2534 Lecture 11: Intro to Social Choice

- Wrap up from last time:
  - briefly: Sandholm and Conitzer’s work on automated mechanism design; Blumrosem, Nisan, Segal: limited communication auctions
  - note: review material on auction design from last week’s slides (we won’t go over in class due to time limitations)
- Intro to Social Choice
- Announcements
  - Make up class next week: Tues, Dec.9, 1-3PM, PT266
  - Assignment 2: marker not quite done (sorry!)
  - Assignment 3 (short): posted today, due Dec.15
  - Projects due on Dec.17

Social Choice

- Social choice
  - more general version of the mechanism design problem
  - assume agents (society, club, …) have preferences over outcomes
  - we have a social choice function that specifies the "right" outcome given the preferences of the population
- Focus is different than mechanism design
  - preferences are usually orderings (qualitative, not quantitative)
  - no monetary transfers considered ("mechanism design w/o money")
  - often focus on design and analysis of aggregation schemes (or "voting rules") that satisfy specific axioms, usually assuming sincere reporting of preferences
  - computational focus: winner determination, approximation, communication complexity, manipulability, …

Social Choice: Basic Setup

- Set of m possible *alternatives (outcomes)* A
- n players
  - each with *preference ordering* $\succ_k$ (or *ranking/vote* v_k) over A
  - assume $\succ_k$ is a *linear order* (no indifference): not a critical assumption
  - let $v = (\succ_1, \ldots, \succ_n)$ denote *preference profile*
  - let L denote the set of linear orderings over A
- Two settings considered
  - A *social choice function (SCF)* $C: L^n \rightarrow A$ (i.e., consensus *winner*)
  - A *social welfare function (SWF)* $C: L^n \rightarrow L$ (i.e., consensus *ranking*)

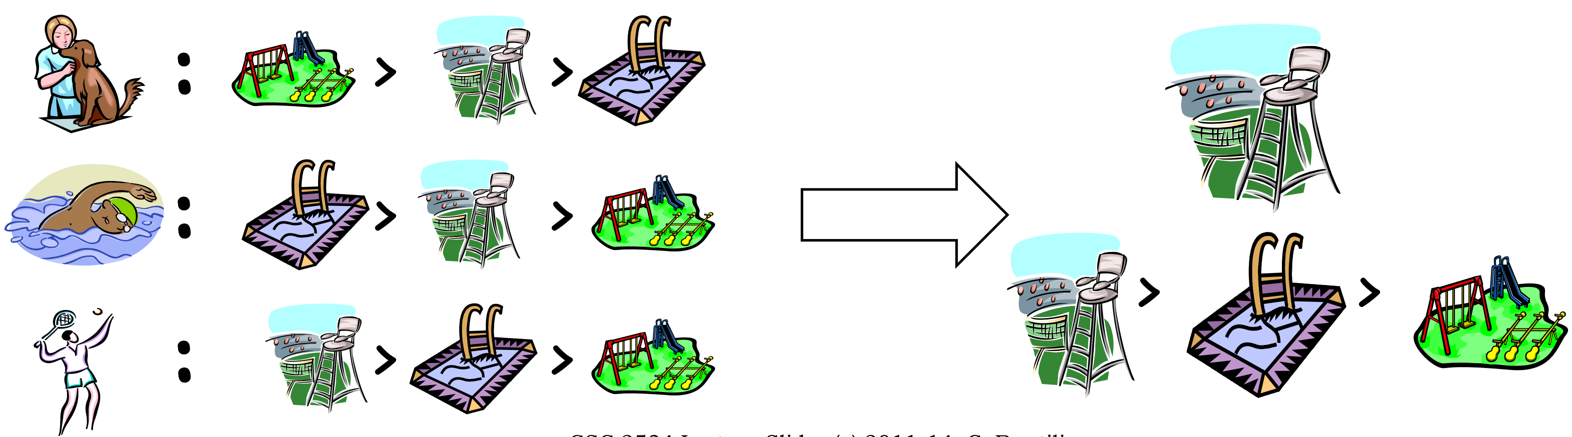

Why Should We Care?

- Computational models/tradeoffs inherently interesting
  - Winner determination, manipulation, approximations, computational/communication complexity
- Decision making/resource allocation in multi-agent systems
- Preference and rank learning in machine learning
  - *Ready availability of preference data* from millions of individuals
  - Web search data, ratings data in recommender systems, …
  - Often implicit; but explicit preferences available at low cost

Voting Rules

- Often SCFs are specified using *voting rules*
  - each player specifies a *vote* (her ranking or some part of it)
  - given vote profile, rule $r: V^n \rightarrow A$ specifies consensus choice
    - distinguish *resolute, irresolute* rules; assume sincere voting
- Three simple rules (with different forms of votes)
  - **plurality vote:** each voter specifies their preferred alternative; winner is candidate with largest number of votes (with some tie-breaking rule)
  - **Borda rule:** each voter specifies ranking; each alternative receives m-1 points for every 1st-place rank, m-2 points for every 2nd-place, etc.; alternative with highest total score wins
  - **approval vote:** each voter specifies a subset of alternatives they “approve of;” a point given for each approval; alternative with highest total score wins (variant: *k*-approval, list exactly *k* candidates)

Notice: each of these can be defined by assigning a *score* to each rank position

Plur:	1	0	0
Borda:	2	1	0
2-Appr:	1	1	0

How do they differ?

- Example preference profile (3 alternatives, bold=approval):
  - **A** ≻ B ≻ C: 5 voters (approve of only top alternative)
  - **C** ≻ B ≻ A: 4 voters (approve of only top alternative)
  - **B** ≻ **C** ≻ A: 2 voters (approve of top two alternatives)
- Winners:
  - plurality: A wins (5 votes)
  - Borda: B wins (scores B: 13; A: 10; C: 10)
  - approval: C wins (scores C: 6; A: 5; B: 2)
- Which is voting rule is better?
  - 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!

Some Voting Systems/Rules

- *Plurality, Borda, k-approval, k-veto*
  - all implementable with *scoring rules*: assign *score* α to each rank position; winner *a* with max total: $\sum_i \alpha(v_i(a))$
  - for two candidates, plurality sometimes called *majority* voting
- *Approval*
  - can't predict how sincere voters will vote based on ranking alone
- *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 > *m/2*
  - Needn't be online: voters can submit rankings once
  - used in Australia, New Zealand, Ireland, …

Small Sampling of Voting Systems/Rules

- *Egalitarian (maxmin fairness)*
  - Winner maximizes min rank: $argmax_a \, min_j \, (m\text{-}v_j\,(a))$
- *Copeland*
  - Let $W(a,b,\mathbf{v}) = 1$ if more voters rank $a \succ b$; 0 if more $b \succ a$; ½ if tied
  - Score $s_c(a,\mathbf{v}) = \sum_{b \neq a} W(a,b,v)$; winner is a with max score
    - *i.e., winner is candidate that wins most pairwise elections*
- *Nanson's rule*
  - Just like STV, but use Borda score to eliminate candidates
- *Tournament/Cup*
  - Arrange a (balanced) tournament tree of pairwise contests
  - Winner is last surviving candidate
- Lots of others!!!

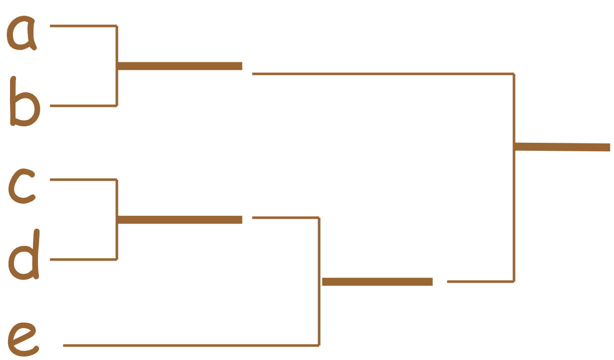

Condorcet Principle

- *Condorcet winner (CW):* an alternative that beats any other in a pairwise majority vote
  - if it exists, must be unique
  - a rule is *Condorcet-consistent* if it selects the Condorcet winner whenever one exists
- *Condorcet paradox:* CW may not exist
  - and pairwise majority preferences may induce cycles in “societal ranking”
  - $A \succ B \succ C$: *m/3* voters
  - $C \succ A \succ B$: *m/3* voters
  - $B \succ C \succ A$: *m/3* voters

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 *any* scoring rule (not just Borda) will choose B if scores strictly decrease with rank
- Nanson, Copeland, Kemeny^{*tba} rules are Condorcet consistent

Consensus Rankings

- May wish to determine a *societal preference order*
  - notice*: any rule that scores candidates* can determine a societal ranking
- Another important rule: *Kemeny rule*
  - Distance measure between rankings—Kendall's τ

$$\tau(r,v) = \sum_{\{c,c'\}} I[r(c) > r(c') \text{ and } v(c') > v(c)]$$

  - *Kemeny ranking $\kappa(V)$:* minimizes sum of distances

$$\kappa(V) = \min_{r} \kappa(r,V); \quad \kappa(r,V) = \sum_{\ell=1}^{n} \tau(r, v_\ell)$$

- Can determine winner too: top of Kemeny ranking
  - Condorcet consistent
  - Example of a voting rule that is hard to compute: NP-hard
  - Other difficult rules include Dodgson's rule, Slater's rule

Also co-inventor of BASIC

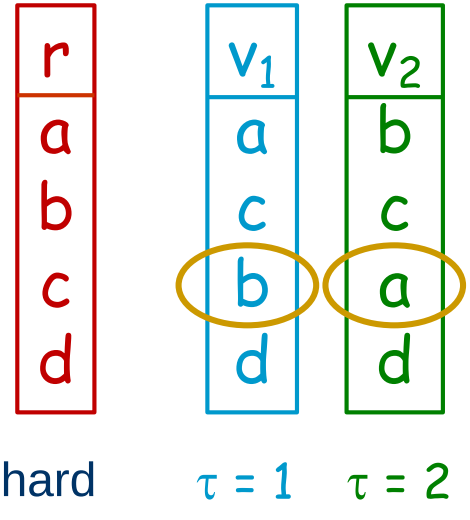

Other Principles

- *Weak monotonicity:* Let profile *V' be identical* to *V* except that some candidate *a* is ranked higher in some *votes.* Then:
  - *Rule:* If $a \in r(V)$ then $a \in r(V')$;
  - *Ranking:* If $a \succ b$ in $r(V)$ then $a \succ b$ in $r(V')$;
  - 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 2-9 voters in BCA group "promote" A to top of ranking, C wins (B, then A eliminated)
  - Lot of rules satisfy it (plurality, Borda, …)

Other Principles

- *Strong monotonicity:* Let $a = r(V)$. Let V' be s.t. for every k, every $b \neq a$, if $a \succ b$ in v_k, *then* $a \succ b$ in v'_k. Then $a = r(V')$.
  - i.e., if no voter "demotes" *a* relative to any other candidate, *a* still wins
  - unlike WeakMon, can reorder non-winning candidates w.r.t. each other
  - Plurality (and many others) violate SM
    - 22 votes: $A \succ B \succ C$
    - 21 votes: $B \succ C \succ A$
    - 20 votes: $C \succ A \succ B$
    - A wins; but if 3 or more BCA voters "promote" C, C wins (even though relative standing of A to B, C unchanged by any voter)

Other Principles

- *Independence of Irrelevant Alternatives (IIA):* *V'* different from *V*, but relative ordering of *a, b*, same in each vote
  - *Rule*: If $a \in r(V)$, $b \notin r(V)$, then $b \notin r(V')$;
    - i.e., if *b* wasn't strong enough to beat *a* given *V*, it shouldn't be given *V'*
  - *Rank*: if $a \succ b$ *in* $r(V)$ then $a \succ b$ *in* $r(V')$;
  - Most rules violate IIA: easy to construct examples

Other Principles (Relatively Uncontroversial)

- In what follows, assume all preference/vote profiles are possible

- *Unanimity:* if all $v \in V$ rank a first, $r(V)=a$; if all rank $a \succ b$, then $a \succ b$ *in* $r(V)$
  - relatively uncontroversial (sometimes called weak Pareto)
- *Weak Pareto:* if all $v \in V$ rank $a \succ b$, then $b \notin r(V)$
  - relatively uncontroversial
- *Non-dictatorial:* there is no voter k s.t. $r(V)=a$ *whenever* k ranks a first
  - for rankings, no k s.t. $a \succ b$ *in* $r(V)$ whenever k ranks $a \succ b$
- *Anonymity:* permuting votes within a profile doesn't change outcome
  - e.g., if all votes identical, but provided by "different" voters
  - implies non-dictatorship
- *Neutrality:* permuting alternatives in a profile doesn't change outcome
  - i.e., result depends on relative position in votes, not identity
  - implies non-imposition (any candidate *can* win, i.e., for *some* profile)

Arrow’s Theorem

- **Arrow’s Theorem (1951):** Assume at least three alternatives. No voting rule can satisfy IIA, unanimity (weak Pareto), and non-dictatorship. Equivalently, there is no SWF that satisfies these properties.
  - *(Recall SWF produces “societal ranking,” not just a winner; c.f. SCF)*
  - Most celebrated theorem in social choice
  - Broadly (perhaps too broadly) interpreted as stating there is no good way to aggregate preferences
- There are a wide variety of alternative proofs around
  - see text for one
  - we’ll consider a simple proof

Brief Proof Sketch

- Fix SWF F; let $\succ_F$ denote social preference order given input profile
- 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 $a \succ_F b$.
- **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.

$S: a \succ b \succ c$
$-S: b \succ c \succ a$ ⟹ $F: a \succ b \succ c$

$S: \mathbf{b} \succ a \succ c$
$-S: c \succ a \succ \mathbf{b}$ ⟹ $F: a \succ c$

- **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

- 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 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.
- **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 tells us: impossible to produce a *societal ranking* satisfying our desired conditions (in a fully general way)
  - Maybe producing a full ranking is too much to ask
  - What if we only want a unique winner?
  - Also not possible…
- **Muller-Satterthwaite Theorem (1977):** Assume at least three alternatives. No resolute voting rule satisfies strong monotonicity, non-imposition, and non-dictatorship. Equivalently, there is no SCF that satisfies these properties.

May’s Theorem

- Should Arrow’s Thm cause complete despair? Not really…
  - dismiss some of the desiderata 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 (which is *majority* in case of two alternatives) is the only voting rule that satisfies anonymity, neutrality, and positive responsiveness (a slight variant of weak monotonicity).

- Social choice has a variety of interesting (and not so interesting) characterizations of this type (we’ll see some more)

Manipulability

- As with mechanism design, 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 sincere plurality vote; in the interest of Nader supporters to vote for Gore. *Notice that Borda, STV would give election to Gore*
- Borda: same example with different numbers
  - 100 votes: Bush $\succ$ Gore $\succ$ Nader
  - 17 votes: Nader $\succ$ Gore $\succ$ Bush
  - 90 votes: Gore $\succ$ Nader $\succ$ Bush
  - Bush wins sincere Borda vote (B:200 pts; G:197pts); in the interest of Nader supporters to rank Gore higher than Nader

Manipulability

- *Strategyproofness* defined for voting procedures just as it is for mechanisms
  - no profiles where insincere report by *k* leads to preferred outcome for *k*
    - *strategyproof: dominant strategy truthful*
    - *incentive compatible: truthful in (voting) equlibrium (e.g., Bayes-Nash)*
- Alternatively, we can define SCFs themselves as being strategyproof
  - there is no profile, agent *k* s.t. $C(\succ_1, \ldots \succ'_k, \ldots \succ_n) \succ_k C(\succ_1, \ldots \succ_k, \ldots \succ_n)$
- Manipulability unavoidable in general (for general SCFs)
  - already seen our old friend GS in the context of mechanism design
- **Thm (Gibbard73, Sattherwaite75):** Let *C* (over *N, O*) be s.t.:
  - (i) $|O| > 2$;
  - (ii) *C* is onto (every outcome is selected for some profile *v*;
  - (iii) *C* is non-dictatorial;
  - (iv) all preference profiles L^n are possible.

  Then *C* cannot be strategy-proof.

Single-peaked Preferences

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

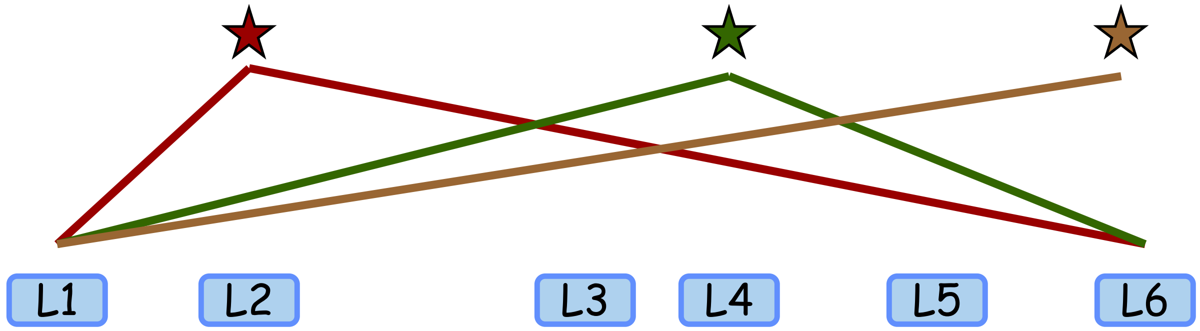

- *k*'s preferences are *single-peaked* (with respect to the *given* ordering of A) if there is alternative *a*[k]* s.t.:
  - *a*[k]* is *k*'s ideal point, i.e., $a^*[k] \succ_k a$ for any $a \neq a^*[k]$
  - $b \succ_k c$ if (a) $c \gg b \gg a^*[k]$ or (b) $a^*[k] \gg b \gg c$

Median Voting

- Suppose all voter's prefs are single-peaked (same domain order!)
- *Median voting scheme:* voter specifies only her peak; winner is median of reported peaks (Black 1948)
  - result is a Condorcet winner (if *n* odd)
  - result is Pareto efficient
  - voting scheme is *strategyproof* (easy to see)

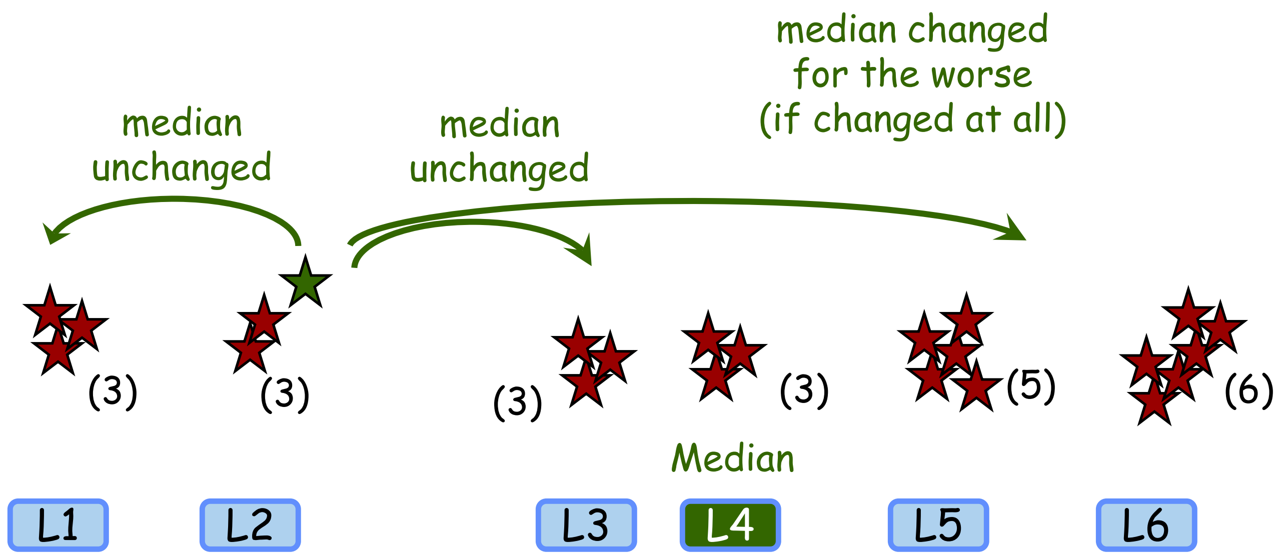

Generalized Median Voting

- Suppose we add *n-1* "phantom voters" with arbitrary peaks
  - announced in advance, chosen for "some purpose"
- Winner is median of the *2n-1* total votes (*n* real, *n-1* phantom)
  - e.g., in example, the phantom votes implement selection of 33rd percentile (or 1/3 quantile) among true peaks
- *Generalized Median:* if preferences are single-peaked, any anonymous, Pareto efficient, strategyproof rule must be a generalized median mechanism (Moulin 1980)
  - some mild generalizations (e.g., multiple dimensions) possible
  - *Recent work:* can you find an axis/axes that render profile *V* SP?
  - … are there natural approximations of SP? how does it impact incentives?

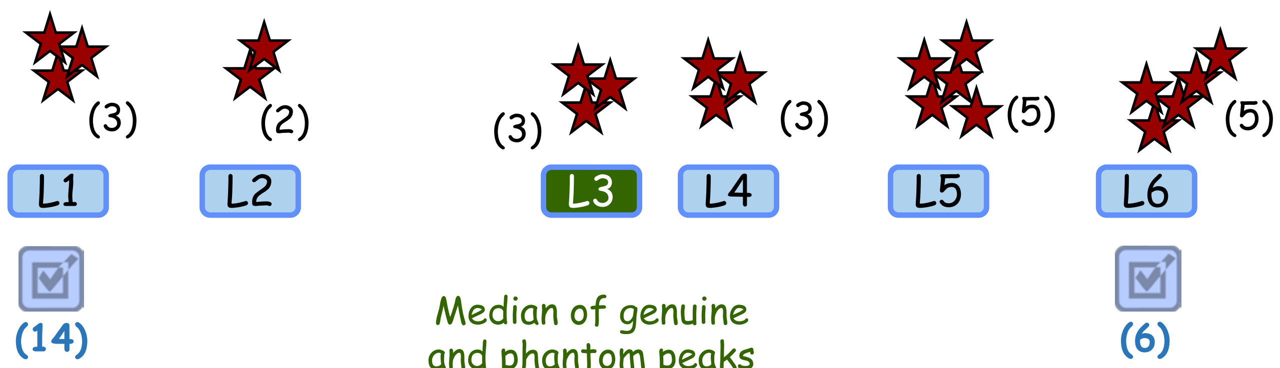

Complexity as Barrier to Manipulation

- Topic of considerable study in CS
  - started with seminal work of Bartholdi, Tovey, Trick (1989, 1991)
  - widely ignored for many years, now well-studied
- Given *n-1* votes, desired candidate *a**: can n^{th} voter ensure *a** wins?
  - *constructive* manipulation; also *destructive* variant (prevent winner)
  - can also consider manipulating *coalitions* (and size needed)
- Decision problem is tractable for some rules
  - plurality: easy, if manipulable, it is accomplished by voting for a*
  - Borda: easy (for single voter): place *a** at top of ballot, greedily add candidates in next positions so they don't "overtake" *a** (if not possible, not manipulable)
- Intractable for others:
  - STV: determining (constructing) manipulating vote NP-hard (BTT91)
  - many voting rules subsequently analyzed this way
- Analysis more nuanced for coalitions, weighted voters, etc.

Complexity as Barrier to Manipulation

- These results should be taken with a grain of salt
  - worst-case manipulation: *some* vote profiles are hard to handle; but doesn't mean *typical* case is (and that's *crucial* for "resistance" claims)
    - increasing work on empirical analysis and avg. case behavior
  - assumptions are beneficial to manipulators: know votes cast by others!
    - hence a conclusion of manipulability under this model may not be very meaningful (too pessimistic, unrealistic)
    - further analysis needed with realistic knowledge constraints (min entropy, sample complexity, etc.)
- Other forms of manipulation
  - control: adding, deleting candidates; setting agenda (tournament); setting up electoral "boundaries" or groups (gerrymandering); …
  - bribery: pay someone to change their vote

Example: Control of Tournament (Cup Rule)

- Set a *balanced* binary tree of pairwise contests
- Person setting the agenda can sometimes choose whichever winner they want (if they know the votes)
  - 35 votes: $A \succ C \succ B$
  - 33 votes: $B \succ A \succ C$
  - 32 votes: $C \succ B \succ A$
  - If *(a,b)* paired first, *c* wins; If *(b,c)* first, *a* wins; If *(a,c)* first, *b* wins
- Complexity of determining if a (dynamic) schedule can make *a* win:
  - known votes: still unknown if polynomial!
  - probabilistic votes: NP-hard (even for $v \in \{0, ½, 1\}$)
- Other interesting questions in this space (esp. for sports, etc):
  - throwing matches, maximizing competitiveness/revenue, etc.

“Complexity” as a barrier to manipulation

- The Doge of Venice:
  - chief magistrate of the Most Serene Republic of Venice c.700-1797
  - elected for life by the city-state's aristocracy
  - concern about the influence of powerful families!
- Voting Protocol in 15th Century *(courtesy Wikipedia via Mike Trick ADT-09)*
  - 30 members of the Great Council are chosen by lot
  - The 30 are reduced by lot to 9
  - The 9 choose 40 representatives
  - The 40 are reduced by lot to 12
  - The 12 choose 20 representatives
  - The 20 twenty are reduced by lot to 9
  - The nine elect 45 representatives
  - The 45 are reduced by lot to 11
  - The 11 choose 41 representatives
  - These 41 actually elect the doge

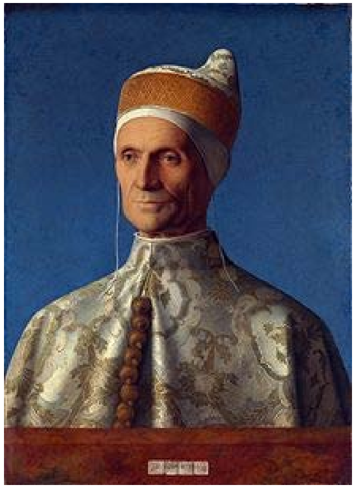

Objective Rankings

- A different perspective: rankings as beliefs (not preferences)
  - suppose there is a true underlying *objective* ranking r^*
    - *e.g., quality of sports teams, ability to lead a nation, impact of policy P on economy, relevance of document/web page to a query, …*
  - agents have opinions on the matter: correlated (noisily) with obj. r*
- Rank aggregation aimed at ascertaining true r^*, not some SCF
- Condorect addressed this in 1785:
  - Suppose *n* voters (e.g., jury) vote on two alternatives (e.g., guilt/innocence). If each votes independently and is correct with $p>½$, then plurality rule gives maximum likelihood estimate of correct alternative, and converges to correct decision as $n \rightarrow \infty$.
  - Young (1995) generalized: if each voter noisily ranks arbitrary pairs (a,b) correctly with probability $p>½$, the *Kemeny consensus* is a maximum likelihood estimate of the true underlying ranking.
  - See Conitzer, Sandholm (2005) for treatment of several other rules (e.g., Borda) using specific noise models tuned to that rule

Other Issues

- Multi-winner elections
  - proportional assemblies, committees, multiple projects, etc.
  - diversity a key consideration: “first *k* past the post” usually a bad idea
- Behavioral social choice
  - designing, analyzing rules based on empirical preferences
  - modeling preference distributions (econometrics, psychometrics)
- Combinatorial preference aggregation
  - preferences over complex domains (multi-issue)
  - appropriate preference rep’ns, aggregation methods, algorithms
- Communication complexity, privacy concerns (*à la* mech. design)
- Preference Elicitation
  - ballot complexity a barrier to wider-spread use of rank-based voting
- Approximation of Social Choice Functions
  - does ability to approximate winner ease burden:
    - communication? computation? privacy?