Evaluating Resistance to False-Name Manipulations in Elections

Bo Waggoner

Lirong Xia

Vincent Conitzer

Thanks to Hossein Azari and Giorgos Zervas for helpful discussions!

Outline

- Background and motivation: Why study elections in which we expect false-name votes?
- Our model
- How to **select** a false-name-limiting method?
- How to **evaluate** the election outcome?
- Recap and future work

Motivating Challenge:

Poll customers about a potential product

A

B

C

Preventing strategic behavior

Deter or hinder misreporting

- Restricted settings (e.g., single-peaked preferences)
- Use computational complexity



False-name manipulation

- False-name-proof voting mechanisms?
- **Extremely** negative result for voting [C., WINE'08]
- Restricting to single-peaked preferences does not help much [Todo, Iwasaki, Yokoo, AAMAS'11]
- Assume creating additional identifiers comes at a cost [Wagman & C., AAAI'08]
- Verify some of the identities [C., TARK'07]
- Use social network structure [C., Immorlica, Letchford, Munagala, Wagman, WINE'10]

Overview article [C., Yokoo, AIMag 2010]

Common factor: false-name-*proof*

Let’s at least put up some obstacles

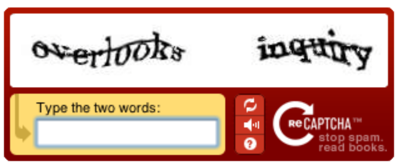

140.247.232.88 jmhzdszx@sharklasers.com

Issues:

1. Some still vote multiple times
2. Some don’t vote at all

Approach

Suppose we can experimentally determine how many identities voters tend to use for each method.

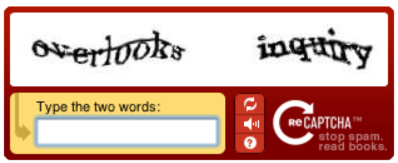

140.247.232.88

jmhzdszx@sharklasers.com

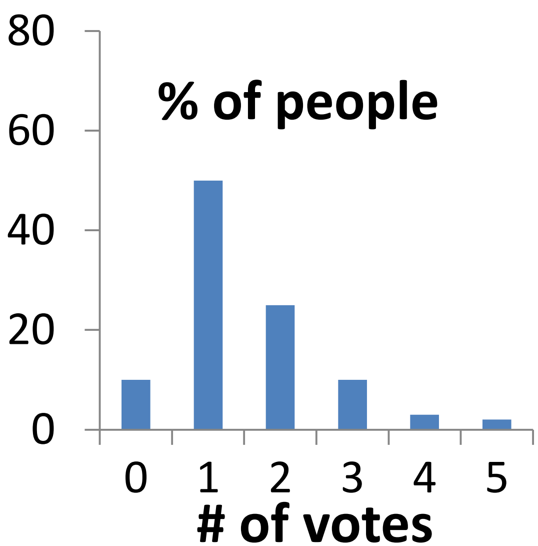

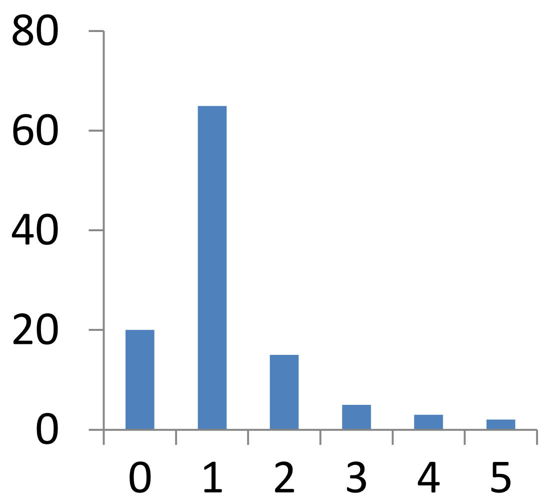

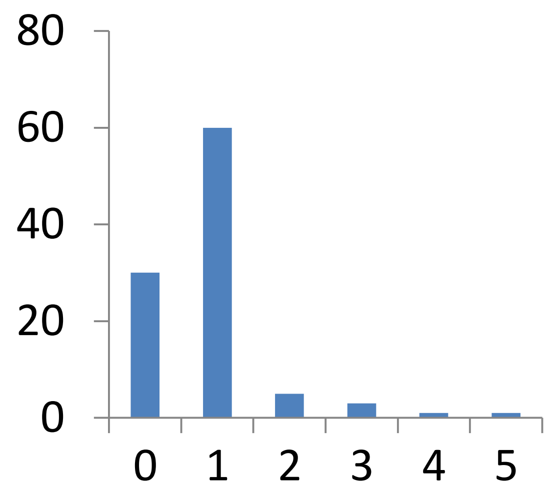

Outline

- Background and motivation: Why study elections in which we expect false-name votes?
- Our model
- How to **select** a false-name-limiting method?
- How to **evaluate** the election outcome?
- Recap and future work

Model

- For each false-name-limiting method, take the individual vote distribution π as given
- Suppose votes are drawn i.i.d.

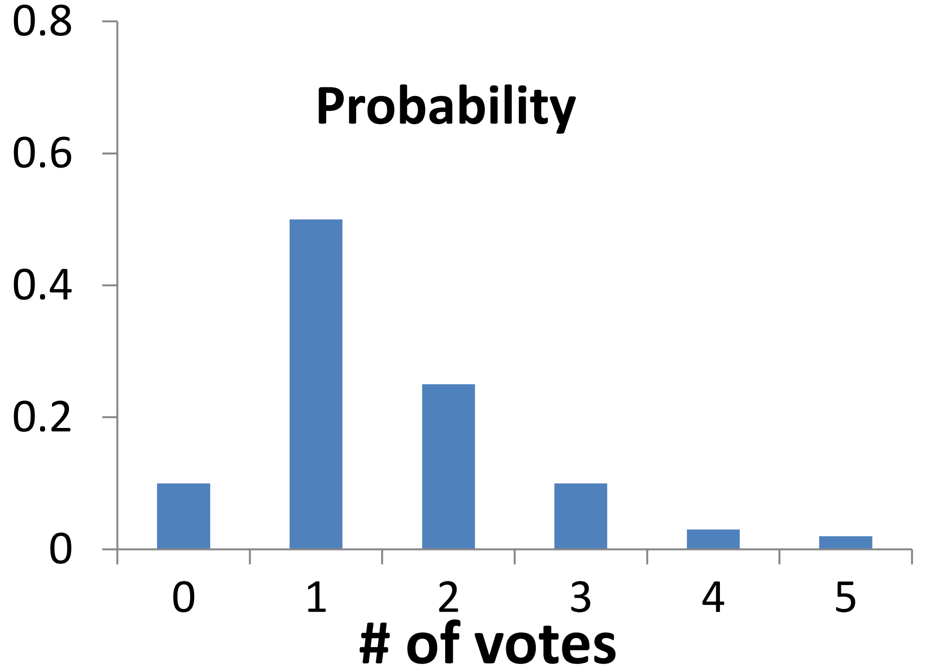

Model

- Single-peaked preferences (here: two alternatives)

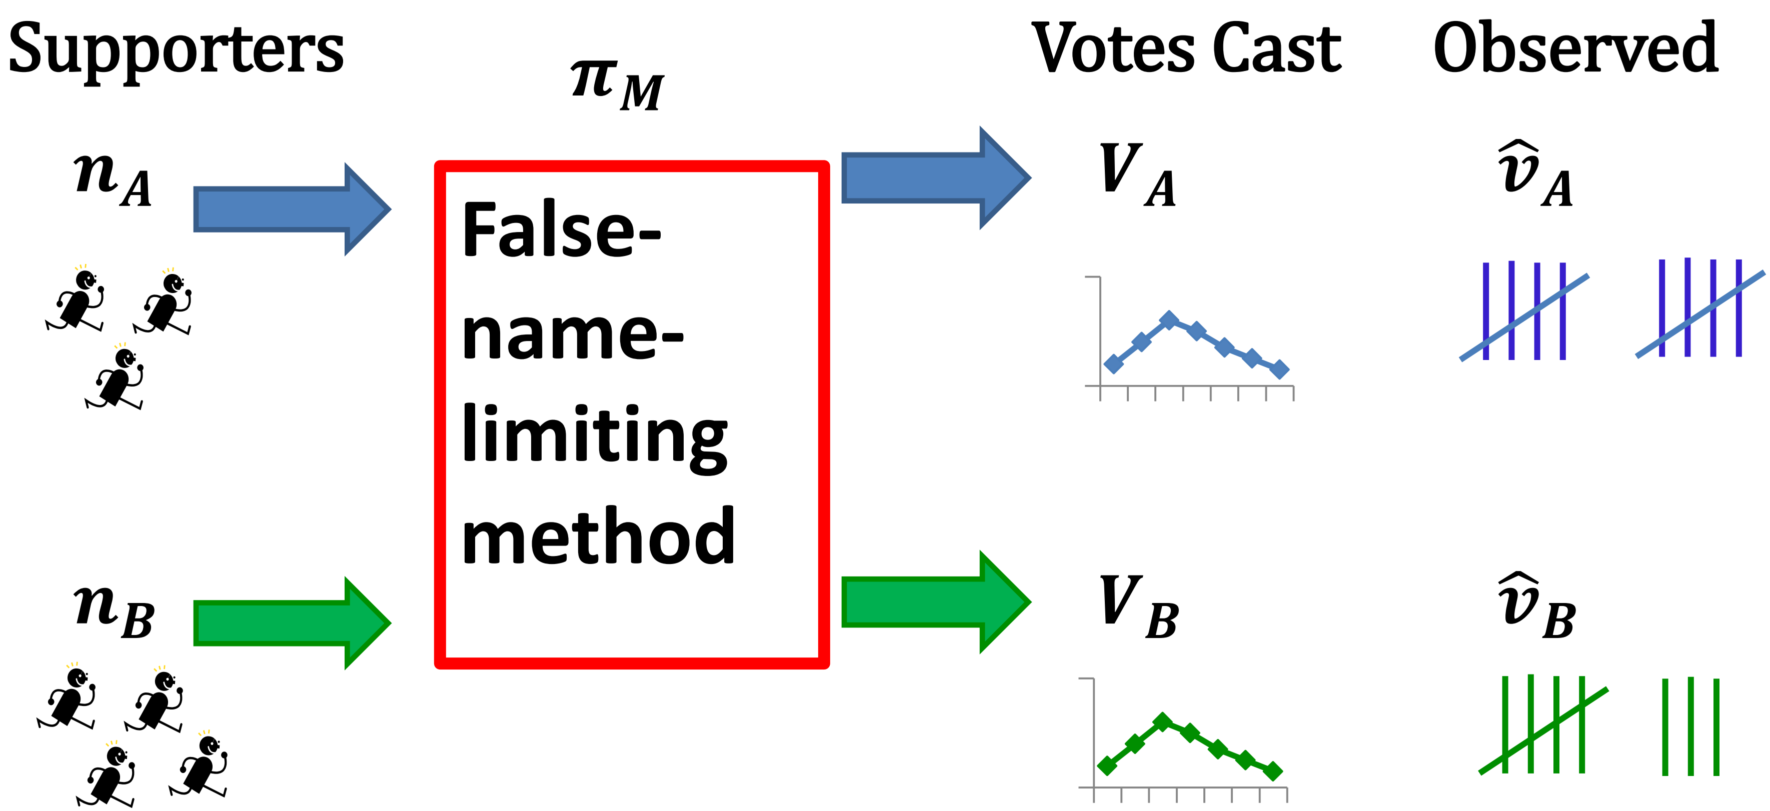

Outline

- Background and motivation: Why study elections in which we expect false-name votes?
- Our model
- How to **select** a false-name-limiting method?
- How to **evaluate** the election outcome?
- Recap and future work

Example

- Is the choice always obvious?
- Individual vote distribution for 2010 U.S. midterm Congressional elections:

Actual (in-person)

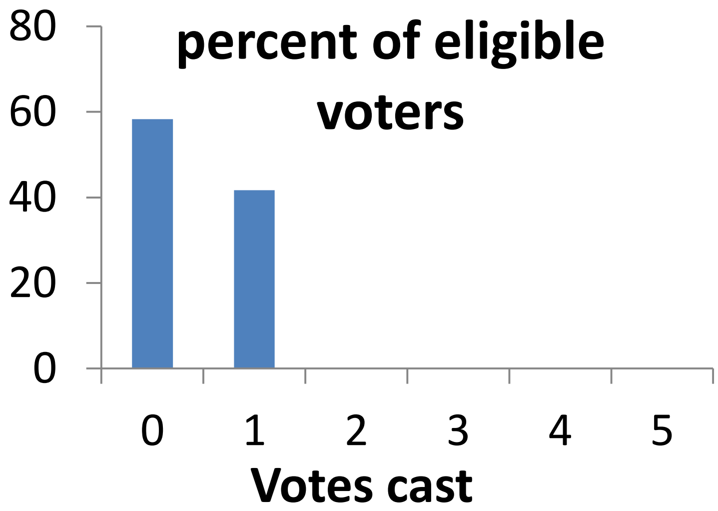

Hypothetical (online)

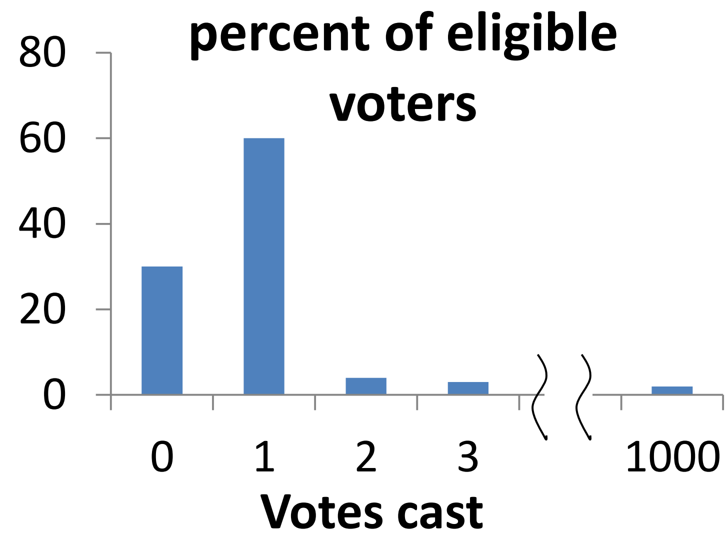

Problem statement

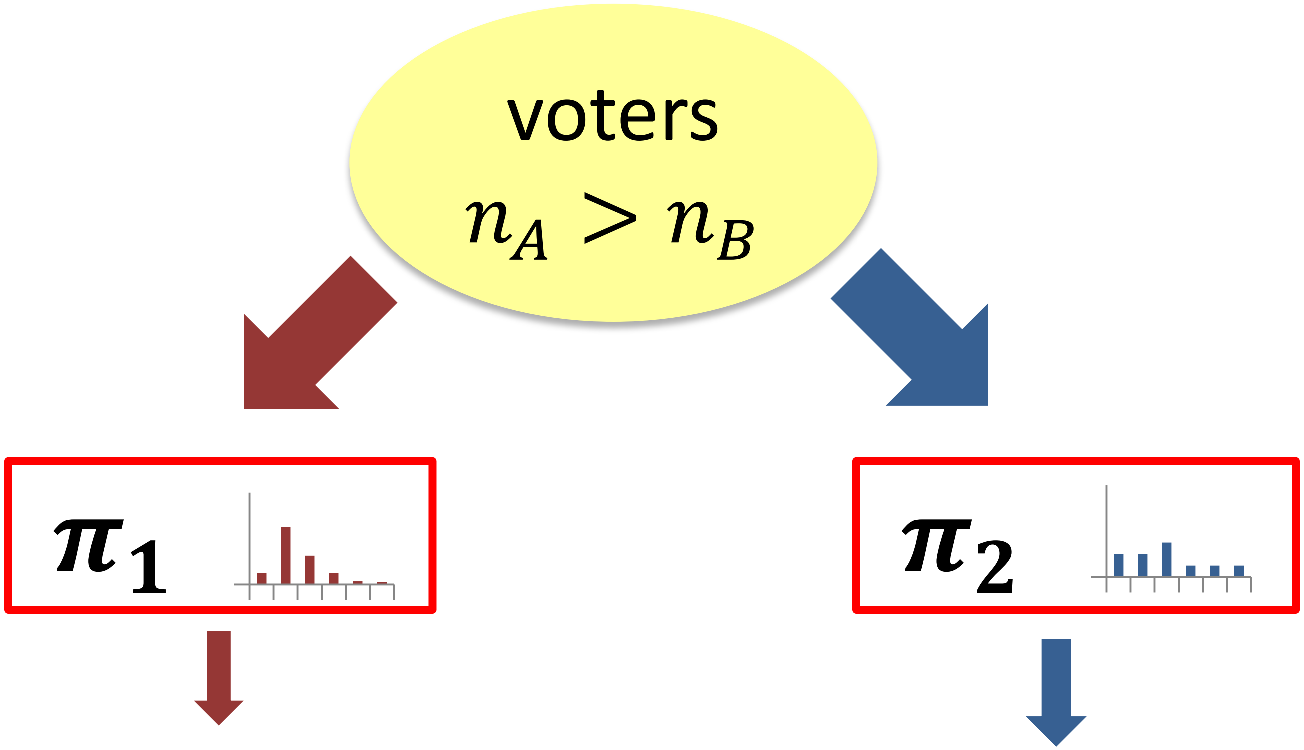

Pr[correct | π_1] **>** Pr[correct | π_2] **?**

(Pr[correct] = Pr[$V_A > V_B$])

Our results

- We show: which of π_1 and π_2 is preferable as elections grow large

- Setting: sequence of growing supporter profiles (n_A, n_B) where:
  1. $n_A - n_B \in O(\sqrt{n})$ (elections are "close")
  2. $n_A - n_B \in \omega(1)$ (but not "dead even")

Selecting a false-name-limiting method

Theorem 1.
Suppose $\frac{\mu_1}{\sigma_1} > \frac{\mu_2}{\sigma_2}$. *Then eventually*
$\Pr[\text{correct} \mid \pi_1] > \Pr[\text{correct} \mid \pi_2]$.

"For large enough elections, the ratio of mean to standard deviation is all that matters."

Selecting a false-name-limiting method

Intuition.

- Distributions approach Gaussians
- Pr[correct] = $\Pr[V_A > V_B] = \Pr[V_A - V_B > 0]$ approaches $\Phi\left(\frac{\mu}{\sigma}\,\frac{n_A - n_B}{\sqrt{n}}\right)$.

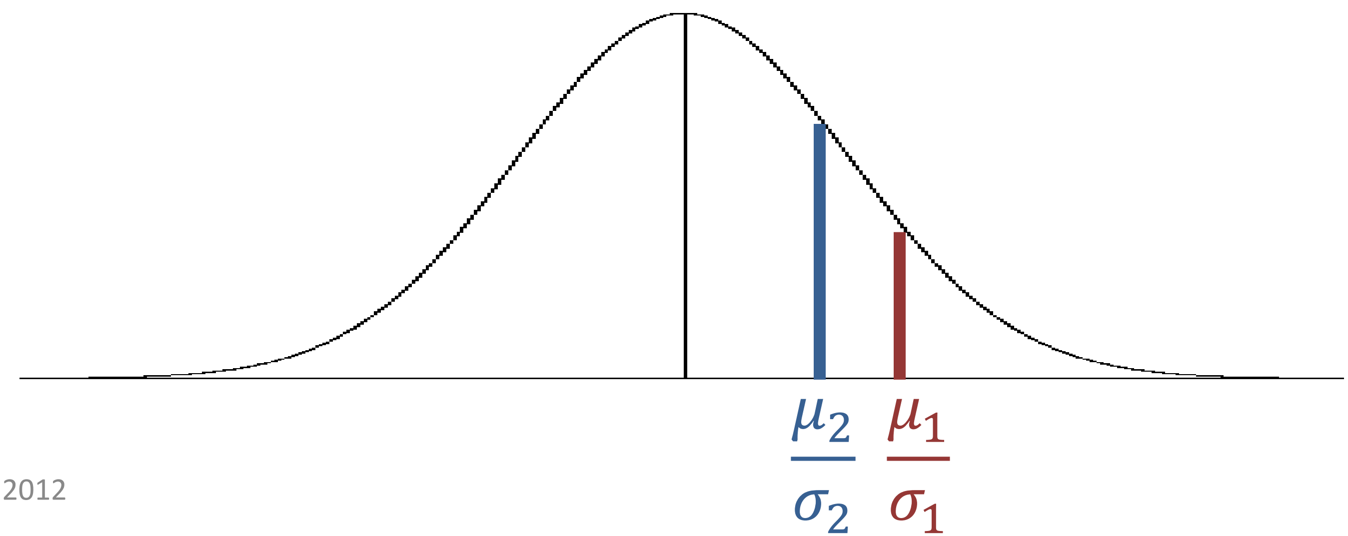

Question 1 Recap

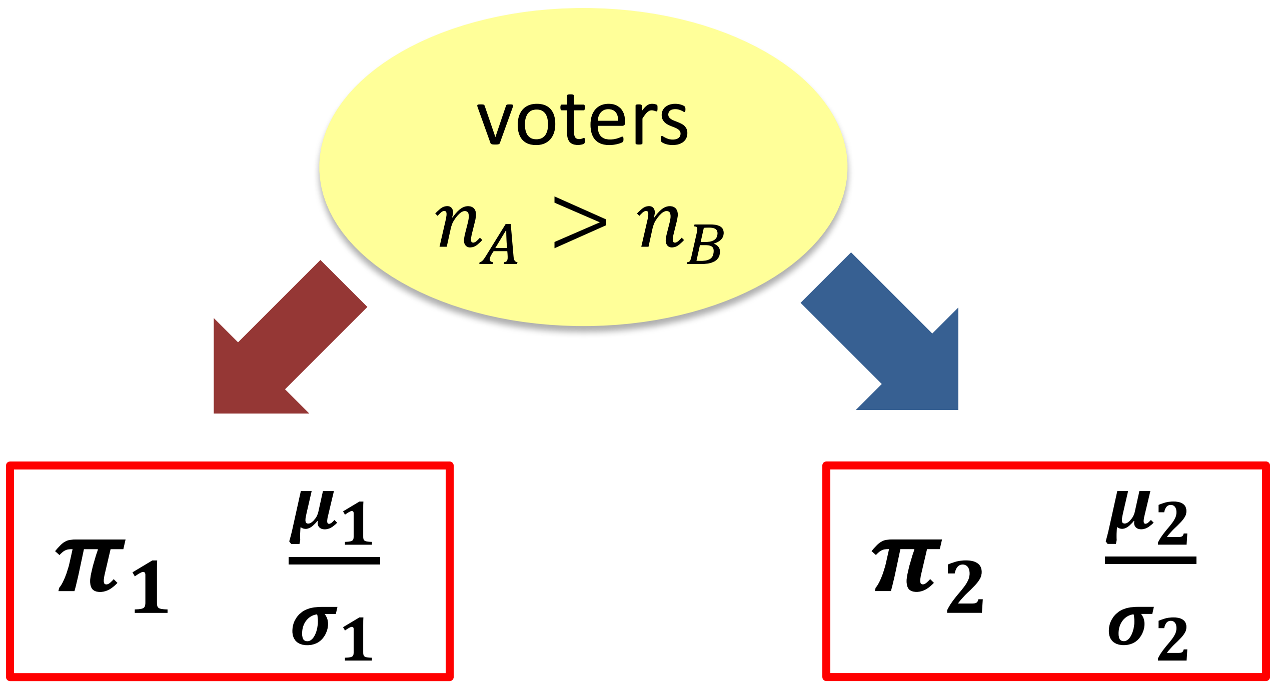

- Takeaway: choose highest ratio!
- Inspiration for new methods?

Outline

- Background and motivation: Why study elections in which we expect false-name votes?
- Our model
- How to **select** a false-name-limiting method?
- How to **evaluate** the election outcome?
- Recap and future work

Analyzing election results

- Observe votes $\hat{v}_A > \hat{v}_B$
- One approach: Bayesian

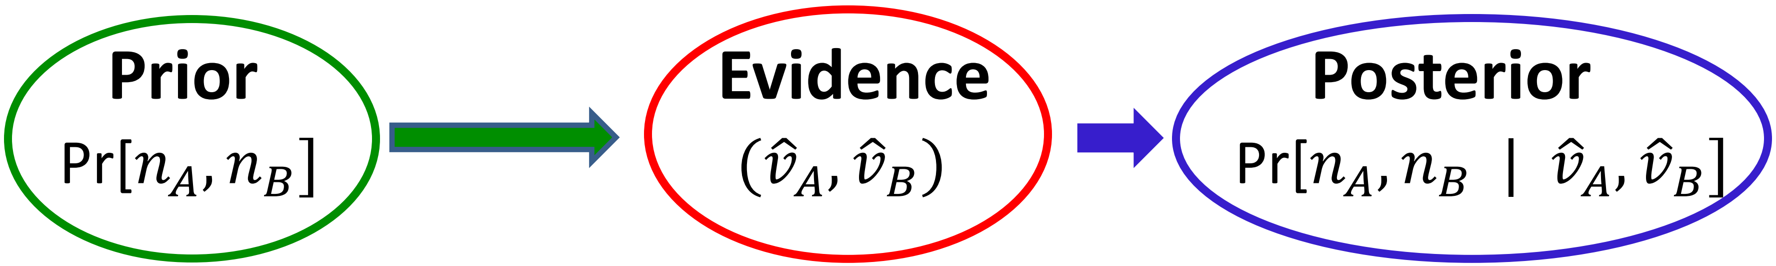

Requires a prior, which may be

- costly/impossible to obtain
- biased or open to manipulation

- Our approach: statistical hypothesis testing

Statistical hypothesis testing

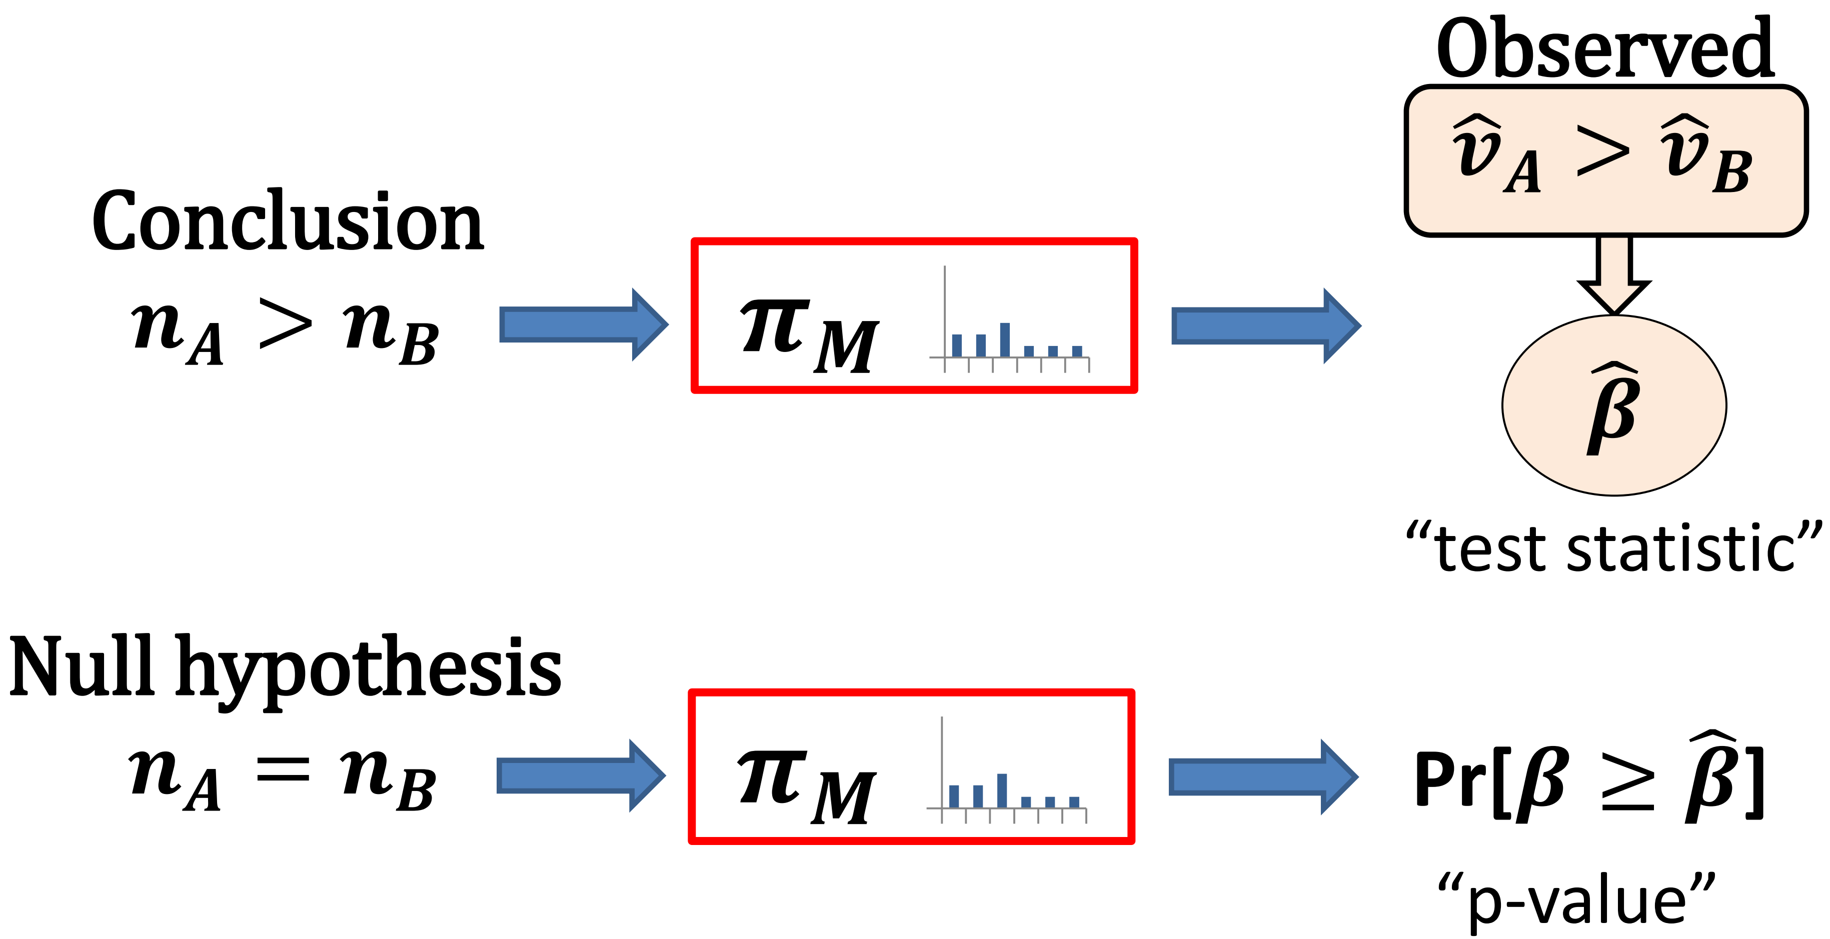

Statistical hypothesis testing

Conclusion

$n_A > n_B$ → π_M 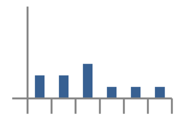 → **Observed** $\hat{\beta}$

Null hypothesis

$n_A = n_B$ → π_M 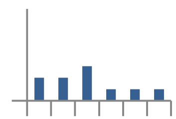 → **p-value** **Pr**$[\beta > \hat{\beta}]$

p-value > .05 → observed is not unlikely under null hypothesis → "accept" null

p-value < .05 → observed is unlikely under null hypothesis → reject null

Complication

Null hypothesis: $n_A = n_B = 1, 2, 3, 4, \cdots$

We can compute a p-value for each one.

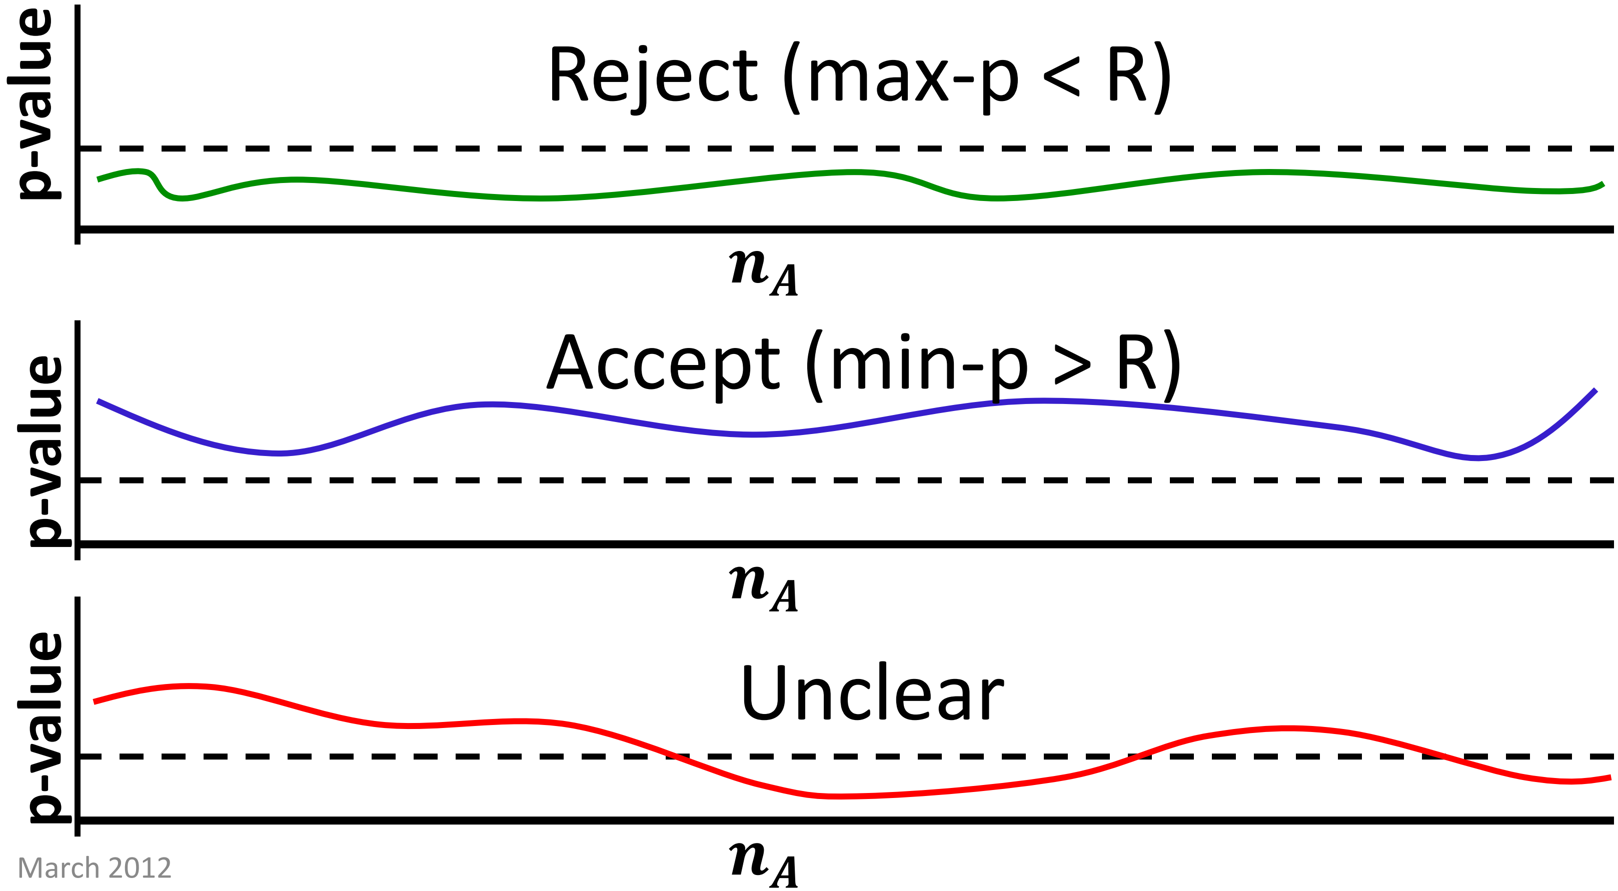

Our statistical test

Procedure:

1. Select significance level R (e.g. 0.05).
2. Observe votes $\hat{v}_A > \hat{v}_B$.
3. Compute $\hat{\beta}$.
4. If $\max\limits_{n_A = n_B} p$-value < R, reject.
5. If $\min\limits_{n_A = n_B} p$-value > R, don't reject.
6. Else, inconclusive whether to reject or not.

Example and picking a test statistic

Supporters	π_M	**Observed**
n_A (?) →	**False-name-limiting method M**	→ 92 = $\hat{v}_A$
n_B (?) →		→ 80 = $\hat{v}_B$

$$\beta(\hat{v}_A, \hat{v}_B) = ?$$

Selecting a test statistic

Observed: $\hat{v}_A = 92, \;\; \hat{v}_B = 80.$

Difference rule: $\hat{\beta} = \hat{v}_A - \hat{v}_B = 12$

Percent rule: $\hat{\beta} = \frac{\hat{v}_A - \hat{v}_B}{\hat{v}} \approx 0.07$

General form: $\hat{\beta} = \frac{\hat{v}_A - \hat{v}_B}{\hat{v}^{\alpha}} = \frac{12}{172^{\alpha}}$

(Adjusted margin of victory)

Test statistics that fail

Theorem 2.

Let the adjusted margin of victory be

$$\beta = \frac{\hat{v}_A - \hat{v}_B}{\hat{v}^{\alpha}}.$$

Then

1. *For any* $\alpha < 0.5$, max-p = ½: *we can never be sure to reject. (Type 2 errors)*
2. *For any* $\alpha > 0.5$, min-p = 0: *we can never be sure to "accept". (Type 1 errors)*

Test statistics for an election

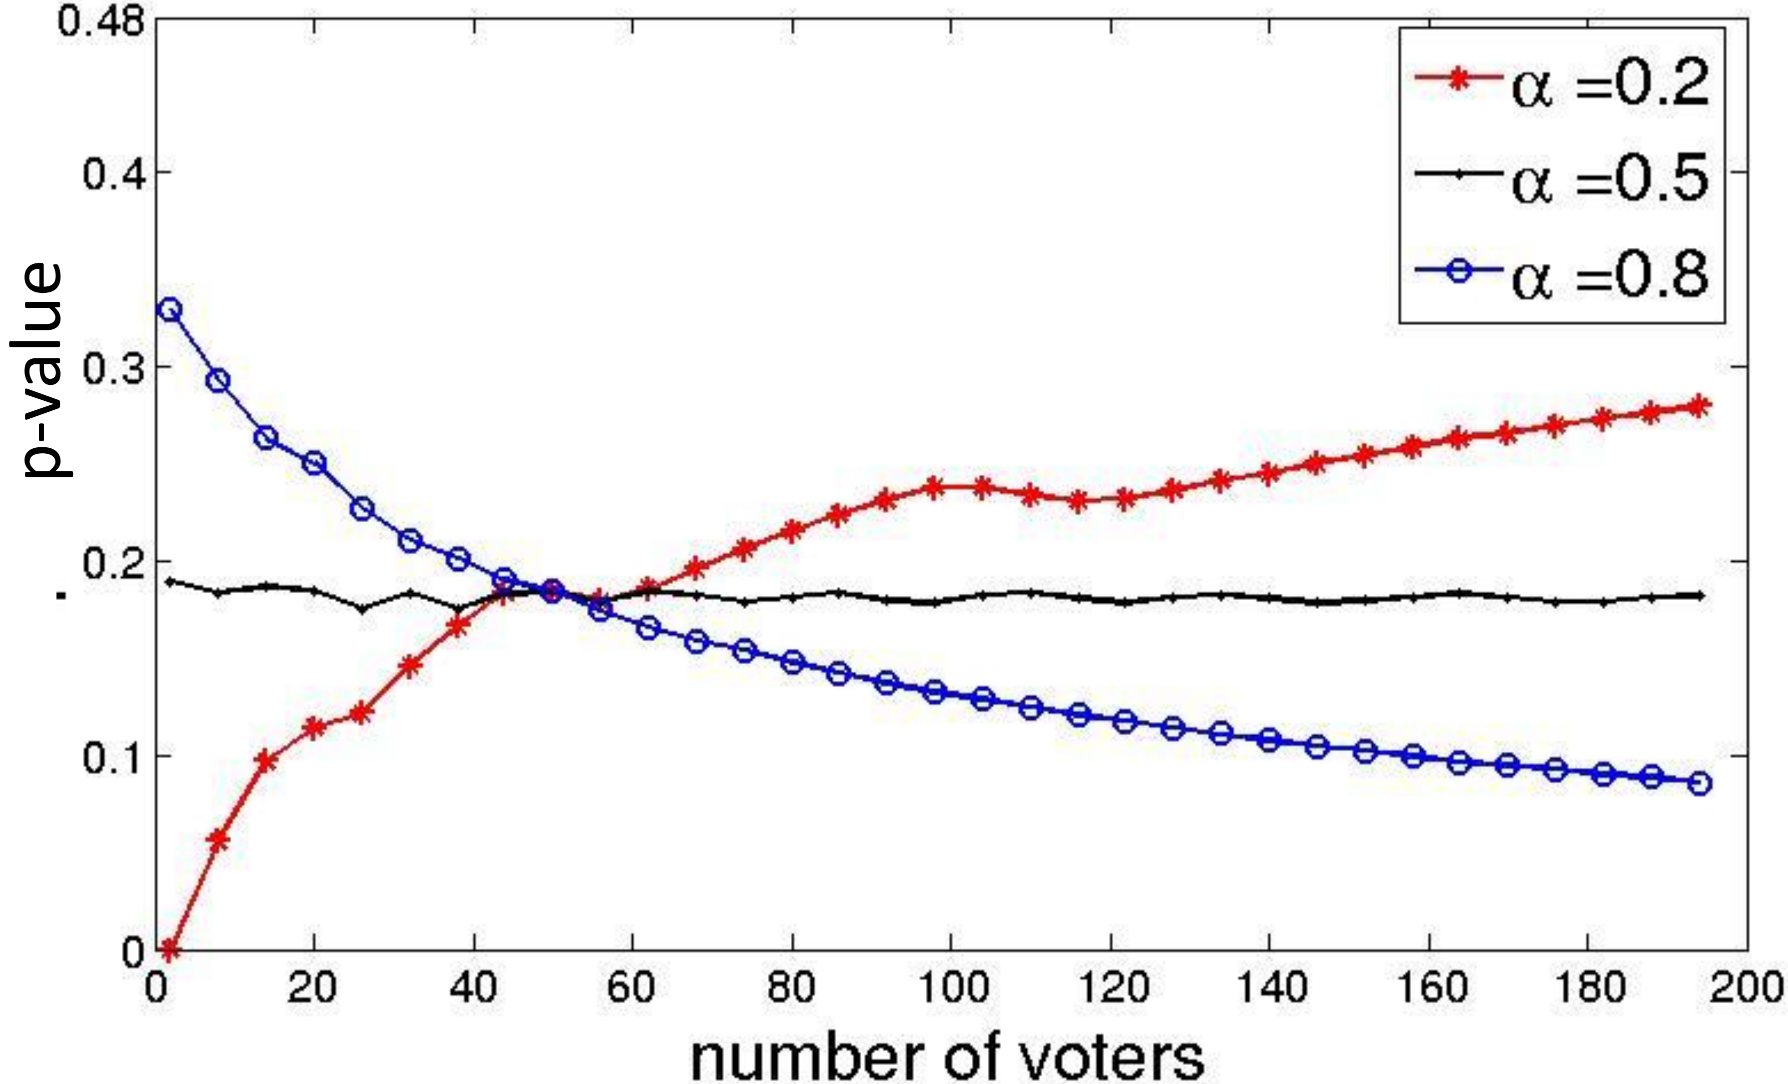

The "right" test statistic

Theorem 3.

Let the adjusted margin of victory formula be

$$\beta = \frac{\hat{v}_A - \hat{v}_B}{\hat{v}^{0.5}}.$$

Then

1. *For a large enough $\hat{\beta}$, we will reject. (Declare the outcome "correct".)*
2. *For a small enough $\hat{\beta}$, we will not reject. (Declare the outcome "inconclusive".)*

Test statistics for an election

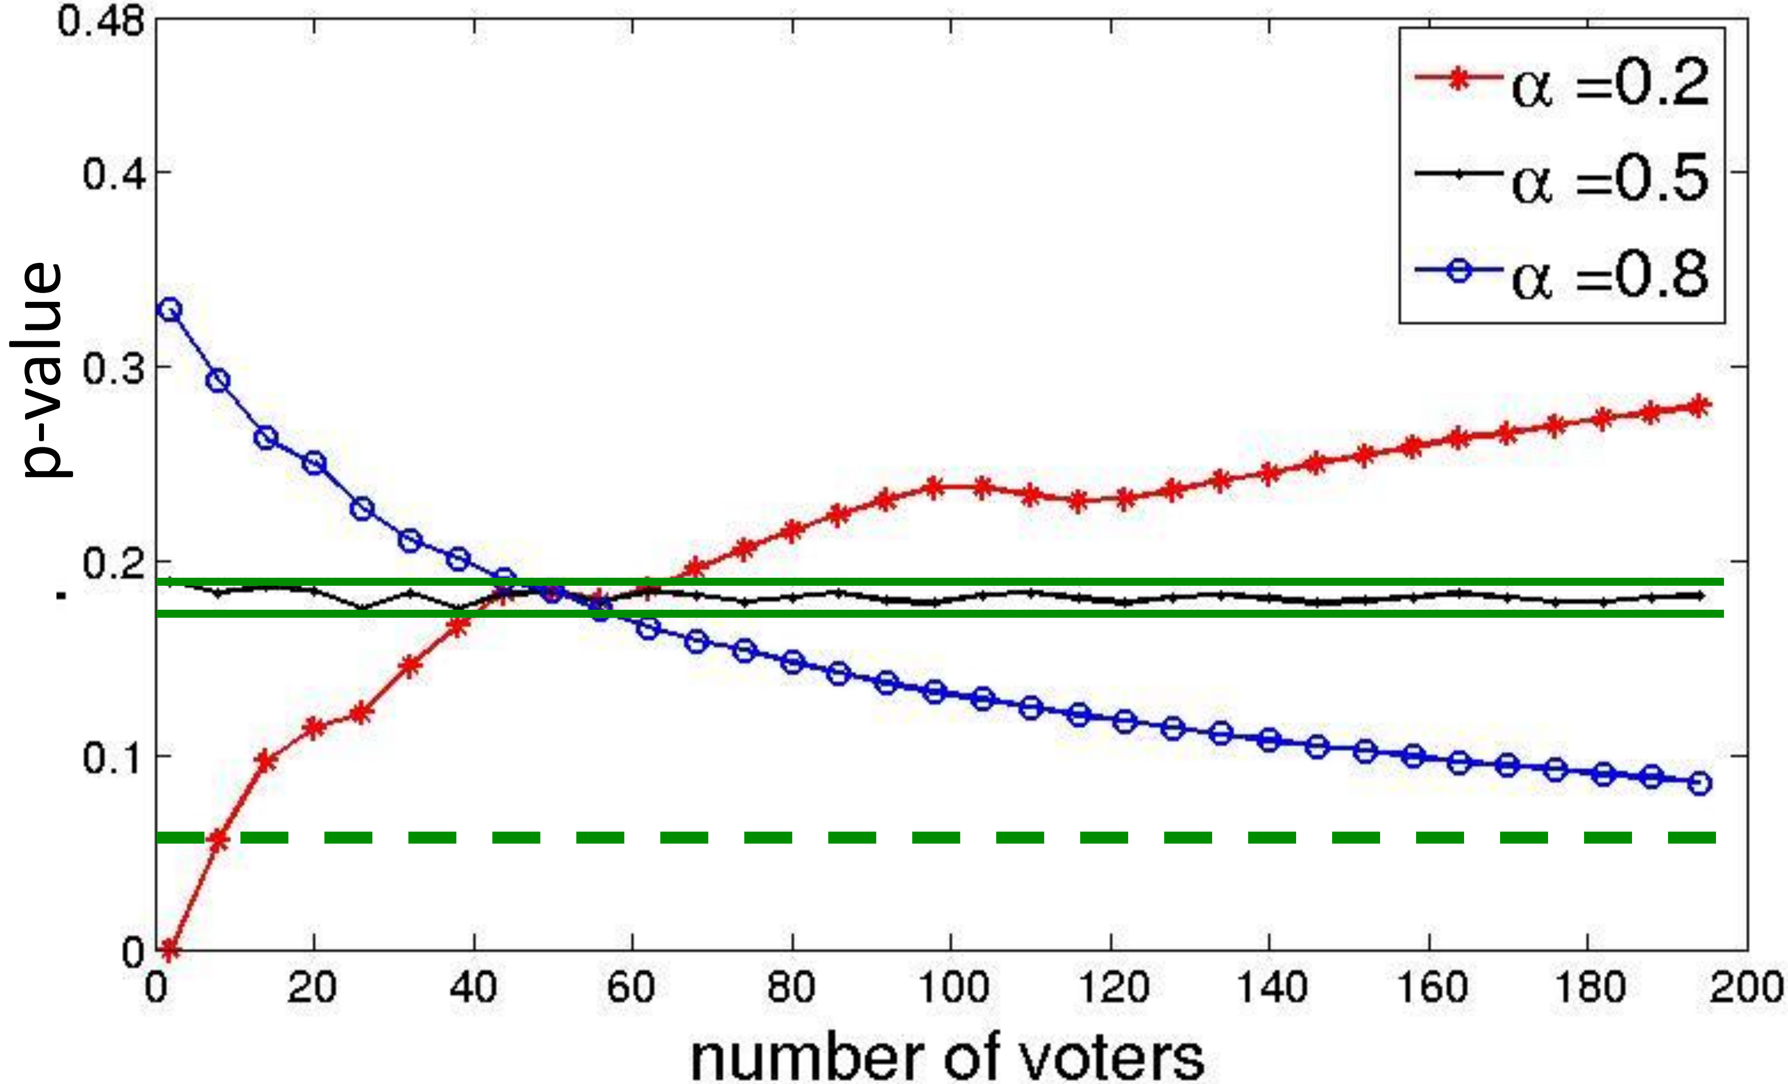

We can usually tell whether to reject or not

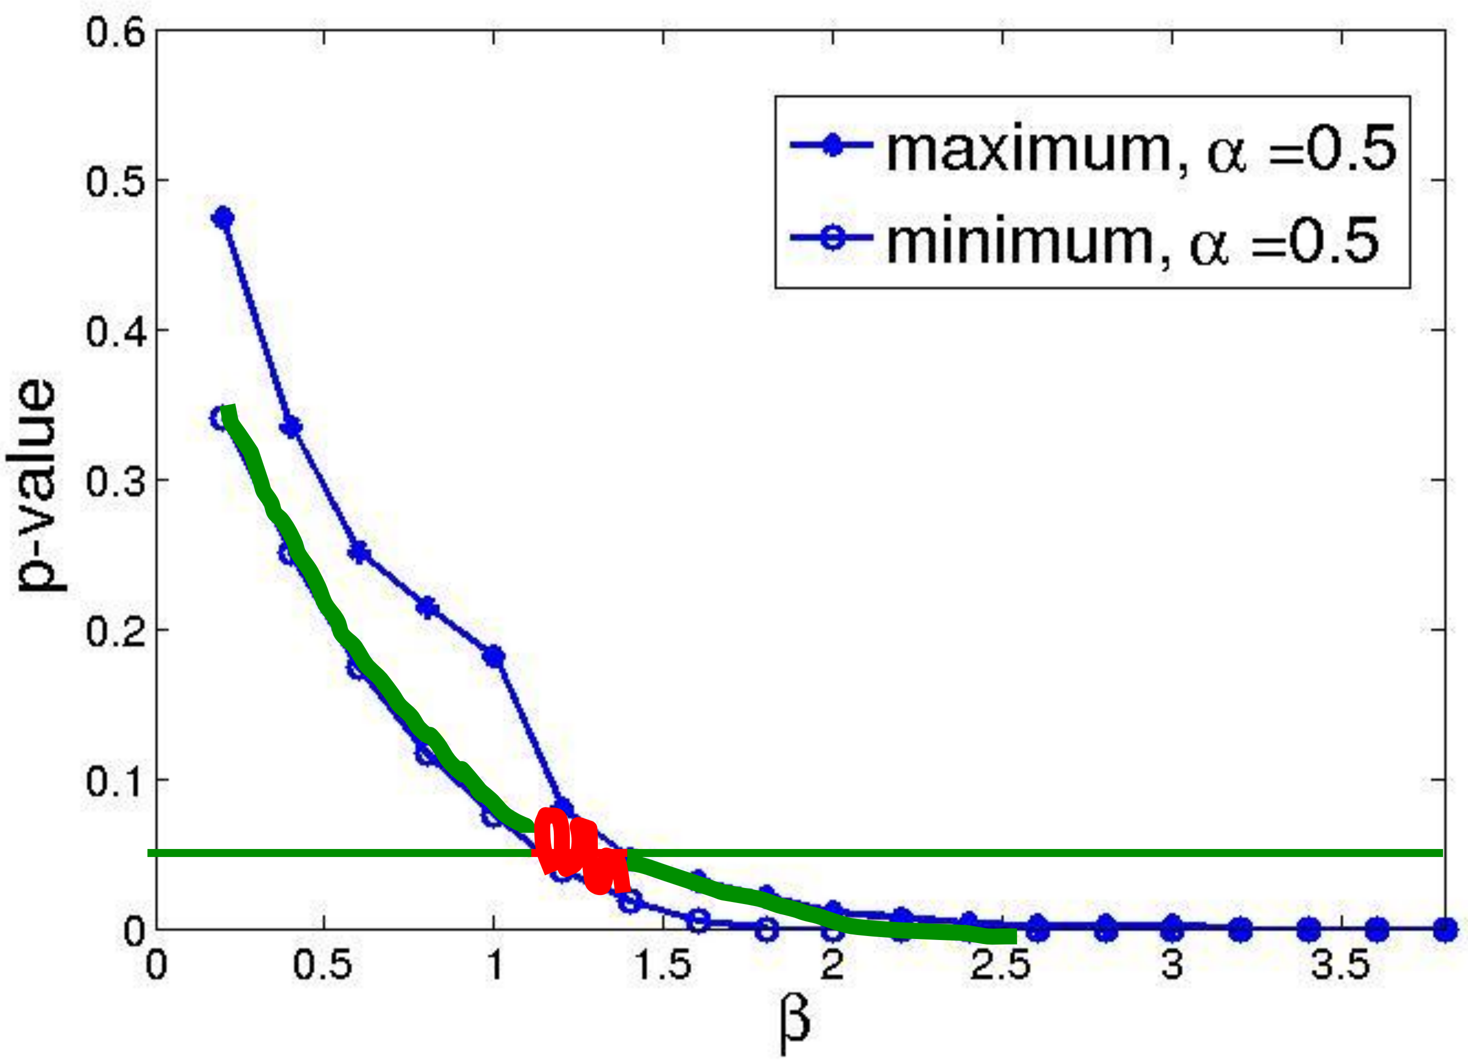

Use this test!

1. Select significance level R (e.g. 0.05).
2. Observe votes $\hat{v}_A > \hat{v}_B$.
3. Compute $\hat{\beta} = \frac{\hat{v}_A - \hat{v}_B}{\hat{\boldsymbol{v}}^{0.5}}$.
4. If $\max_{n_A = n_B}$ p-value < R, reject: high confidence.
5. If $\min_{n_A = n_B}$ p-value > R, don't: low confidence.
6. Else, inconclusive whether to reject or not.

   (rare!)

Outline

- Background and motivation: Why study elections in which we expect false-name votes?
- Our model
- How to **select** a false-name-limiting method?
- How to **evaluate** the election outcome?
- Recap and future work

Summary

- Model: take π as given, draw votes i.i.d.

- How to **select** a false-name-limiting method?

  A: Pick the method with the highest $\frac{\mu}{\sigma}$.

- How to **evaluate** the election outcome?

  A: Statistical significance test with

$$\hat{\beta} = \frac{\hat{v}_A - \hat{v}_B}{v^{0.5}}$$

  using max p-value and min p-value.

Future Work

- Single-peaked preferences (done)
- Application to real-world problems
- Other models or weaker assumptions
- How to actually produce distributions π?
  - Experimentally
  - Model agents and utilities

Thanks!