Voting in pursuit of the truth

Vincent Conitzer
conitzer@cs.duke.edu

Objectives of voting

- **OBJ1:** Compromise among subjective preferences

- **OBJ2:** Reveal the "truth"

A model with two alternatives

- One of the two alternatives {R, B} is the "correct" winner; this is not directly observed
- Each voter votes for the correct winner with probability $p > ½$, for the other with $1-p$ (i.i.d.)
- The probability of a given profile in which r is the number of votes for R and b that for B ($r+b=n$)...
  - ... given that R is the correct winner is $p^r(1-p)^b$
  - ... given that B is the correct winner is $p^b(1-p)^r$

Condorcet Jury Theorem [1785]

- **Theorem.** In the model on the previous slide, if we use the majority rule, the probability that the correct winner is chosen goes to 1 as the number of votes goes to infinity.
- *Proof*: by law of large numbers the fraction of votes for the correct winner will converge to p
- Lots of variants: voters of unequal skill, correlated voters, strategic voters, …

Statistical estimation

- We have some parameters of our statistical model with unknown values
  - for example, the mean and variance of people's height
  - … in our case, which alternative is the correct winner
- We have some data
  - for example, some people's height
  - … in our case, the votes
- Want to estimate the values of the parameters

Maximum likelihood estimation

- Choose the parameter values q that maximize the likelihood of the data

- I.e., choose arg max P(data | q)

Example of MLE (**not** our main model)

- *p* of the voters vote for R; we don't know *p*, try to estimate it
- From a poll, we observe votes for R, B, B, R, B
- Let's call our estimate *q*
- If *q* were correct, the probability of our profile would be $q^2(1-q)^3$
- Differentiating with respect to *q*:

  $2q(1-q)^3 - 3q^2(1-q)^2 = 0$ or $2(1-q)-3q=0$ or $q=2/5$
- ... which is the sample fraction of R (generally true)

Back to our main model

- One of the two alternatives {R, B} is the "correct" winner; this is not directly observed
- Each voter votes for the correct winner with probability $p > ½$, for the other with $1-p$ (i.i.d.)
- The probability of a given profile in which r is the number of votes for R and b that for B ($r+b=n$)...
  - ... given that R is the correct winner is $p^r(1-p)^b$
  - ... given that B is the correct winner is $p^b(1-p)^r$
- Maximum likelihood estimate: correct winner = whichever has more votes (majority rule)

What if voters have different skill?

[Nitzan & Paroush 1982, Shapley and Grofman 1984]

- Each voter i votes for the correct winner with probability $p_i > ½$, for the other with $1-p_i$ (independently)
- The probability of a given profile in which r is the number of votes for R and b that for B ($r+b=n$)...
  - ... given that R is the correct winner is

    $(\prod_{i \text{ for R}} p_i)(\prod_{i \text{ for B}} (1-p_i))$ →**take log**→ $\Sigma_{i \text{ for R}} \log p_i + \Sigma_{i \text{ for B}} \log(1-p_i)$
    $= \Sigma_{i \text{ for R}} (\log p_i - \log(1-p_i)) + \Sigma_{i \text{ in } \textbf{all votes}} \log(1-p_i)$
  - ... given that B is the correct winner is

    $(\prod_{i \text{ for B}} p_i)(\prod_{i \text{ for R}} (1-p_i))$ →**take log**→ $\Sigma_{i \text{ for B}} \log p_i + \Sigma_{i \text{ for R}} \log(1-p_i)$
    $= \Sigma_{i \text{ for B}} (\log p_i - \log(1-p_i)) + \Sigma_{i \text{ in } \textbf{all votes}} \log(1-p_i)$
- Hence if we do weighted majority with weight $\log p_i - \log(1-p_i) = \log [p_i / (1-p_i)]$ we will get the MLE!

A bit more about probability

- Example: roll two dice
- Random variables:
  - X = value of die 1
  - Y = value of die 2
- Outcome is represented by an ordered pair of values (x, y)
  - E.g., (6, 1): X=6, Y=1
  - Atomic event or sample point tells us the **complete** state of the world, i.e., values of **all** random variables
- Exactly one atomic event will happen; each atomic event has a ≥0 probability; sum to 1
  - E.g., P(X=1 and Y=6) = 1/36

Y \ X	1	2	3	4	5	6
6	1/36	1/36	1/36	1/36	1/36	1/36
5	1/36	1/36	1/36	1/36	1/36	1/36
4	1/36	1/36	1/36	1/36	1/36	1/36
3	1/36	1/36	1/36	1/36	1/36	1/36
2	1/36	1/36	1/36	1/36	1/36	1/36
1	1/36	1/36	1/36	1/36	1/36	1/36

- An event is a proposition about the state (=subset of states)
  - X+Y = 7
- Probability of event = sum of probabilities of atomic events where event is true

Conditional probability

- We might know something about the world – e.g., "X+Y=6 or X+Y=7" – given this (and **only** this), what is the probability of Y=5?
- Part of the sample space is eliminated; probabilities are renormalized to sum to 1

Y \ X	1	2	3	4	5	6
6	1/36	1/36	1/36	1/36	1/36	1/36
5	1/36	1/36	1/36	1/36	1/36	1/36
4	1/36	1/36	1/36	1/36	1/36	1/36
3	1/36	1/36	1/36	1/36	1/36	1/36
2	1/36	1/36	1/36	1/36	1/36	1/36
1	1/36	1/36	1/36	1/36	1/36	1/36

Y \ X	1	2	3	4	5	6
6	1/11	0	0	0	0	0
5	1/11	1/11	0	0	0	0
4	0	1/11	1/11	0	0	0
3	0	0	1/11	1/11	0	0
2	0	0	0	1/11	1/11	0
1	0	0	0	0	1/11	1/11

- P(Y=5 | (X+Y=6) or (X+Y=7)) = 2/11

Facts about conditional probability

- P(A | B) = P(A and B) / P(B)

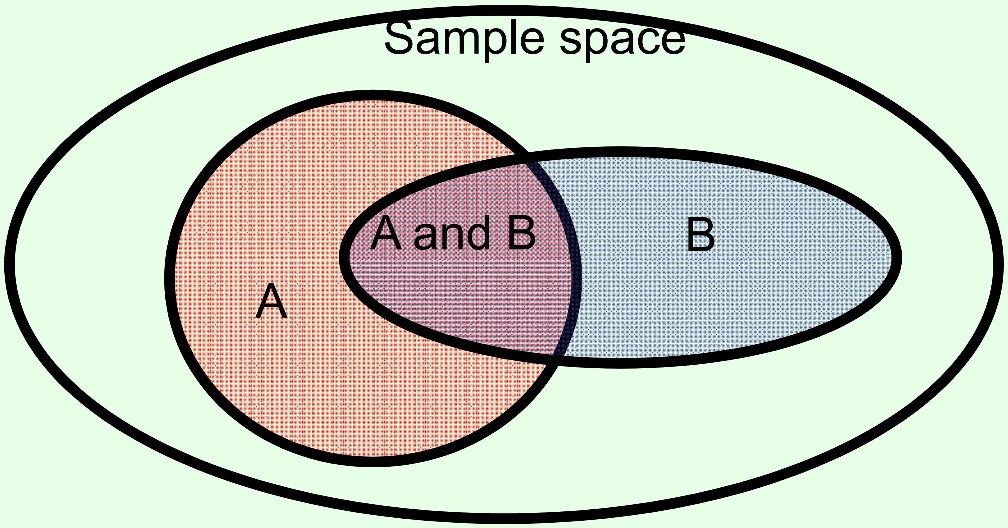

- P(A | B)P(B) = P(A and B) = P(B | A)P(A)
- P(A | B) = P(B | A)P(A)/P(B)
  - Bayes' rule

Maximum a posteriori

- Maybe what we really want is
  arg max P(q | data)

- But by Bayes' rule,
  P(q | data) = P(data | q) P(q) / P(data)

- So **if** P(q) is uniform,
  arg max P(q | data) = arg max P(data | q)
  (MLE)

More generally: The MLE approach to voting

[can be said to date back to Condorcet 1785]

- Given the "correct outcome" c
  - each vote is drawn conditionally independently given c, according to $\Pr(V|c)$

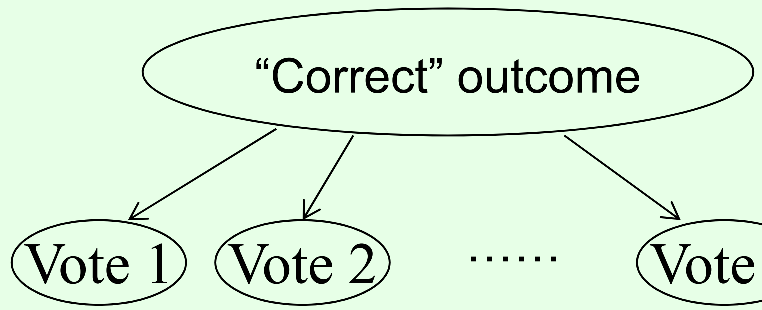

- The MLE rule: For any profile $P = (V_1, \ldots, V_n)$,
  - The likelihood of P given c: $L(P|c)=\Pr(P|c)=\prod_{V\in P}\Pr(V|c)$
  - The MLE as a rule is defined as

$$\text{MLE}_{\Pr}(P)=\text{argmax}_c\prod_{V\in P}\Pr(V|c)$$

- Body of work characterizing these with >2 alternatives [Young '88, '95, C. & Sandholm '05, C., Rognlie, Xia '09, Xia & C. '11, Procaccia, Reddi, Shah '12 …]

A noise model for >2 alternatives

[dating back to Condorcet 1785]

- Correct outcome is a ranking W , $p>1/2$

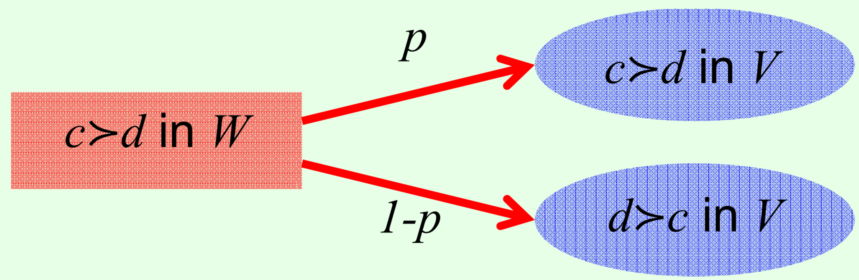

$$\Pr(\, b \succ c \succ a \mid a \succ b \succ c\,) = \; ?\; (1\text{-}p)p)^2$$

- MLE = Kemeny rule [Young '88, '95]
  - $\Pr(P|W) = p^{nm(m-1)/2-K(P,W)}\,(1\text{-}p)^{K(P,W)} = p^{nm(m-1)/2}\left(\frac{1-p}{p}\right)^{K(P,W)}$
  - The winning rankings are insensitive to the choice of p $(>1/2)$

A variant for partial orders

[Xia & C. IJCAI-11]

- Parameterized by $p_+ > p_- \geq 0$ $(p_+ + p_- \leq 1)$
- Given the "correct" ranking W, generate pairwise comparisons in a vote V_{PO} independently

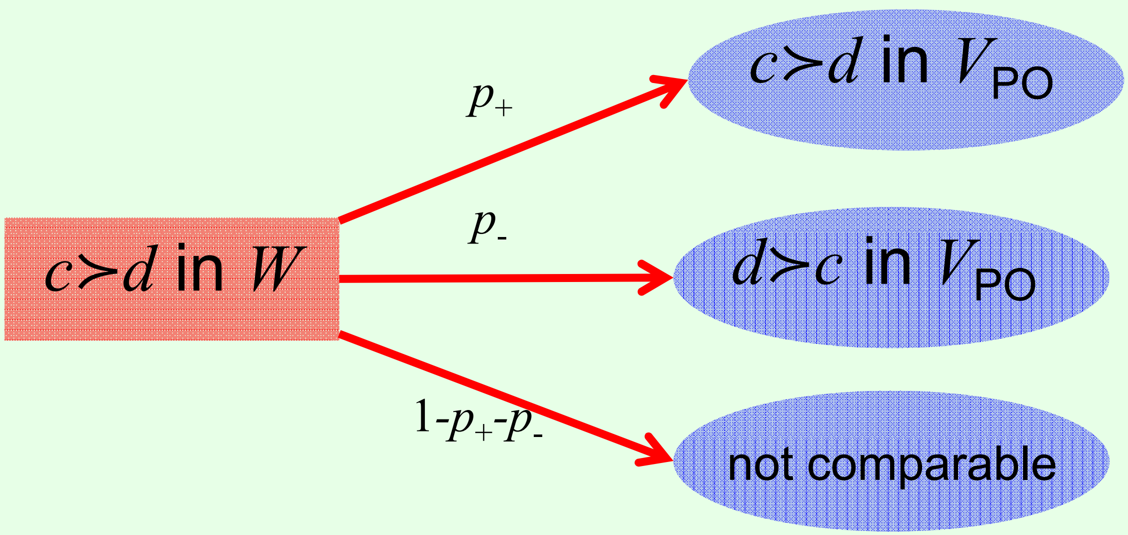

MLE for partial orders…

[Xia & C. IJCAI-11]

- In the variant of Condorcet's model
  - Let T denote the number of pairwise comparisons in P_{PO}
  - $\Pr(P_{\text{PO}}|W) = (p_+)^{T\text{-}\mathrm{K}(P_{\text{PO}},W)} (p_-)^{\mathrm{K}(P_{\text{PO}},W)} (1\text{-}p_+\text{-}p_-)^{\mathrm{nm(m\text{-}1)/2}\text{-}T}$

$$= (1-p_+-p_-)^{nm(m-1)/2-T}(p_+)^T\left(\frac{p_-}{p_+}\right)^{K(P_{PO},W)}$$

  - The winner is $\mathrm{argmin}_W \mathrm{K}(P_{\text{PO}},W)$

Which other common rules are MLEs for some noise model?

[C. & Sandholm UAI'05; C., Rognlie, Xia IJCAI'09]

- Positional scoring rules
- STV - kind of…
- Other common rules are provably not
- Reinforcement: if $f(V_1) \cap f(V_2) \neq \emptyset$ then $f(V_1+V_2) = f(V_1) \cap f(V_2)$ (f returns rankings)
- *Every MLE rule must satisfy reinforcement!* (why?)
- Incidentally: Kemeny uniquely satisfies neutrality, reinforcement, and Condorcet property [Young & Levenglick 78]

Correct **alternatives**

- Suppose the ground truth outcome is a correct alternative (instead of a ranking)
- Positional scoring rules are still MLEs
- Reinforcement: if $f(V_1) \cap f(V_2) \neq \emptyset$ then $f(V_1+V_2) = f(V_1) \cap f(V_2)$ (but now f produces a winner)
- Positional scoring rules* are the only voting rules that satisfy anonymity, neutrality, and reinforcement! [Smith '73, Young '75]
  - * Can also break ties with another scoring rule, etc.
- Similar characterization using reinforcement for ranking?

Independence assumption ignores social network structure

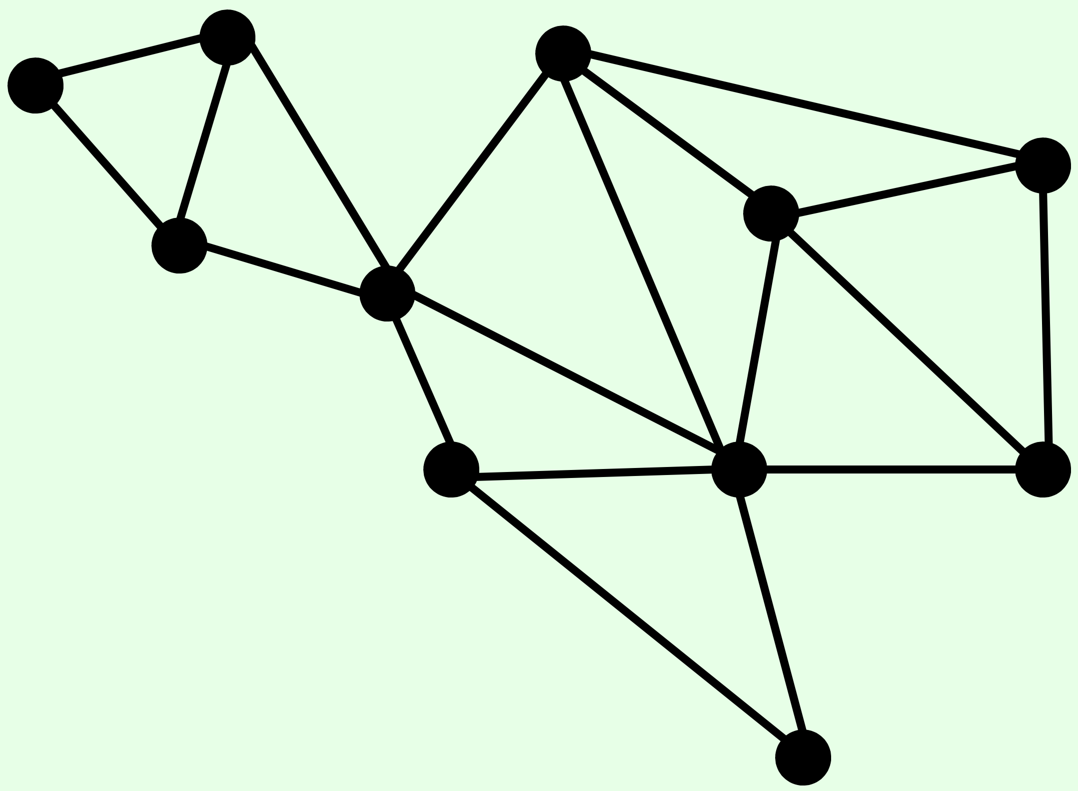

Voters are likely to vote similarly to their neighbors!

What should we do if we know the social network?

- Argument 1: *"Well-connected voters benefit from the insight of others so they are more likely to get the answer right. They should be weighed more heavily."*
- Argument 2: *"Well-connected voters do not give the issue much independent thought; the reasons for their votes are already reflected in their neighbors' votes. They should be weighed less heavily."*
- Argument 3: *"We need to do something a little more sophisticated than merely weigh the votes (maybe some loose variant of districting, electoral college, or something else...)."*

Factored distribution

- Let A_v be v's vote, N(v) the neighbors of v
- Associate a function $f_v(A_v, A_{N(v)} \mid c)$ with node v (for c as the correct winner)
- Given correct winner c, the probability of the profile is $\Pi_v\ f_v(A_v, A_{N(v)} \mid c)$
- **Assume:**

  $f_v(A_v, A_{N(v)} \mid c) = g_v(A_v \mid c)\ h_v(A_v, A_{N(v)})$

  - Interaction effect is independent of correct winner

Example
(2 alternatives, 2 connected voters)

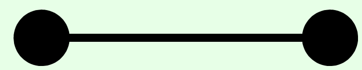

- $g_v(A_v=c \mid c) = .7$, $g_v(A_v= -c \mid c) = .3$
- $h_{vv'}(A_v=c, A_{v'}=c) = 1.142$,

  $h_{vv'}(A_v=c, A_{v'}=-c) = .762$
- $P(A_v=c \mid c) =$

  $P(A_v=c, A_{v'}=c \mid c) + P(A_v=c, A_{v'}=-c \mid c) =$
  $(.7*1.142*.7*1.142 + .7*.762*.3*.762) = .761$

- (No interaction: h=1, so that $P(A_v=c \mid c) = .7$)

Social network structure does not matter!

- **Theorem.** The maximum likelihood winner does not depend on the social network structure. (So for two alternatives, majority remains optimal.)
- *Proof.*

  $\arg\max_c \Pi_v f_v(A_v, A_{N(v)} \mid c) =$

  $\arg\max_c \Pi_v g_v(A_v \mid c)\, h_v(A_v, A_{N(v)}) =$

  $\arg\max_c \Pi_v g_v(A_v \mid c).$

The independent conversations model

- Edges are associated with alternatives
  - i.i.d., $p>0.5$ of correct alternative
- Voters go with the majority of their edges

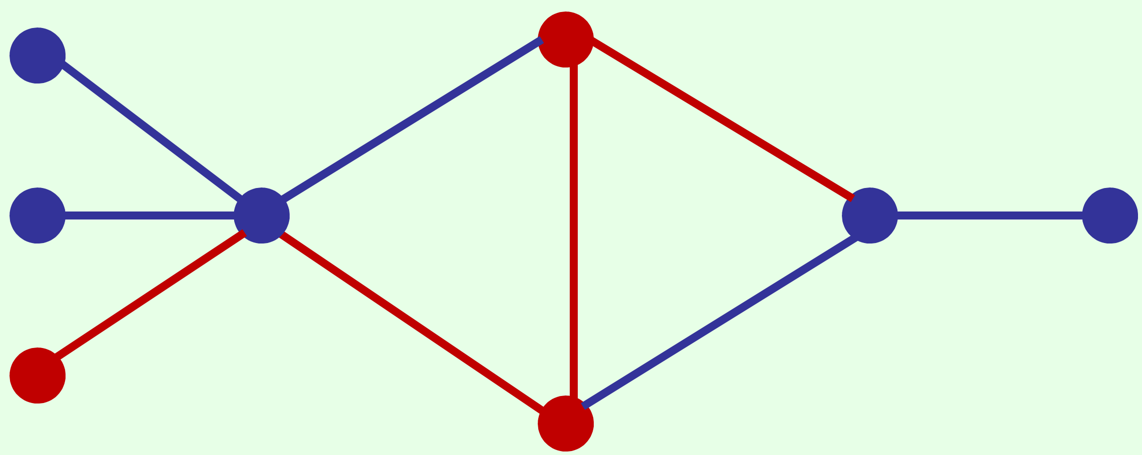

- If blue is correct, the probability of this configuration is $p^5(1-p)^4$; if red is correct, it is $p^4(1-p)^5$
- ... but we don’t actually know the edge colors!

Attempt #2

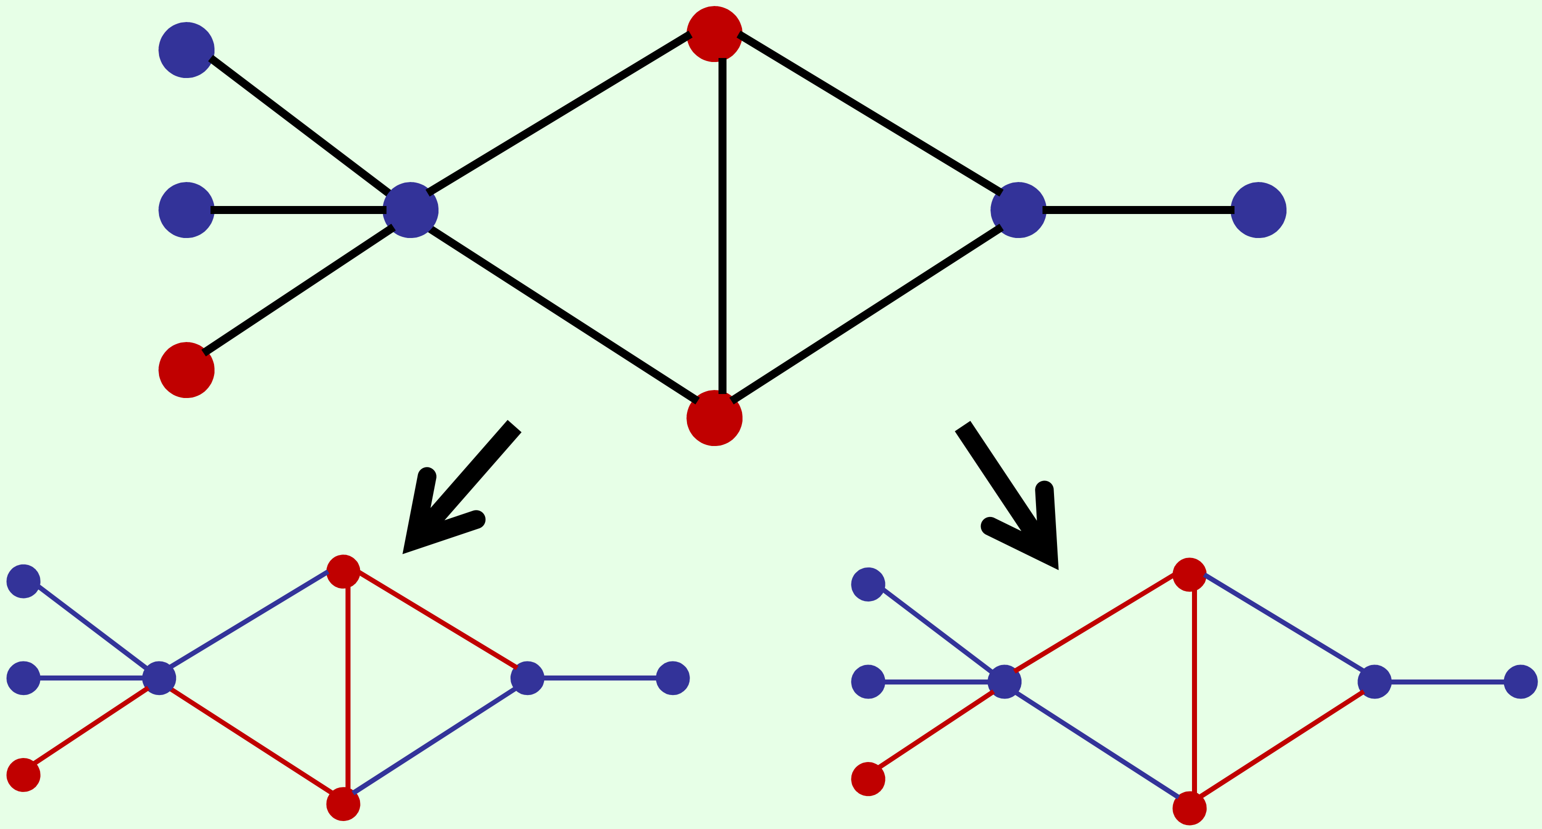

If blue is correct, the probability of the votes is $2p^5(1-p)^4$; if red is correct, it is $2p^4(1-p)^5$

How about this one?

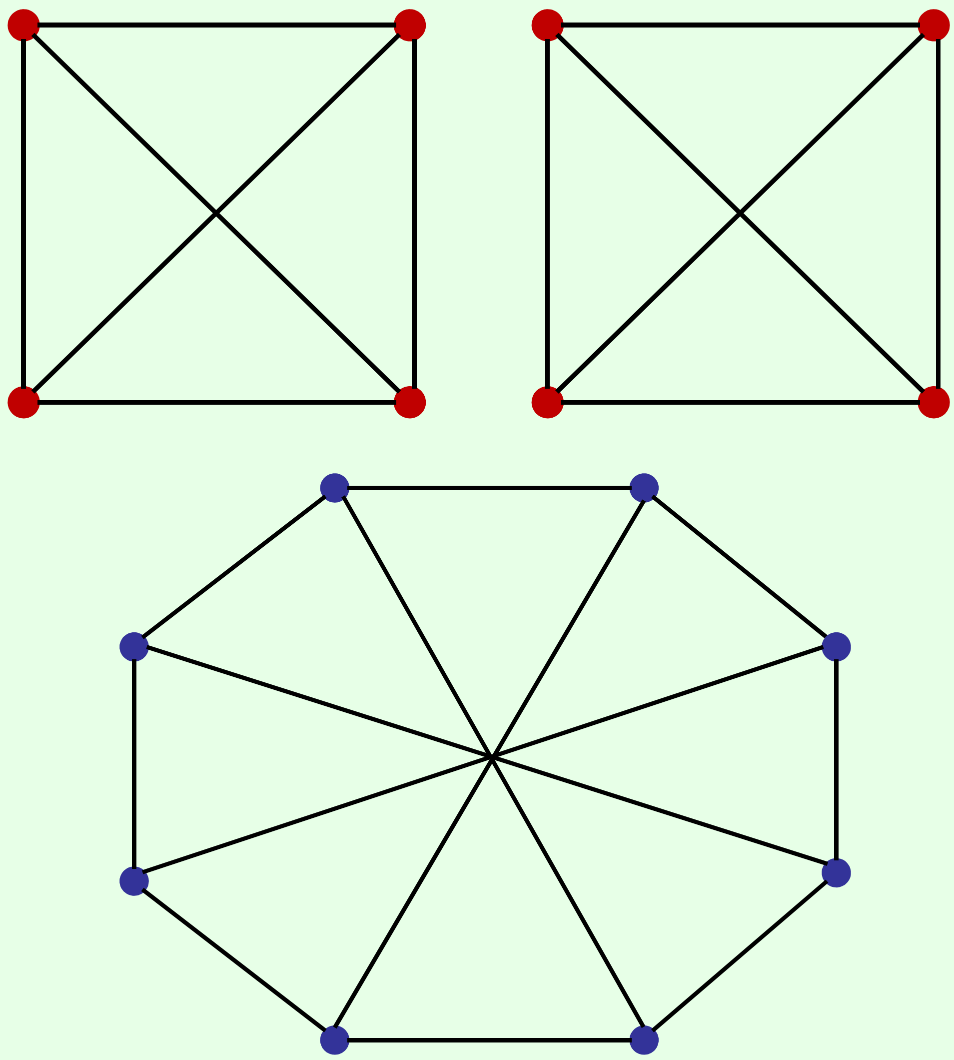

- Opposite colored edges must be a matching
- Matchings for red part: 1 of 0, 12 of 1, 42 of 2, 36 of 3, 9 of 4
- Matchings for blue part: 1 of 0, 12 of 1, 42 of 2, **44** of 3, **7** of 4
- If p=0.9, blue wins
- If p=0.6, red wins!

#P-hardness

- **Theorem.** Computing $P(A_V \mid c)$ is #P-hard.
- *Proof.* Reduction from PERMANENT (counting the number of perfect bipartite matchings)

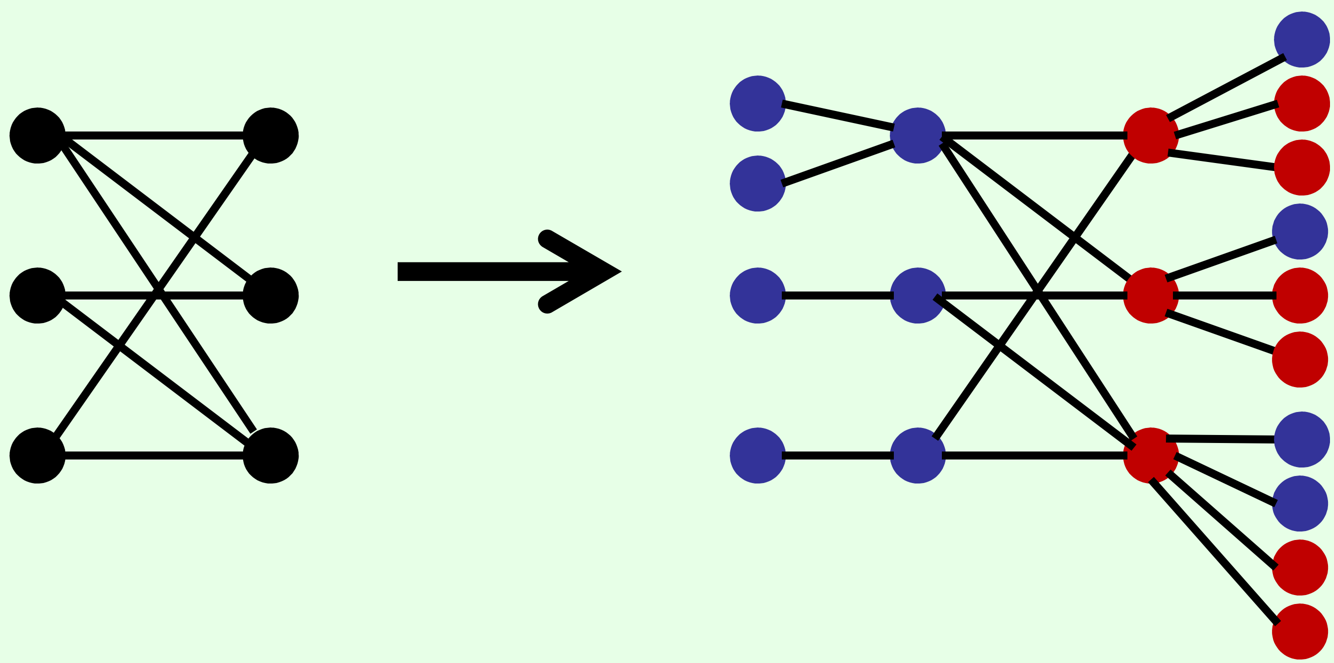

Estimating c and A_E together

- **Theorem.** An element of
  $\arg\max_{c,A_E} P(A_V, A_E \mid c)$
  can be found in polynomial time (even if edges have different p_e)
- *Proof.* For given c, reduction to a general version of *b*-matching
  - Choosing an edge = it votes for c
  - Each vertex has lower and upper bounds on number of matches
  - Weight on e is $\log p_e - \log(1-p_e)$
    - Goal is to max $\Sigma_{e \text{ votes for } c} \log p_e + \Sigma_{e \text{ votes for } -c} \log(1-p_e)$

Conclusions

- In some voting contexts we might think there's a correct outcome and votes are noisy perceptions of the truth
- MLE model allows us to
  - formalize this intuition,
  - justify certain voting rules as optimal, but also:
  - identify the assumptions going into those justifications (and modify them if we don't like them)
- General methodology that can also be applied to voting in combinatorial domains, judgment aggregation, etc.