# Voting Protocols and Arrow's Theorem

Subhash Suri

March 4, 2020

## 1 Strategic Voting

- Voting is about *aggregating* individual wishes (preferences) into a collective *societal preference*. A central decision question in voting is *which* collective choice or outcome is the most logical one when the individuals of the society have conflicting preferences.
- Many natural examples come to mind when we think of voting, such as choosing a president in national elections, hiring a faculty member, or choosing community projects (swimming pool, ice rink, library, etc) to fund under limited budget.
- Voting also comes up quite naturally in many computer science contexts.
  1. **Rank Aggregation.** Consider a prediction problem where the goal is to create a *ranked list*, such as the Web pages in response to a search query. We may have several different heuristics, for instance, one based on anchor text, another on link structure, and yet another on page content, which may produce different results for the same search query. It is natural to try to combine these different and *possibly conflicting* rankings into a "consensus ranking." This is precisely the voting problem: each heuristic is a "voter". Our hope is that the consensus ranking is somehow better (more accurate) than any of the individual ranking. For example perhaps the best pages are high on all lists, while the "spam page" is high on only a few of the lists and thus will be get pushed down in the consensus ranking.
  2. **Crowd Sourcing.** Systems such as Amazon's Mechanical Turk or MOOC education platforms use "workers or graders" to solve ranking problems: e.g., rank order several candidate user interfaces. Again, it is hoped that the "voting" will lead to an overall better (wisdom of the crowds) outcome.
  3. **Participatory Democracy.** The goal is to increase participation in government decisions, especially at the local level. Voters weigh in on budgetary decisions and tradeoffs: spend money on improving parks, schools, or public housing.

Technology enables new types of voting. The current system typically use $k$-*approval voting*–each voter is told the overall budget (e.g., $1 million) and a list of projects with cost estimates. Each voter picks their $k$ favorite projects, with no ordering between them. Projects are then sorted in decreasing order of the number of votes received, and funded in this order until the budget runs out.

A drawback with this system is that voters need not take into account projects' costs—the $k$ projects' combined costs are allowed to exceed the budget), which results in more expensive projects being overrepresented. A new voting prototype is "knapsack voting," where a voter is allowed to approve any number of projects, as long as their total cost is at most the budget. This forces voters to account for project costs, in that voting for more expensive projects decreases the number of projects that you can vote for. It's hard to imagine implementing knapsack voting with a paper ballot at a polling station, and a computerized platform for large-scale voting is essential for its viability.

# 2 Modeling the Problem

- The point of a voting protocol is to make a decision as a function of what people want (their preferences). We need a formal model.

  1. A set of $N$ voters (agents, individuals) whose preferences are to be aggregated.
  2. A set of $n$ Candidates (outcomes, alternatives).
  3. An individual's preference is a rank ordering of all the candidates (ties allowed).
  4. The **goal** is to find a *social ordering* of the candidates that is *consistent* with the individual preferences.

- A voting rule can then have two different forms, depending on whether its output is a full ranked list or just a single alternative (the "winner"). The former is what we want in the rank aggregation application, and perhaps also in the crowdsourcing application. The latter suffices for political elections. We will focus on the general problem, of producing a full rank order.

# 3 Some Well-Known Voting Methods.

1. **Majority Rule Voting.** The voting problem is fairly straightforward *if there are only 2 candidates*—majority voting works. (Elect the alternative that appears first in the largest number of voter's lists.) It's hard to think of any other voting rule that is better when there are two alternatives.

2. The problem, however, becomes surprisingly more complex and subtle with 3 or more candidates, and that is the main subject of this lecture. We list some of the most well-known voting rules below.

3. **Plurality:** the most preferred candidate (ranked #1 by most voters) wins.

4. **Plurality with Runoff.** (Used in Presidential Elections in France.) In first round, each voter casts a single vote. Eliminate all but the top 2, then choose between those using Plurality.

5. **Single Transferable Vote (Instant Runoff).** (Used in UK.) Each agent votes for their most-preferred candidate. If there is no majority, the candidate with fewest votes is eliminated. Each voter who voted for the eliminated candidate transfers their vote to their most-preferred candidate among the remaining candidates. Repeat until a winner.

6. **Approval Voting.** (Used by Mathematical Association of America, etc.) Each voter can cast a single vote for as many candidates as he wants. The candidate with the most votes is selected.

7. **Borda Count:** Each voter rank orders all $n$ candidates. The most preferred candidate receives $n-1$ points, the next $n-2$ points, and the last 0. Add up these voter points for each candidate. The one with the maximum number of points wins.

8. **Pairwise Elimination:** Pair candidates with a *schedule*. The candidate who is preferred by a minority of voters is deleted Repeat until only one candidate is left

   (Example: suppose the order is $A, B, C$. Then, first we do $A$ vs $B$. Eliminate the loser. The winner goes against $C$.)

# 4 Counter-intuitive Behavior of Voting Protocols.

- We first note that different voting protocols may produce very different outcomes *on the same input*, and also exhibit counter-intuitive behavior.
- We will use the following example (preference lists)..

| | | |
|---|---|---|
| 499 | prefer | $A > B > C$ |
| 3 | prefer | $B > C > A$ |
| 498 | prefer | $C > B > A$ |

1. Who is the pairwise winner? It's $B$ because

   - $B$'s score is $1003 = (3 + 498)$ vs $A$ and $(499 + 3)$ vs $C$
   - $A$'s score is $998 = 499$ vs $B$ + 499 vs $C$
   - $C$'s score is $999 = 3$ vs $A$ and $(498 + 498)$ vs $A$ and $B$

2. Who is the Plurality Vote winner? It's $A$!

3. Who wins under Plurality with elimination (or single transferable)? It's $C$!
   ($B$ gets the least #1 votes in round 1 and is eliminated, leaving $A$ vs $C$ election, which $C$ wins because $B$'s supporters switch to $C$.)

4. Another example to show Non-Robustness and Sensitivity

| | | |
|---|---|---|
| 35 | prefer | $A > C > B$ |
| 33 | prefer | $B > A > C$ |
| 32 | prefer | $C > B > A$ |

   (a) The Plurality winner is $A$. The Borda winner also is $A$. Calculation:

| $A$ | $B$ | $C$ |
|---|---|---|
| 70 | 0 | 35 |
| 33 | 66 | 0 |
| 0 | 32 | 0 |

   (b) Now suppose $C$ drops out. How does that affect the outcome? The new preferences

| | | |
|---|---|---|
| 35 | prefer | $A > B$ |
| 33 | prefer | $B > A$ |
| 32 | prefer | $B > A$ |

   So, now under both Borda and Plurality, $B$ wins!

   (c) Who wins the pairwise elimination, with the order $(A, B, C)$? $C$ wins!
   ($A$ loses against $B$. $B$ loses against $C$.)

(d) Who wins the pairwise elimination, with the order $A, C, B$? $B$ wins!
($A$ wins against $C$. $A$ loses to $B$.)

(e) Who wins the pairwise elimination, with the order $B, C, A$? $A$ wins!!

5. **Another Example:**

| | | |
|---|---|---|
| 1 | prefer | $B > D > C > A$ |
| 1 | prefer | $A > B > D > C$ |
| 1 | prefer | $C > A > B > D$ |

Pairwise Election with ordering: $A, B, C, D$.

Who wins? $D$!

Notice something strange. Every single voter prefers $B$ to $D$, yet $D$ wins.

## 4.1 Extended Condorcet Criterion and Consensus Ranking

The XCC condition is the following:

> If $X, Y$ is a partition of the set of candidates such that *every* $x \in X$ *beats every* $y \in Y$ in pairwise order, then all of $X$ must precede all of $Y$ in the social order.

For the problematic cycle order of Condorcet, $XCC$ does not apply, and *any arbitrary ordering of candidates* will do. For this relaxed criterion, one can show the following nice property.

**Theorem.** *For any input ranking, there is a consensus ordering satisfying* $XCC$.

**Proof.** Consider a directed graph $G$ whose vertices are the candidates $V$, with an edge from $x$ to $y$ if $x$ beats $y$ in pairwise comparison. This is a *tournament* graph, and therefore admits a Hamiltonian path (see the theorem below). We rank the candidates in the order specified by this path, and claim that it satisfies $XCC$.

For the sake of contradiction, suppose the ranking fails $XCC$. Let $(X, Y)$ be a partition for which some $y \in Y$ comes before some $x \in X$ even though all $x \in X$ beat all $y \in Y$. There must be an *adjacent pair* $(y, x)$ in the Hamiltonian ordering, with $y$ coming just before $x$. But then $(y, x)$ must be an edge in $G$, contradicting the assumption about the $XCC$ partition $(X, Y)$ that all $x$ beat all $y$.

**Theorem.** *Every tournament graph* $G$ *contains a Hamiltonian path.*

**Proof.** By induction on the number of vertices $n$. For $n = 1$, the claim is trivial. For $n \geq 2$, let $x \in V$ be an arbitrary vertex, and consider the Hamiltonian path $P = (x_1, x_2, \ldots, x_{n-1})$ in the graph $G \backslash x$. If either $(x, x_1)$ or $(x_{n-1}, x)$ is the edge in $G$, we can extend $P$ by inserting $x$ at the beginning or the end.

Otherwise, the edges from $x$ are directed away from $x_1$ and into $x_{n-1}$). Then, in the sequence of edges incident to the vertices $x_1, x_2, \ldots, x_{n-1}$, there must be an index $k$ such that $(x_k, x)$ is directed into $x$ but $(x, x_{k+1})$ is directed out of it. In that case, we insert $x$ into $P$ between $x_k$ and $x_{k+1}$ obtaining the Hamiltonian path for $G$.

# 5 Arrow's Axiomatic Approach.

Given the counter-intuitive behavior (occasional unnatural outcomes) of voting protocols, it is reasonable to ask: *is there is a protocol that does not exhibit these troubling behaviors?* Against this backdrop, we arrive at a famous result in economics and social choice theory, called **Arrow's Impossibility Theorem.**

Kenneth Arrow, who won the Nobel Prize in Economics, investigated an axiomatic approach to rank-order voting protocols, and set out the following **3 Fairness Conditions**, which *every useful* protocol must satisfy:

1. If every voter prefer $X$ to $Y$, then Social choice should also prefer $X$ to $Y$.
2. If every voter's preference between $X$ and $Y$ is unchanged, then the social order between $X$ and $Y$ should remain unchanged as well. In other words, voter's preference changes over other candidates, $U, V, Z$ etc, should not affect $X$ vs $Y$ preference.
3. There is no **dictator**: a single voter whose preferences always determines the social preference.

**Remarks on Conditions:**

- Condition 1 is also called **Pareto Efficiency** (PE): when all agents agree on the rank order of two candidates, then the society should respect that order. If everyone has $A > B$, then social order must also rank $A > B$.
- Condition 2 is called **Independence of Irrelevant Alternatives** (IIA): as long as all voter's relative ranking of $X$ and $Y$ remains the same, their relative order (who comes first) should also remain the same in the social choice.

- The dictatorship conditions rules out the presence of a voter who can dictate the results, meaning that there is no single voter whose preferred order is adopted by the society **no matter** what the preferences of the remaining voters.

**Arrow's Impossibility Theorem.** *No protocol (social choice function) can satisfy all 3 conditions: If we have PE and IIA, then the only protocol is dictatorial!*

This is both a beautiful and a highly unexpected result.

## 5.1 Proof of the Theorem.

- A **profile** is the rank ordering of the candidates by all voters (one preference list per voter).
- Intuitively, the proof uses the fact that the axioms of PE and IIA constrain how the social preference order can change as individual voter preferences change.
- We will carefully construct a set of profiles, with a very special structure, to bring out the logical inconsistency of our 3 fairness axioms.
- The proof exploits the fact that the protocol must produce outcomes that are consistent with our 3 axioms *for all possible voter profiles.* Starting with a highly structured profile whose outcome is indisputable, we will generate other profiles whose outcomes under the constraints of the axioms lead to contradictions.

We begin with the following Lemma that uses Pareto Efficiency.

**Pareto Lemma.** *Pick any candidate $b$. If every voter ranks $b$ either at the very top, or at the very bottom, then social order $>$ must put $b$ either at the very top or the very bottom. (No middle place.)*

**Proof of Pareto Lemma.** Consider voter profiles satisfying the lemma for which $b$ is not extreme, so there are candidates both above and below in the rank order. Let $a, c$ be such that $a > b$ and $b > c$.

We now slightly modify each voter's profile by moving $c$ just above $a$, leaving the rest unchanged. Since in each voter's list, $b$ is at the top or the bottom, this switch does not affect the $(a, b)$ or the $(b, c)$ order. Therefore, by IIA, the social order $>$ must continue to maintain $a > b$ and $b > c$, and by transitivity $a > b > c$.

But each voter now has $c > a$, so by Pareto Efficiency, the social order must also have $c > a$. A contradiction!

**Pivotal Voter Lemma.** *Pick any candidate $b$. There is voter $n^*$ that is pivotal for $b$. That is, while the other candidates keep their ranking of $b$, the agent $n^*$ can move $b$ from bottom to top by changing his vote.*

**Proof of Pivotal Voter Lemma.**

1. List the voters in some arbitrary order $1, 2, ...$
2. Start with a profile in which every voter puts $b$ at the bottom; remaining preferences being arbitrary.
3. By PE, the social order must also rank $b$ at the bottom.
4. Now let the voters $1, 2, ...$ one by one modify their profile by moving $b$ to the top, without changing any other ranking.
5. There must be one earliest voter $n^*$ whose switch causes $b$ to move to the top in the social ranking—because when all voters move $b$ to the top, social order will also have $b$ at the top, so the change must occur at some point.
6. Define $P_1$ as the profile in which the first $(n^* - 1)$ voters have $b$ at the top; rest at the bottom.
7. Define $P_2$ as the profile in which first $n^*$ voters have $b$ at the top. (This is the tipping profile.)
8. Social order for $P_1$ has $b$ at the bottom, but for $P_2$ has $b$ at the top.
9. Agent $n^*$ is the pivotal voter for $b$, and this completes the proof.

**Lemma:** *$n^*$ is pivotal for all candidate pairs $(a, c)$ not involving $b$.*

**Proof.** Without loss of generality, let's choose $a$ as candidate to consider. We construct a new profile $P_3$ by making small modifications to $P_2$, as follows.

1. First, move $a$ to the top of $n^*$'s ranking, without changing anything else. Thus, in $n^*$ ranking we will have $a > b > c$.
2. Second, let the remaining voters arbitrarily rearrange their relative rankings of $(a, c)$, **but leaving $b$ in its extremal position**. (That is, we want to argue no matter what their relative preference is, $n^*$ will dictate.)
3. This is our profile $P_3$.

Now let's see how the social ordering changes.

1. In $P_1$, $b$ was at the bottom of social order, and so $a > b$.
2. In $P_3$, the relative ordering of $a, b$ are the same for all voters as in $P_1$. Therefore, by IIA, we must also have $a > b$ under $P_3$.
3. In $P_2$, $b$ was at the top of the social order, and so $b > c$.
4. In $P_3$, the relative ordering of $b, c$ are the same for all voters as in $P_2$. Therefore, by IIA, we must also have $b > c$.

Taken together and by transitivity, $P_3$ must have $a > c$ as social order. Thus, $n^*$ has the power to make an arbitrary candidate a more preferred than $c$ **irrespective** of the ranking of everyone else. This is the dictator, and the proof of the lemma is complete.

**Unique Dictator Lemma.** *$n^*$ is also pivotal for any pair $(a, b)$.*

**Proof.** The earlier discussion centered around a particular candidate $b$, but we could chosen any candidate to play the role of $b$. So, if we want to prove the pivotal role of $n^*$ on $b$, let's consider some other arbitrary candidate $c$.

By the same reasoning as before, we must have a pivotal voter $n^{**}$ for $c$, who has the power to control the ranking for any pair $(x, y)$ not involving $c$. But $(a, b)$ is such a pair.

On the other hand, we also know that the old pivot $n^*$ is able to affect the $a, b$ ranking—for instance, it was able to change the $a > b$ ranking in profile $P_1$ to $b > a$ ranking in profile $P_2$. This is precisely how $n^*$ was identified as pivotal. So, $n^{**}$ and $n^*$ must be the same–there cannot be two dictators!

**Remark:** With $b$ as a fixed candidate, the pivot $n^*$ provide a vantage point of the preference space from $n^*$ perspective. It has the ability to dictate the rank order of $(a, c)$ for any pair.

$n^{**}$ has a different vantage point from $c$, and it has the power to dictate the relative order of any pair such as $(a, b)$. But we know that $n^{**}$ cannot make $a > b$ under the profile $P_2$. So, it must be the case that $n^{**} = n^*$.

Thus, if social choice function obeys PE and IIA, it is dictatorial.

## 5.2 Interpretation and Consequences

- Non-mathematical statements of the theorem such as *No voting method is fair*, or *The only voting method that isn't flawed is a dictatorship*, are simplifications. What the theorem does say is that a deterministic preferential voting mechanism (where a preference order is the only information in a vote, and any possible set of votes gives a unique result) cannot comply with all of the conditions given above simultaneously.

- Many suggest weakening the IIA criterion as a way out of the paradox. For instance, some contend IIA is an unreasonably strong criterion. It is the one breached in most useful voting systems (trivially implied by the possibility of cyclic preferences).
- So, what Arrow's theorem really shows is that any majority-wins voting system is a non-trivial game, and that game theory should be used to predict the outcome of most voting mechanisms. This could be seen as a discouraging result, because a game need not have efficient equilibria, e.g., a ballot could result in an alternative nobody really wanted in the first place, yet everybody voted for. (For more, see Wiki.)
- Arrow's impossibility theorem applies to voting rules that produce a full ranked list; also called social welfare functions. Gibbard-Satterthwaite theorem proves a similar impossibility for voting rules that produce a single winner (social choice functions).

# 6 A Tractable Special Case: single-peaked preferences

- The Arrow's Theorem dashes any hopes of having a general unimpeachable voting rule. However, the need for voting still remains, and some rule must be chosen.
- Just as computational intractability results (e.g. $NP$-hardness) force us to compromise our goals, approximation algorithms, computationally tractable special cases, etc., impossibility theorems in voting theory imply that compromises must be made.
- In the following, we discuss one interesting restricted class of preferences for which a natural strategyproof voting rule works.
- Suppose the set of alternatives is the unit interval $[0, 1]$. For example, the interval could represent the political spectrum (from most liberal to most conservative).
- A voter $i$ has *single-peaked preferences* if there is a "peak" $p_i \in [0, 1]$ such that whenever $z$ is farther from $p_i$ than $y$, the voter prefers $y$ to $z$. That is, if $z < y < p_i$ or $z > y > p_i$, the voter prefers $y$ to $z$.
- One can imagine scenarios where single-peaked preferences are a reasonable first-cut approximation of voter's preferences: electing a politician, locating a school, etc..
- Suppose each voter votes by providing a reported peak $x_i \in [0, 1]$. Which alternative should we choose?

- One idea is to choose the $\frac{1}{n}\sum_{i=1}^{n} x_i$. But this is not strategyproof: a voter might be able to pull the chosen outcome closer to her peak by reporting an overly extreme peak. *Think of real-life examples of this phenomenon.*

- A second and better idea is to choose the *median* of the reported peaks; for simplicity, assume there is an odd number of voters.

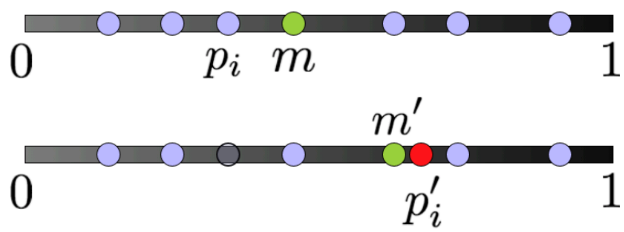


Figure 2: Median $m$ of reported peaks (with $n = 7$). A voter $i$ with peak $p_i$ cannot change the median unless she overshoots to the other side by misreporting $p'_i$; this would move the median $m'$ even further from $i$'s true preference $p_i$.

- The median voting rule is strategyproof. The only way a voter can manipulate the median is to report a peak on the opposite side of the median from her true peak, but this can only pull the median farther away from her true peak, a worse outcome for her.