased in terms of a rule producing a ranking.
- There are a wide variety of alternative proofs
  - Karlin and Peres provide a proof in section 13.7
  - An especially simple proof is given in the next two slides for those who are interested.

Brief Proof Sketch

- A coalition $S \subseteq N$ is *decisive* for a over b if, whenever $a \succ_k b, \forall k \in S$, and $a \nsucc_j b, \forall j \notin S$, we have $a \succ_F b$.
- Fix SWF F; let $\succ_F$ denote social preference order given input profile
- **Lemma 1:** if S is decisive for a over b then, for any c, S is decisive for a over c and c over b.
- **Sketch:** Let S be decisive for a over b.
  - Suppose $a \succ_k b \succ_k c, \forall k \in S$ and $b \succ_j c \succ_j a$, $\forall j \notin S$.
  - Clearly, $a \succ_F b$ by decisiveness.
  - Since $b \succ_j c$ for all j, $b \succ_F c$ (by unanimity), so $a \succ_F c$.
  - If b placed anywhere in ordering of *any agent,* by
  - IIA, we must still have $a \succ_F c$.
  - Hence S is decisive for a over c.
  - Similar argument applies to show S is decisive for c over b.
- **Lemma 2:** If S is decisive for a over b, then it's decisive for every pair of alternatives $(c,d) \in A^2$
- **Sketch:** By Lemma 1, S decides c over b. Reapplying Lemma 1, S decides c over d.

Brief Proof Sketch continued

- So now we know a coalition S is either *decisive* for all pairs or for no pairs.
- Notice that *entire group N is decisive for any pair of outcomes* (by unanimity)
- **Lemma 3:** For any $S \subseteq N$, and any partition (T,U) of S. If S is decisive then either or T is decisive or U is decisive.
- **Sketch:** Let $a \succ_k b \succ_k c$ for $k \in T$; $b \succ_j c \succ_j a$ for $j \in U$;
  $c \succ_q a \succ_q b$ for $q \in N \backslash S$;
  - Social ranking has $b \succ_F c$ since S is decisive.
  - Suppose social ranking has $a \succ_F b$, which implies $a \succ_F c$ (by transitivity).
    - Notice only agents in T rank $a \succ c$, and those in U, $N \backslash S$ rank $c \succ a$.
    - But if we reorder prefs for any other alternatives (keeping $a \succ c$ in T, $c \succ a$ in U and $N \backslash S$), by IIA, we must still have $a \succ_F c$ in this new profile.
    - Hence T is decisive for a over c (hence decisive for all pairs).
  - Suppose social ranking has $b \succ_F a$
    - Since only agents in U rank $b \succ a$, similar argument shows U is decisive.
  - So either T is decisive or U is decisive.

End of proof

- **Proof of Theorem:** Entire group N is decisive. Repeatedly partition, choosing the decisive subgroup at each stage. Eventually we reach a singleton set that is decisive for all pairs… the dictator!

Muller-Satterthwaite Theorem

- Arrow's theorem: impossible to produce a societal ranking satisfying our desired conditions
  - What if we only want a unique winner?
  - Also not possible…

- **Muller-Satterthwaite Theorem (1977):** Assume at least three alternatives. No resolute (one that doesn't produce ties) voting rule can satisfy strong monotonicity, non-imposition (unanimity), and non-dictatorship.

May’s Theorem

- Complete despair? Not really. We could either:
  - dismiss some of the axioms/properties as too stringent
  - live with “general” impossibility, but use rules that tend to (in practice) give desirable results (behavioral social choice)
  - look at restrictions on the assumptions (number of alternatives, all possible preference/vote profiles, …)

- Here’s a positive result (and characterization)…

- **May’s Theorem (1952):** Assume *two* alternatives. Plurality is the only voting rule that satisfies anonymity, neutrality, and positive responsiveness (a slight variant of weak monotonicity).

Manipulation of Elections

- Recall our discussion of mechanism design (e.g., auctions)
  - we needed special mechanisms (e.g., VCG mechanism, 2nd-price auction) to ensure that people would report their valuations truthfully
  - these mechanisms relied on carefully crafted payments
  - in other settings (e.g., 1st-price auction), true valuations are not declared
- In voting (social choice) we don’t usually consider payments such as
  - if we go to your restaurant, you need to pick up the bar tab; or if your candidate wins an election, we increase your property taxes 0.3%
  - aside: it’s worth noting that VCG was motivated in some circles as a means for taxing for public projects (the “Clarke tax”)
- So is it possible for a voter to get a better outcome by misreporting their preferences?

Examples of Manipulability

- Most voting rules provide positive incentive to misreport preferences to get a more desirable outcome
  - political phenomena such as *vote splitting* are just one example
- Plurality:
  - 100 votes: Bush $\succ$ Gore $\succ$ Nader
  - 12 votes: Nader $\succ$ Gore $\succ$ Bush
  - 95 votes: Gore $\succ$ Nader $\succ$ Bush
  - Bush wins truthful plurality vote; Nader supporters are better off voting for Gore! *Notice that Borda, STV would give election to Gore*
- Borda: same example with different numbers
  - 100 votes: Bush $\succ$ Nader $\succ$ Gore
  - 17 votes: Nader $\succ$ Gore $\succ$ Bush
  - 90 votes: Gore $\succ$ Nader $\succ$ Bush
  - Bush wins truthful Borda vote (B:200 pts; G:197pts); Nader supporters better off ranking Gore higher than Nader! Bush supporters were better off ranking Gore last.

Gibbard-Satterthwaite Theorem

- *Strategyproofness (truthfulness) is* defined for voting rules just as for mechanisms
  - Informally, a voting tule is strategy-proof is there no preference profiles where an insincere report by any voter *k* (i.e., reporting something other than his true ranking) leads to outcome that is preferred by *k* to the result obtained from his true report
- Manipulability unavoidable in general (for general SCFs)

- **Thm (Gibbard73, Sattherwaite75):** Let *r* be a voting rule (over voters *N*, alternatives *A*) s.t.:
  - (i) $|A| > 2$;
  - (ii) *r* is *onto* (every outcome is selected for some vote profile *V*);
  - (iii) *r* is non-dictatorial;
  - (iv) all preference profiles (combinations of rankings) are possible.

  Then *r* cannot be strategy-proof.

Are we doomed to possible manipulation?

- Unlike the prevous impossibility theorems, the axioms in the Gibbard Satherthwaite Theorem seem very reasonable.
- But the theorem does imply that all preference profiles are possible which in many applications is not the case.
- Moreover, one of the insights of algorithmic social choice is that while certain voting rules can be manipulated, it may be computationally hard to determine how this manipulation can be done.

Single-peaked Preferences

- Special class of preferences for which GS Theorem is circumvented
- Let $\gg$ denote some "natural" ordering over alternatives A
  - e.g., order political candidates on left-right spectrum
  - e.g., locations of park, warehouse on a line (e.g., position on a highway)

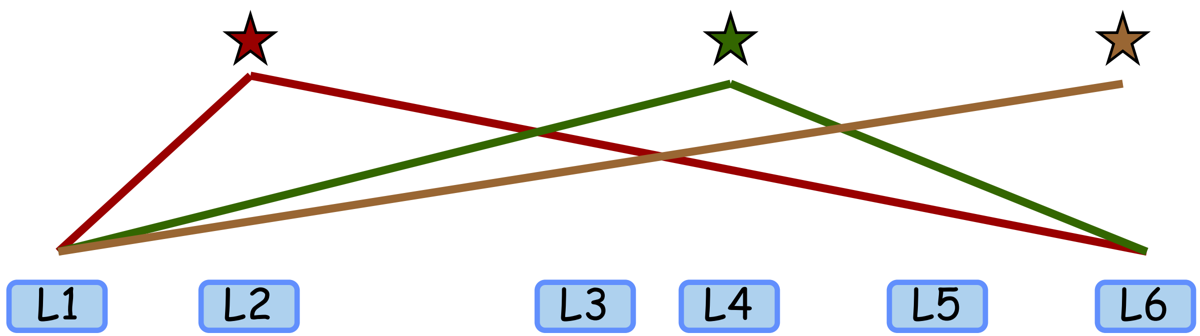

- Voter k's preferences are *single-peaked* if there is an *ideal* alternative, $a^*[k]$, that k likes best, and that as you move away from $a^*[k]$ in the ordering $\gg$, alternatives become less and less preferred by k; that is:
  - $a^*[k] \succ_k a$ for any $a \neq a^*[k]$
  - $b \succ_k c$ if either: (1) $c \gg b \gg a^*[k]$; or (2) $a^*[k] \gg b \gg c$
- In figure: green voter (★) prefers $L4 \succ L3 \succ L2 \succ L1$ and $L4 \succ L5 \succ L6$

Median Voting

- Suppose all voter's prefs are single-peaked
  - they must be single-peaked w.r.t. the same domain ordering ≫
  - but you can use any ordering you want (as long as it creates SP'ed prefs)
- *Median voting scheme:* voter specifies only her peak; winner is *median* of the reported peaks (Black 1948)

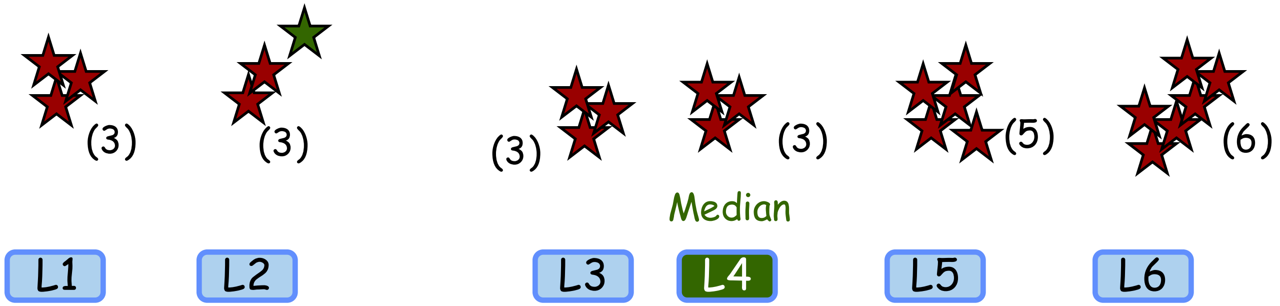

What's Special about Median Voting?

- Assume all voters have single-peaked preferences and we use median voting to determine the winner
- One property: voters don't report full rankings, just peaks (or favorite)
- Another critical property: the voting scheme is *strategyproof*
  - easy to see, let's look at example
  - intuition: if you "lie about your peak" you either report something:
    - on the same side of median as your peak: median unchanged
    - …or on opposite side of median as peak: median moves further away

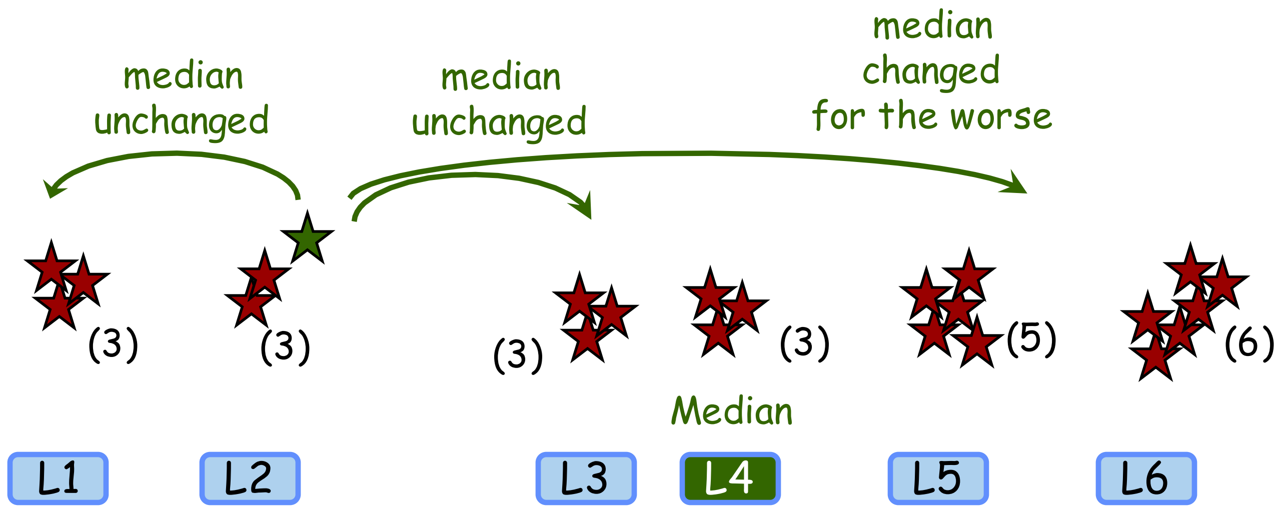

What's Special about Median Voting?

- Assume single-peaked preferences and use median voting
- The winner *W* is Pareto efficient *(in example L4)*
  - *no other choice is better for one person without hurting someone else*
- The winner *W* is a Condorcet winner (if *n* odd): Why?
  - at least *(n+1)/2* voters prefer *W* to anyone *left* of *W* (more if there is more than one voter's peak at the median)
  - at least *(n+1)/2* voters prefer *W* to anyone *right* of *W* (more if there is more than one voter's peak at the median)
  - ***so W wins a majority election against any other candidate***
- Known as the *Median Voter Theorem*

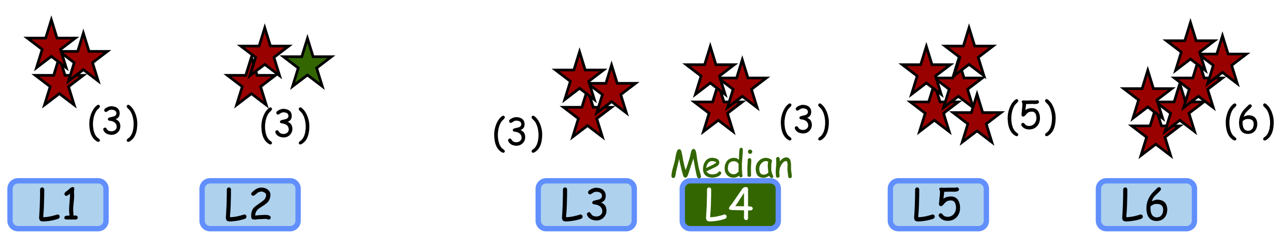

What's Special about Median Voting?

- Can take Median Voter Theorem a step further, imagine following procedure:
  - place *W* at top of societal ranking, then remove it from candidate set
  - repeat process to find median winner among *remaining* candidates
    - there again must be a Condorcet winner (!)
    - *in example: peaks for all voters stays the same except for those who voted for L4 (those voters each have a new peak, either L3 or L5)*
  - remove and repeat until you've ranked all candidates
- Societal ranking must be complete and transitive and respects majoritarian preferences: if A >B in ranking, the majority prefer A to B
  - breaks the Condorcet paradox