

# Networks and Markets — Lecture 10: Social Choice, Voting, and Information Aggregation

Nikhil Garg, Cornell Tech

This lecture is partly based on the social-choice material in Easley and Kleinberg’s *Networks, Crowds, and Markets*, but it also includes modern applications, including Community Notes and participatory budgeting.

## Notes outline:

- Voting as aggregation of ranked preferences
- Why majority rule can fail with more than two candidates
- Borda count, ranked-choice voting, and impossibility results
- Wisdom of crowds and when information aggregation breaks
- Modern applications: Community Notes, participatory budgeting, deliberation, and AI evaluation

## Two kinds of aggregation

Social choice is about combining many people’s inputs into a collective decision. There are two conceptually different tasks:

- **Preference aggregation:** people have different tastes, priorities, or values, and we want a rule that turns these preferences into a social choice.
- **Information aggregation:** people are trying to estimate the same underlying fact, state of the world, or best action, and we want to combine noisy signals.

In information aggregation, we would like “experts” to have more weight if they exist, but this is generally undesirable in preference aggregation.

## 1 Voting

Suppose there is a set of alternatives

$$A = \{a_1, a_2, \dots, a_m\}.$$

Each voter  $i$  has a ranking over the alternatives, written  $\succ_i$ . For example,

$$a \succ_i b \succ_i c$$

means that voter  $i$  prefers  $a$  to  $b$  and  $b$  to  $c$ .

The collection of all voters' rankings is called a *preference profile*. A voting rule takes a profile as input and outputs a **winner**, a **set of winners**, or a **social ranking** of all alternatives.

When there are only two alternatives, majority rule is extremely natural. With three or more alternatives, things become much harder.

## 1.1 Majority rule and pairwise comparisons

Suppose there are only two candidates,  $A$  and  $B$ , and an odd number of voters. Then majority rule says:

- $A$  wins if more than half of voters rank  $A$  above  $B$ ,
- $B$  wins otherwise.

This rule has several attractive features:

- It treats voters equally.
- It treats the two alternatives equally.
- If a voter changes their mind in favor of  $A$ , that cannot hurt  $A$ . (This is called *monotonicity*.)
- If everyone prefers  $A$  to  $B$ , then  $A$  wins. (This is called *unanimity*).

So with two candidates, majority rule is intuitive and hard to beat. Extending “the majority idea” with more than 2 candidates is hard.

**Pairwise majority comparisons.** With three or more alternatives, one natural idea is to compare candidates pairwise. For any two alternatives  $x$  and  $y$ , say that

$$x \succ_M y$$

if a majority of voters prefer  $x$  to  $y$ .

This pairwise-majority relation is useful because it lets us ask whether there is a candidate who beats every other candidate one-on-one.

**Condorcet winner.** An alternative  $x$  is a *Condorcet winner* if, for every other alternative  $y$ ,

$$x \succ_M y.$$

If a Condorcet winner exists, it looks like a compelling candidate for the winner: in head-to-head comparisons, it defeats everyone else. Does a Condorcet winner always exist?

## 1.2 Failures of majority vote with three or more candidates

### 1.2.1 Condorcet paradox.

Here is the classic example with three voters and three candidates:

Voter 1	Voter 2	Voter 3
$A$	$B$	$C$
$B$	$C$	$A$
$C$	$A$	$B$

Now compare candidates pairwise.

- $A$  beats  $B$ , because voters 1 and 3 prefer  $A$  to  $B$ .
- $B$  beats  $C$ , because voters 1 and 2 prefer  $B$  to  $C$ .
- $C$  beats  $A$ , because voters 2 and 3 prefer  $C$  to  $A$ .

So majority preferences cycle:

$$A \succ_M B, \quad B \succ_M C, \quad C \succ_M A.$$

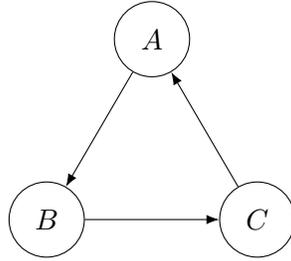


Figure 1: A majority cycle:  $A$  beats  $B$ ,  $B$  beats  $C$ , and  $C$  beats  $A$ .

This is the *Condorcet paradox*: even when each voter has a perfectly sensible ranking, the social majority relation can be intransitive. This means that there is no Condorcet winner, and pairwise majority comparisons do not yield a clear winner. What are properties that we may want with such preferences?

### 1.2.2 Strategic voting with multiple candidates.

Another failure of simple voting rules is that people may vote strategically rather than sincerely with more than 2 candidates.

**Spoiler effect example under *plurality* voting.** Suppose there are three candidates:

- $L_1$ : a left candidate,
- $L_2$ : another left candidate,
- $R$ : a right candidate.

Consider the following preference profile:

33 voters	32 voters	35 voters
$L_1$	$L_2$	$R$
$L_2$	$L_1$	$L_2$
$R$	$R$	$L_1$

Under plurality voting, first-choice totals are:

$$L_1 : 33, \quad L_2 : 32, \quad R : 35.$$

So  $R$  wins. So the presence of two similar left candidates allows  $R$  to win, even though a clear majority prefers either left candidate to  $R$ . This is the *spoiler effect*.

**Strategic response.** If voters who sincerely prefer  $L_2$  believe  $L_2$  cannot win, they may strategically vote for  $L_1$  to avoid electing  $R$ .

So plurality voting can punish sincere voting and reward tactical coordination.

### 1.2.3 A compromise candidate can still lose.

Suppose preferences are:

45 voters	40 voters	15 voters
$A$	$C$	$B$
$B$	$B$	$A$
$C$	$A$	$C$

Plurality first-choice totals are:

$$A : 45, \quad B : 15, \quad C : 40.$$

So the plurality elects  $A$ . Pairwise comparisons suggest perhaps  $B$  should win.

- $B$  beats  $A$  by 55–45 (the last two groups prefer  $B$  to  $A$ ).
- $B$  beats  $C$  by 60–40 (the first and third groups prefer  $B$  to  $C$ ).

So  $B$  is the Condorcet winner and a natural “compromise” candidate, but plurality ignores that information and chooses  $A$ . This profile will be useful later when we compare voting rules.

### 1.2.4 Approval voting with two winners – winner takes all dynamics

Approval voting lets each voter approve as many candidates as they want. For a single winner, that can soften some spoiler problems. But with two seats and a simple “top two approval scores win” rule, it can become winner-take-all. If 60% of voters approve candidates  $A$  and  $B$ , while 40% approve  $C$  and  $D$ , then  $A$  and  $B$  win both seats. So the majority bloc gets both winners, even if a more proportional outcome might have been “one seat for each side.”

## 1.3 Impossibility theorems: Arrow and Gibbard-Satterthwaite

The above are not just quirks of a few badly designed rules.

Arrow’s impossibility theorem says, very roughly, that with three or more alternatives there is no perfect rank-order voting rule that simultaneously handles arbitrary preferences, respects unanimity, ignores irrelevant alternatives, and avoids dictatorship. So there is no flawless procedure that converts everyone’s rankings into a social ranking while satisfying all of our favorite axioms at once.

Gibbard-Satterthwaite makes a parallel point about incentives: with three or more outcomes, any deterministic non-dictatorial rule that can choose different winners on different profiles is manipulable somewhere. In other words, strategic voting is a generic feature of deterministic social choice. The practical lesson from both theorems is that every real voting system has tradeoffs, so we should ask which imperfections we are willing to tolerate rather than searching for a perfect rule.

## 1.4 Alternative voting rules

Since plurality and pairwise majority can fail, people have designed other rules that use more of each ranking. Two of the most widely discussed are the Borda count and ranked-choice voting.

### 1.4.1 Borda count

With  $m$  candidates, the Borda count gives:

- $m - 1$  points to a first-place vote,
- $m - 2$  points to a second-place vote,
- and so on down to 0 points for last place.

The candidate with the largest total score wins.

**Worked example.** Return to the profile

45 voters	40 voters	15 voters
$A$	$C$	$B$
$B$	$B$	$A$
$C$	$A$	$C$

With three candidates, Borda scores are 2 for first, 1 for second, and 0 for third. So:

$$\text{score}(A) = 45 \cdot 2 + 40 \cdot 0 + 15 \cdot 1 = 105,$$

$$\text{score}(B) = 45 \cdot 1 + 40 \cdot 1 + 15 \cdot 2 = 115,$$

$$\text{score}(C) = 45 \cdot 0 + 40 \cdot 2 + 15 \cdot 0 = 80.$$

Thus Borda elects  $B$ .

This feels appealing here because  $B$  was the compromise candidate and also the Condorcet winner.

#### Benefits of the Borda count.

- It uses more information than plurality.
- It rewards candidates with broad support, not just intense first-choice support.
- It can favor compromise candidates that plurality misses.
- It produces a full ranking, not just a winner.

The Borda count is still manipulable. For example, voters may *bury* a strong rival by ranking them artificially low, even if that rival is not actually their least favorite option. Similarly, spoiler candidates can still exist.

### 1.4.2 Ranked-choice voting (instant runoff)

Ranked-choice voting (RCV), also called *instant-runoff voting*, works by repeated elimination:

1. count only first-choice votes;
2. if some candidate has a majority, stop;
3. otherwise eliminate the candidate with the fewest first-choice votes;
4. transfer those ballots to the next remaining candidate on each ballot;
5. repeat.

**Worked example.** Use the same profile again:

45 voters	40 voters	15 voters
<i>A</i>	<i>C</i>	<i>B</i>
<i>B</i>	<i>B</i>	<i>A</i>
<i>C</i>	<i>A</i>	<i>C</i>

Round 1 first-choice totals:

$$A : 45, \quad B : 15, \quad C : 40.$$

Candidate *B* is eliminated. Those 15 ballots transfer to *A* (their next remaining choice), giving

$$A : 60, \quad C : 40.$$

So RCV elects *A*.

### Benefits of ranked-choice voting.

- It lets voters rank candidates rather than naming only one.
- It can reduce some spoiler effects relative to plurality.
- It simulates a runoff without requiring a second election.
- It sometimes encourages candidates to seek second-choice support.

However, note that the compromise candidate *B* was everyone’s second choice except for one bloc who put *B* first. Under RCV, because *B* had too few first-choice votes, *B* never survived to the final round.  $\implies$  So RCV can still fail to elect the Condorcet winner.

**Comparing rules on one profile.** The same profile can produce different winners under different rules:

Rule	Winner on the 45/40/15 profile
Plurality	<i>A</i>
Pairwise majority / Condorcet	<i>B</i>
Borda count	<i>B</i>
Ranked-choice voting	<i>A</i>

This table is a good reminder that a voting rule is not merely a neutral way of “reading off” the people’s will. Different rules encode different ideas about what should matter: first-place support, pairwise strength, broad average ranking.

## 2 Information aggregation and the wisdom of crowds

So far we have been aggregating *preferences*. Now switch to a different setting: there is an underlying binary state of the world, and voters are trying to guess it.

- Is this factual claim true or false?

- Is the defendant guilty or not guilty?
- Is this note helpful or not helpful?
- Which of two diagnoses is correct?

Here there may be a meaningful notion of a correct answer, and some people may be more informed than others. This is where the “wisdom of crowds” idea enters.

## 2.1 Jury theorem and wisdom of crowds

Suppose:

- there is a correct binary answer,
- each voter gets a private signal,
- each voter is correct independently with probability  $p$ ,
- and  $p > \frac{1}{2}$ .

If there are  $n$  voters and  $n$  is odd, then the probability that majority vote is correct is

$$\sum_{j=(n+1)/2}^n \binom{n}{j} p^j (1-p)^{n-j}.$$

**Condorcet jury theorem.** If signals are independent and each voter is individually better than random ( $p > \frac{1}{2}$ ), then as the number of voters grows, the probability that majority vote is correct converges to 1.

If instead  $p < \frac{1}{2}$ , then adding more voters makes majority vote converge to the *wrong* answer. So even mildly informative independent signals can produce a highly accurate group decision.

## 2.2 When the wisdom of crowds fails

The jury theorem depends on assumptions that are often unrealistic. The crowd is wise only under specific conditions.

**Correlated errors.** If everyone reads the same bad article, uses the same flawed model, or copies the same misleading source, then many votes do not create many independent pieces of evidence.

In statistics one sometimes summarizes this by saying that the crowd’s *effective sample size* is much smaller than the number of voters.

So a crowd of 1000 highly correlated voters may contain far less information than a crowd of 50 diverse independent voters.

**Herding and information cascades.** A related failure mode is that people watch what others do and infer information from earlier actions.

For example, suppose there are 2 restaurants on a street, and the first restaurant is definitely better. Suppose each person receives a private signal about which restaurant is better – the signal indicates the first restaurant with probability  $p$ . The first person chooses a restaurant based on their signal. The second person uses both their private signal and the first person’s choice, and so on. If a long line ends up forming outside the second restaurant, future people might join it, even if their private signal points to the first restaurant.

### 2.3 Beyond majority vote: richer aggregation

Just like in preference aggregation, we can ask people for more than a single vote.

- ask for a confidence level,
- ask for a full ranking,
- ask for a probability distribution,
- or ask what they think other people will say.

**Surprisingly popular.** The *surprisingly popular* method asks each respondent two questions:

1. What do *you* think is the best answer?
2. What fraction of other respondents do you think will choose each answer?

The rule then looks for an answer that receives *more support than respondents predicted it would*.

Suppose there is a knowledgeable minority who know the correct answer, and they also know that most people will miss it. Then that answer may look “surprisingly popular” relative to the crowd’s own forecasts.

**Toy example.** What is the capital of Pennsylvania? Suppose there are two possible answers, *Harrisburg* and *Philadelphia*. Imagine:

- 40% of respondents choose *Harrisburg*, but on average they predict only 25% will choose *Harrisburg*;
- 60% choose *Philadelphia*, but on average they predict 75% will choose *Philadelphia*.

Then *Harrisburg* is the surprisingly popular answer:

$$40\% - 25\% = 15\%$$

above its predicted support, while *Philadelphia* is

$$60\% - 75\% = -15\%$$

below its predicted support.

This type of method is most useful when: there is a latent correct answer or better judgment, informed people are a minority, those informed people know they are a minority, and people in the majority overestimate how much they are in the majority.

**Other aggregation tricks.** There are many other ways to improve aggregation:

- **Confidence weighting:** ask for confidence as well as an answer.
- **Calibration weighting:** give more weight to people with better track records.
- **Prediction markets:** use prices to aggregate beliefs under incentives.
- **Peer prediction and information elicitation:** reward people based on how informative or predictive their reports are.

All of these approaches share the same basic idea: If we can extract richer information from each participant, we may aggregate better.

## 3 Applications

Community Notes is a particularly interesting modern application because it mixes social choice, information aggregation, incentives, and network structure.

### 3.1 Community Notes

Platforms see a huge number of misleading or false claims, that it might want to identify. A simple centralized moderation system does not scale well and often triggers complaints about bias or opacity. At the same time, a pure majority-rule system can be captured by brigading or by one large political bloc. Community Notes uses a “bridging” system to try to identify notes that are helpful across groups, not just inside one faction.

Very roughly, the workflow is:

1. a contributor writes a proposed note on a post;
2. other contributors rate whether that note is helpful;
3. the scoring system decides whether the note should be shown publicly.

**Bridging and estimated groups.** The key idea is *bridging*: a note should be stronger if it is rated helpful by users who often disagree on other notes. In the open-source system, these groups are not hand-labeled by party or ideology. Instead, they are estimated from the historical matrix of note ratings: users who tend to agree on which notes are helpful end up near one another in a latent space, and users who often disagree end up farther apart. So the system is trying to infer disagreement structure from behavior, not from self-reported identity.

**Numerical example.** Suppose there are two latent groups of raters,  $A$  and  $B$ . Note 1 gets 62 helpful ratings from group  $A$  and 5 from group  $B$ , while Note 2 gets 32 helpful ratings from group  $A$  and 32 from group  $B$ . If we only counted raw helpful votes, Note 1 would win 67 to 64. But a bridging score could still prefer Note 2, because Note 1 is supported almost entirely by one side while Note 2 gets substantial support from both sides of the disagreement.

### 3.2 Other applications

**Participatory budgeting.** Participatory budgeting asks citizens to help allocate a public budget across projects such as parks, transit stops, or school improvements. It is a social-choice problem over a *bundle* of public goods, so the mechanism must aggregate preferences while respecting a budget constraint and some notion of fairness across groups or neighborhoods.

**Deliberative democracy.** Deliberative democracy puts discussion before voting, as in citizens’ assemblies or mini-publics. The hope is that deliberation improves information and judgment, but it can also introduce status hierarchies or social pressure, so the discussion process is itself part of the aggregation mechanism.

**AI alignment, data labeling, and “LLM as judge”.** Modern AI systems create a new version of the same old problem: how do we aggregate many noisy human or model judgments into one training signal or evaluation? Data labeling, preference learning, and “LLM as judge” all raise familiar social-choice questions about disagreement, bias, weighting, and whether simple majority is the right aggregation rule.