Some Game-Theoretic Aspects of Voting

Vincent Conitzer, Duke University

Conference on Web and Internet Economics (WINE), 2015

comsoc mailing list: https://lists.duke.edu/sympa/subscribe/comsoc

ec16

17th ACM CONFERENCE ON ECONOMICS AND COMPUTATION
JULY 24-28 2016 | MAASTRICHT, THE NETHERLANDS

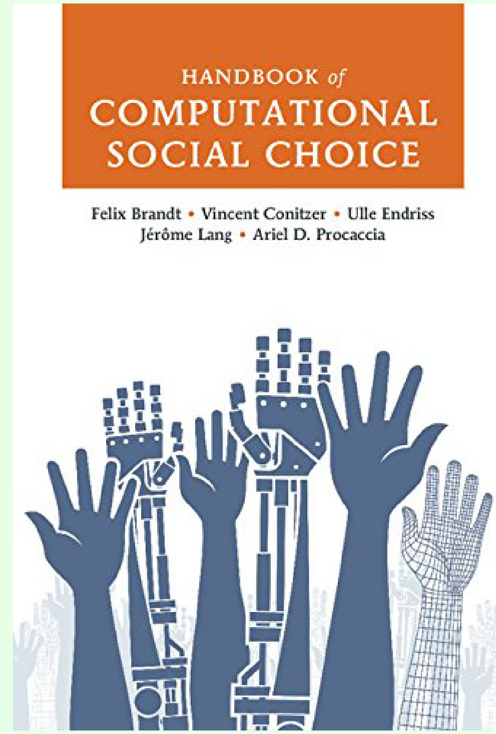

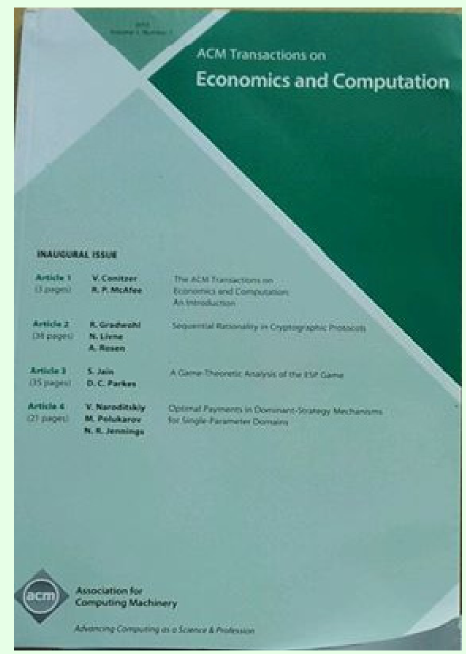

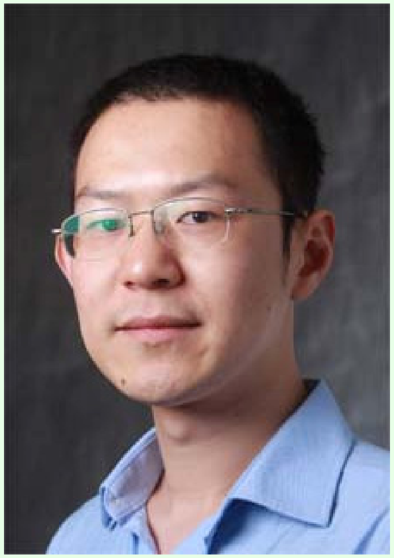

Lirong Xia
(Ph.D. 2011,
now at RPI)

Markus Brill
(postdoc 2013-
2015, now at
Oxford)

Rupert
Freeman
(Ph.D. student
2013 - ?)

Voting

n voters…

… each produce a ranking of *m* alternatives…

$b \succ a \succ c$

$a \succ c \succ b$

$a \succ b \succ c$

… which a social preference function (or simply voting rule) maps to one or more aggregate rankings.

$a \succ b \succ c$

Plurality

1 0 0

$b \succ a \succ c$

$a \succ c \succ b$

$a \succ b \succ c$

2 1 0

$a \succ b \succ c$

Borda

2 1 0

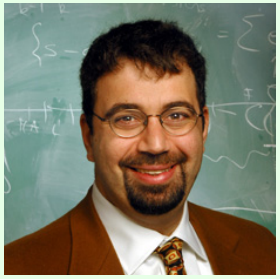

$b \succ a \succ c$

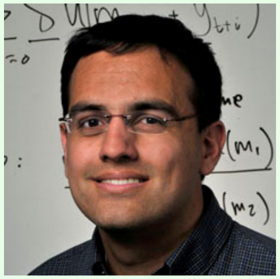

$a \succ c \succ b$

$a \succ b \succ c$

$a \succ b \succ c$

5 3 1

Kemeny

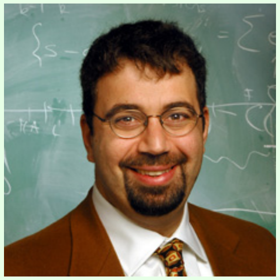

$b \succ a \succ c$

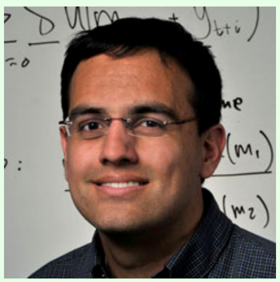

$a \succ c \succ b$

$a \succ b \succ c$

$a \succ b \succ c$

2 disagreements

$\leftrightarrow$

3*3 - 2 = 7 agreements (maximum)

- The unique SPF satisfying neutrality, consistency, and the Condorcet property [Young & Levenglick 1978]
- Natural interpretation as maximum likelihood estimate of the “correct” ranking [Young 1988, 1995]

Ranking Ph.D. applicants
(briefly described in C. [2010])

- Input: Rankings of subsets of the (non-eliminated) applicants

  $\succ$ 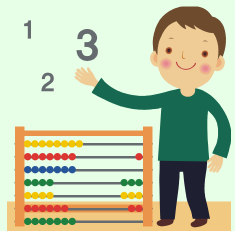 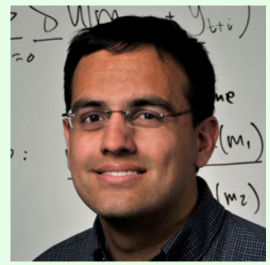 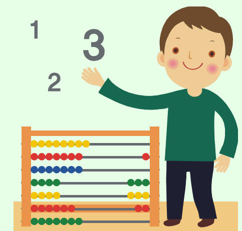 $\succ$ 

  $\succ$ 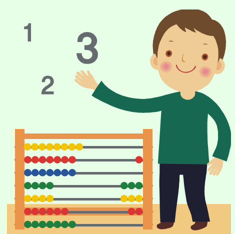 $\succ$ 

- Output: (one) Kemeny ranking of the (non-eliminated) applicants

 $\succ$ 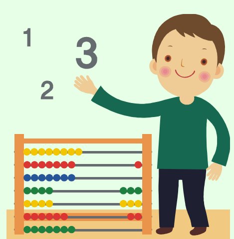 $\succ$ 

Instant runoff voting / single transferable vote (STV)

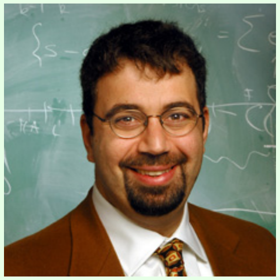

$b \succ a \succ c$

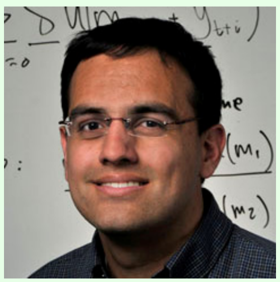

$a \succ b \succ b$

$a \succ b \succ c$

$a \succ b \succ c$

- The unique SPF satisfying: independence of bottom alternatives, consistency at the bottom, independence of clones (& some minor conditions) [Freeman, Brill, C. 2014]
- NP-hard to manipulate [Bartholdi & Orlin, 1991]

Manipulability

- Sometimes, a voter is better off revealing her preferences insincerely, aka. manipulating
- E.g., plurality
  - Suppose a voter prefers a > b > c
  - Also suppose she knows that the other votes are
    - 2 times b > c > a
    - 2 times c > a > b
  - Voting truthfully will lead to a tie between b and c
  - She would be better off voting, e.g., b > a > c, guaranteeing b wins

Gibbard-Satterthwaite impossibility theorem

- Suppose there are at least 3 alternatives
- There exists no rule that is simultaneously:
  - non-imposing/onto (for every alternative, there are some votes that would make that alternative win),
  - nondictatorial (there does not exist a voter such that the rule simply always selects that voter's first-ranked alternative as the winner), and
  - nonmanipulable/strategy-proof

Computational hardness as a barrier to manipulation

- A (successful) manipulation is a way of misreporting one's preferences that leads to a better result for oneself
- Gibbard-Satterthwaite only tells us that for some instances, successful manipulations exist
- It does not say that these manipulations are always easy to find
- Do voting rules exist for which manipulations are computationally hard to find?

A formal computational problem

- The simplest version of the manipulation problem:
- CONSTRUCTIVE-MANIPULATION:
  – We are given a voting rule *r*, the (unweighted) votes of the other voters, and an alternative *p*.
  – We are asked if we can cast our (single) vote to make *p* win.
- E.g., for the Borda rule:
  – Voter 1 votes A > B > C
  – Voter 2 votes B > A > C
  – Voter 3 votes C > A > B
- Borda scores are now: A: 4, B: 3, C: 2
- Can we make B win?
- Answer: YES. Vote B > C > A (Borda scores: A: 4, B: 5, C: 3)

Early research

- **Theorem.** CONSTRUCTIVE-MANIPULATION is NP-complete for the second-order Copeland rule. [Bartholdi, Tovey, Trick 1989]
  - Second order Copeland = alternative's score is sum of Copeland scores of alternatives it defeats
- **Theorem.** CONSTRUCTIVE-MANIPULATION is NP-complete for the STV rule. [Bartholdi, Orlin 1991]
- Most other rules are easy to manipulate (in P)

Ranked pairs rule [Tideman 1987]

- Order pairwise elections by decreasing strength of victory
- Successively "lock in" results of pairwise elections unless it causes a cycle

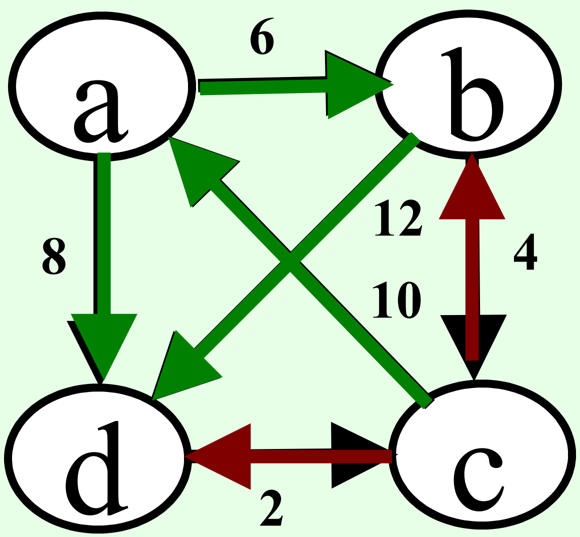

Final ranking:
c>a>b>d

- **Theorem.** CONSTRUCTIVE-MANIPULATION is NP-complete for the ranked pairs rule [Xia et al. IJCAI 2009]

Many manipulation problems…

	unweighted votes, constructive manipulation		weighted votes, constructive				weighted votes, destructive		
# alternatives			2	3	4	≥5	2	3	≥4
# manipulators	1	≥ 2							
plurality	P	P	P	P	P	P	P	P	P
plurality with runoff	P	P	P	NP-c	NP-c	NP-c	P	NP-c	NP-c
veto	P	P	P	NP-c	NP-c	NP-c	P	P	P
cup	P	P	P	P	P	P	P	P	P
Copeland	P	P	P	P	NP-c	NP-c	P	P	P
Borda	P	NP-c	P	NP-c	NP-c	NP-c	P	P	P
Nanson	NP-c	NP-c	P	P	NP-c	NP-c	P	P	NP-c
Baldwin	NP-c	NP-c	P	NP-c	NP-c	NP-c	P	NP-c	NP-c
Black	P	NP-c	P	NP-c	NP-c	NP-c	P	P	P
STV	NP-c	NP-c	P	NP-c	NP-c	NP-c	P	NP-c	NP-c
maximin	P	NP-c	P	P	NP-c	NP-c	P	P	P
Bucklin	P	P	P	NP-c	NP-c	NP-c	P	P	P
fallback	P	P	P	P	P	P	P	P	P
ranked pairs	NP-c	NP-c	P	P	P	NP-c	P	P	?
Schulze	P	P	P	P	P	P	P	P	P

Table from: C. & Walsh, Barriers to Manipulation, Chapter 6 in *Handbook of Computational Social Choice*

STV manipulation algorithm

[C., Sandholm, Lang JACM 2007]

Runs in
$O(((1+\sqrt{5})/2)^m)$ time
(**worst case**)

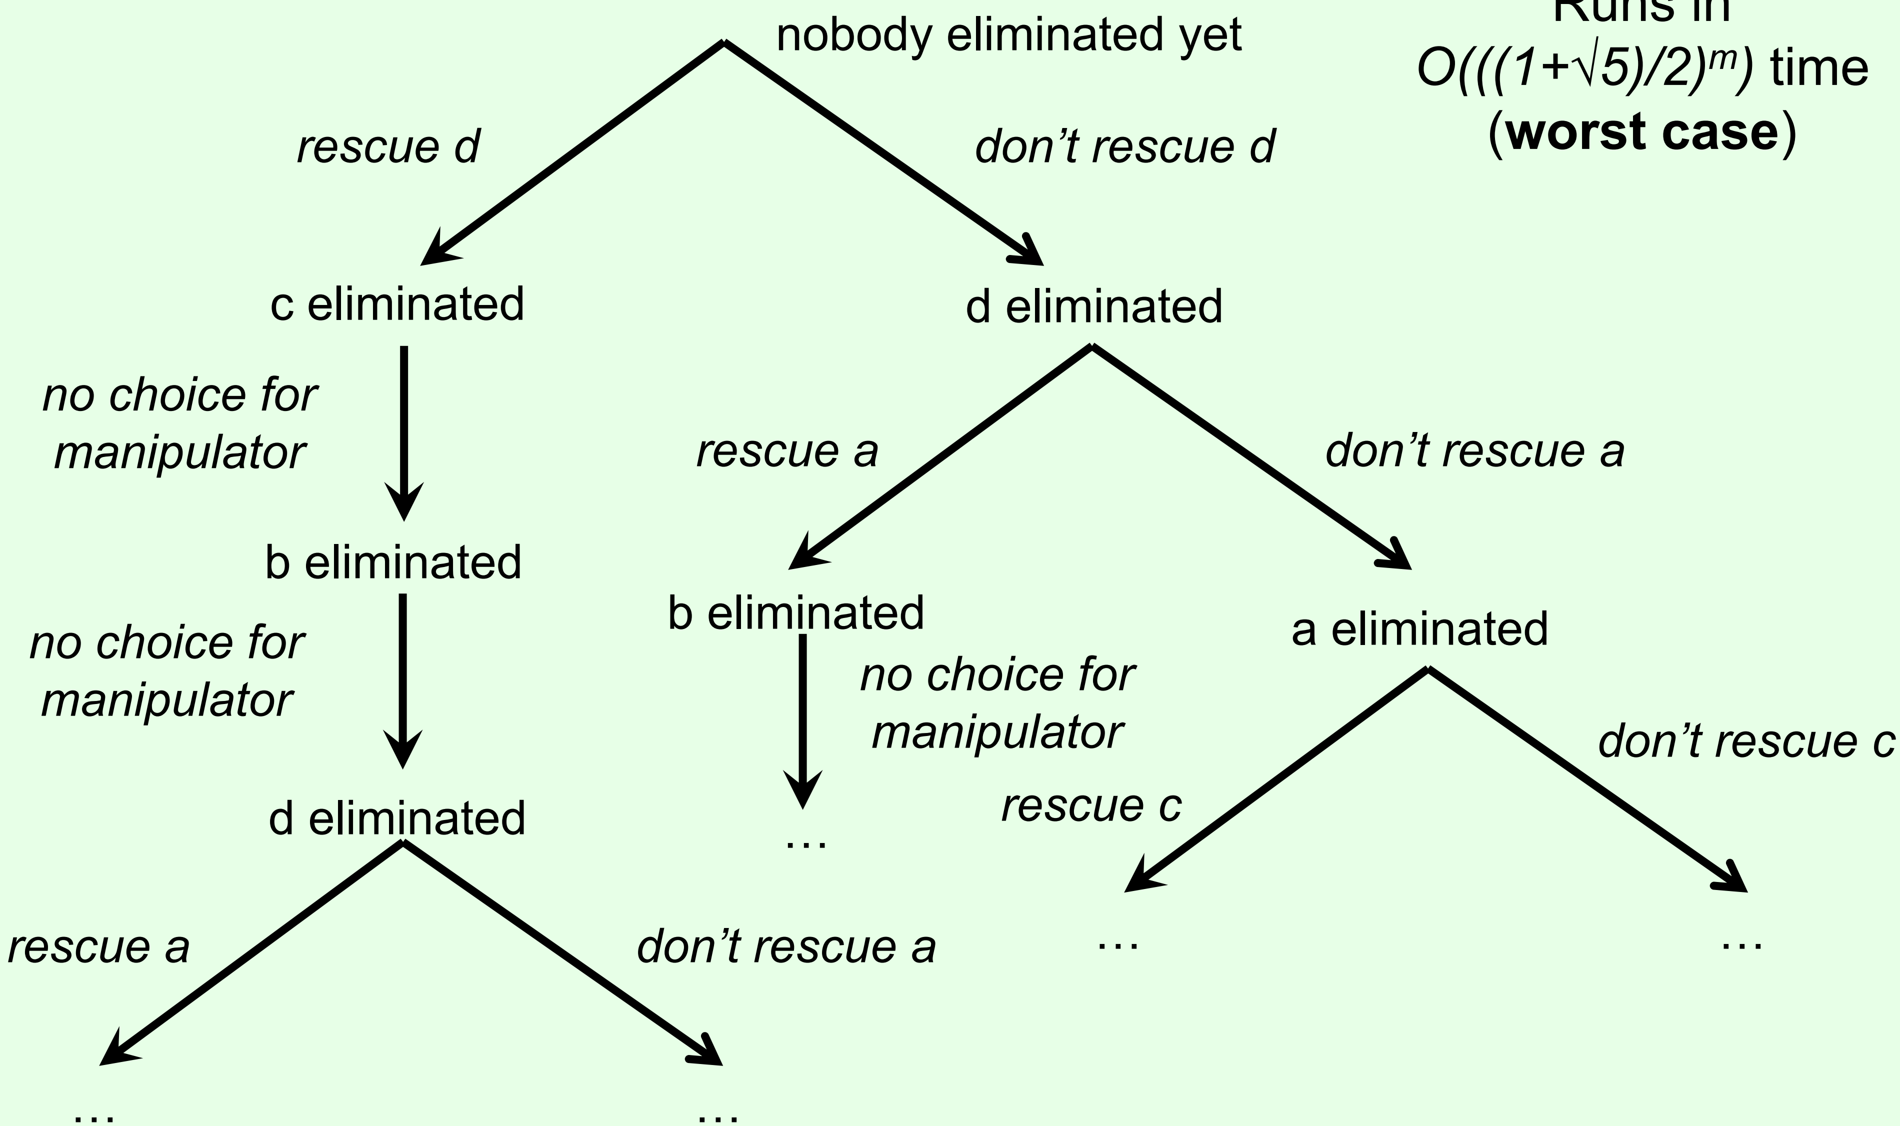

Runtime on random votes [Walsh 2011]

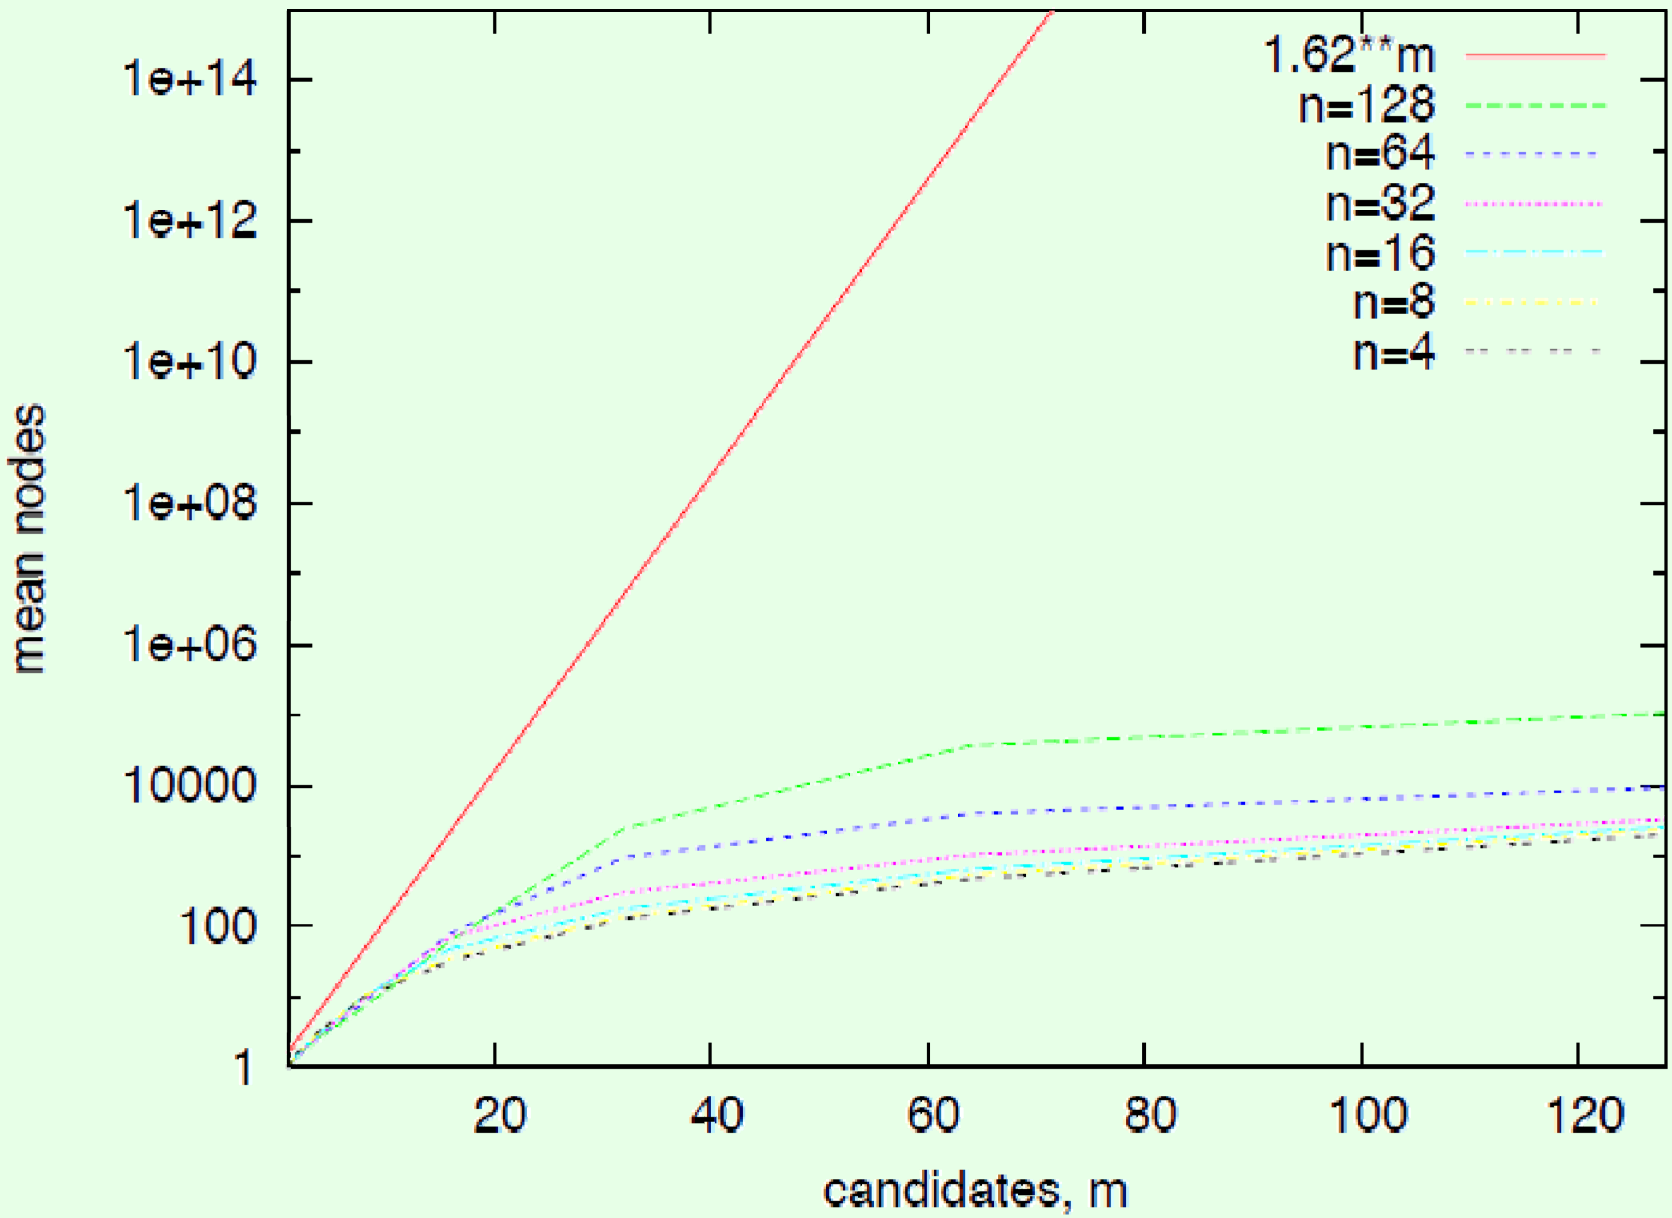

Fine – how about another rule?

- Heuristic algorithms and/or experimental (simulation) evaluation [C. & Sandholm 2006, Procaccia & Rosenschein 2007, Walsh 2011, Davies, Katsirelos, Narodytska, Walsh 2011]
- Quantitative versions of Gibbard-Satterthwaite showing that under certain conditions, for some voter, even a random manipulation on a random instance has significant probability of succeeding [Friedgut, Kalai, Nisan 2008; Xia & C. 2008; Dobzinski & Procaccia 2008; Isaksson, Kindler, Mossel 2010; Mossel & Racz 2013]

"for a social choice function f on $k \geq 3$ alternatives and n voters, which is ϵ-far from the family of nonmanipulable functions, a uniformly chosen voter profile is manipulable with probability at least inverse polynomial in n, k, and ϵ^{-1}."

Simultaneous-move voting games

- *Players:* Voters $1,\ldots,n$
- *Preferences:* Linear orders over alternatives
- *Strategies / reports:* Linear orders over alternatives
- *Rule:* $r(P')$, where P' is the reported profile

Voting: Plurality rule

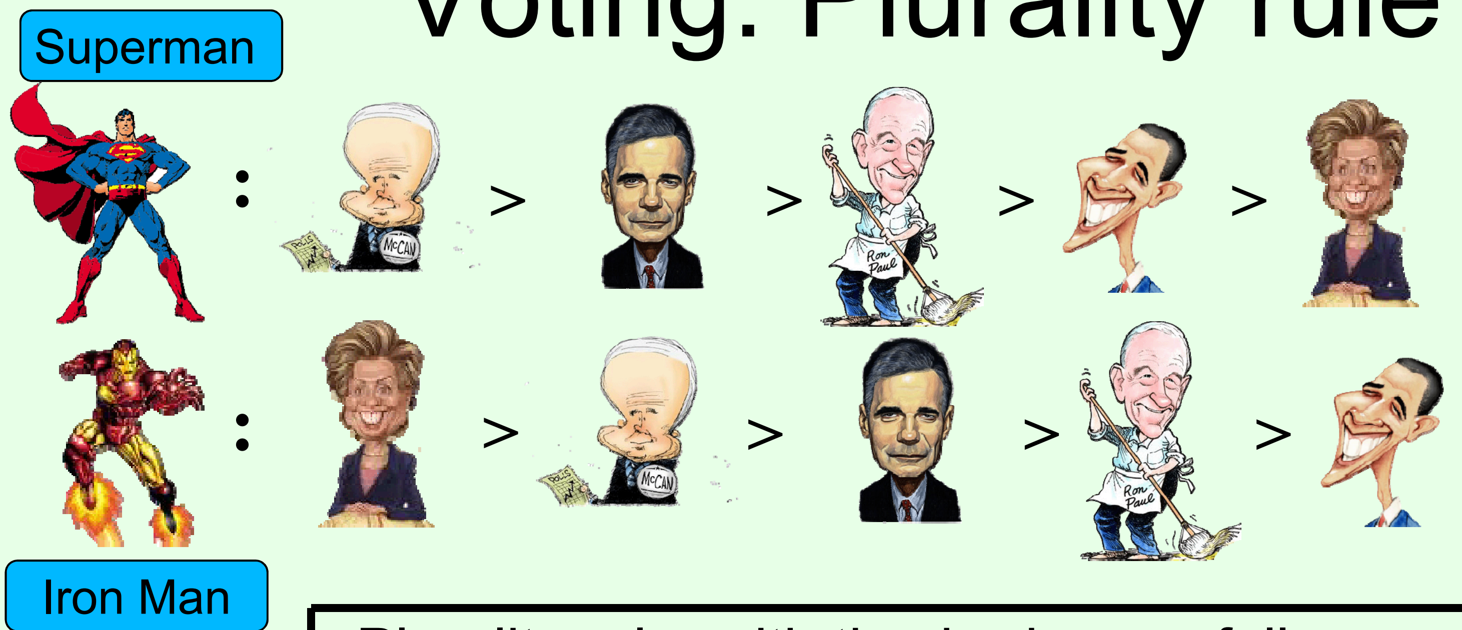

Plurality rule, with ties broken as follows:

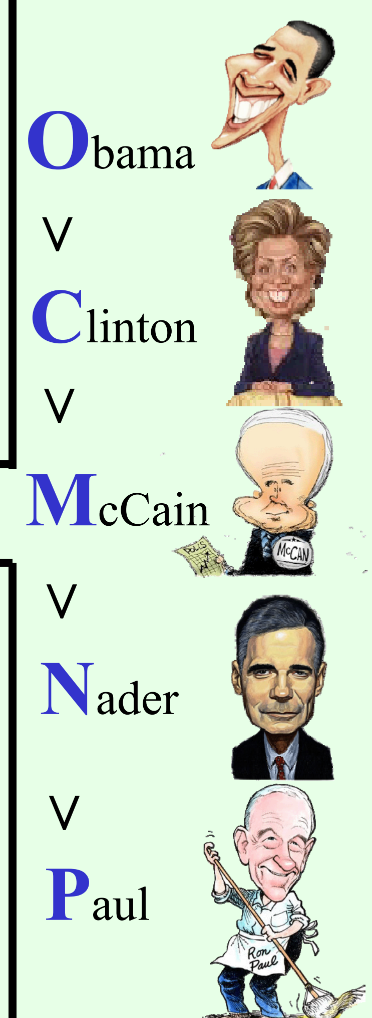

Many bad Nash equilibria…

- Majority election between alternatives a and b
  - Even if everyone prefers a to b, everyone voting for b is an equilibrium
  - Though, everyone has a weakly dominant strategy
- Plurality election among alternatives a, b, c
  - In equilibrium everyone might be voting for b or c, even though everyone prefers a!
- Equilibrium selection problem
- Various approaches: laziness, truth-bias, dynamics… [Desmedt and Elkind 2010, Meir et al. 2010, Thompson et al. 2013, Obraztsova et al. 2013, Elkind et al. 2015, …]

Voters voting sequentially

Duke CS TGIF* Movie Night

Do you plan to attend the next movie night?

[Yes, count me in!]

[Yes, count me in! (Vegetarian)]

Current count: 29

Current top films:

1. Inception
2. Eternal Sunshine of the Spotless Mind
3. Pulp Fiction

4 vote(s) *You have voted for this film.*

Title: Willow

Description: [link] This epic Lucasfilm fantasy serves up enough magical adventure to satisfy fans of the genre, though it treads familiar territory. With abundant parallels to Star Wars, the story (by George Lucas) follows the exploits of the little farmer Willow (Warwick Davis), an aspiring sorcerer appointed to deliver an infant princess from the evil queen (Jean Marsh) to whom the child is a crucial threat. Val Kilmer plays the warrior who joins Willow's campaign with the evil queen's daughter (Joanne Whalley, who later married Kilmer). Impressive production values, stunning locations (in England, Wales, and New Zealand) and dazzling special effects energize the routine fantasy plot, which alternates between rousing action and cute sentiment while failing to engage the viewer's emotions. A parental warning is appropriate: director Ron Howard has a light touch aimed at younger viewers, but doesn't shy away from grisly swordplay and at least one monster (a wicked two-headed dragon) that could induce nightmares.

Trailer: http://matttrailer.com/willow_1988

10 vote(s) *You have voted for this film.*

Title: Pulp Fiction

Description: [link] http://www.youtube.com/watch?v=0AHETuK70Sc The lives of two mob hit men, a boxer, a gangster's wife, and a pair of diner bandits intertwine in four tales of violence and redemption

3 vote(s) *You have voted for this film.*

Title: Tom yum goong

Description: [link] In Bangkok, the young Kham was raised by his father in the jungle with elephants as members of their family. When his old elephant and the baby Kern are stolen by criminals, Kham finds that the animals were sent to Sidney. He travels to Australia, where he locates the baby elephant in a restaurant owned by the evil Madame Rose, the leader of an international Thai mafia. With the support of the efficient Thai sergeant Mark, who was involved in a conspiracy, Kham fights to rescue the animal from the mobsters.

2 vote(s) *You have voted for this*

Title: Dogville

Description: [link] Dogville is a 2003 philosophical drama written and directed by Lars von Trier, and starring Nicole Kidman. It is a parable that uses an extremely minimal, stage-like set to tell the story of Grace Mulligan (Kidman), a woman hiding from mobsters, who arrives in the small mountain town of Dogville and is provided refuge in return for physical labor. Because she has to win and keep the acceptance of every single one of the inhabitants of the town to be allowed to stay, any attempt by her to do things her own way or to put a limit on her service risks driving her back out into the

Our setting

- Voters vote sequentially and strategically
  - voter $1 \rightarrow$ voter $2 \rightarrow$ voter $3 \rightarrow \ldots$ etc
  - states in stage i: all possible profiles of voters $1,\ldots,i$-1
  - any terminal state is associated with the winner under rule r
- At any stage, the current voter knows
  - the order of voters
  - previous voters' votes
  - true preferences of the later voters (complete information)
  - rule r used in the end to select the winner
- We call this a Stackelberg voting game
  - Unique winner in SPNE (not unique SPNE)
  - the subgame-perfect winner is denoted by $SG_r(P)$, where P consists of the true preferences of the voters

Voting: Plurality rule

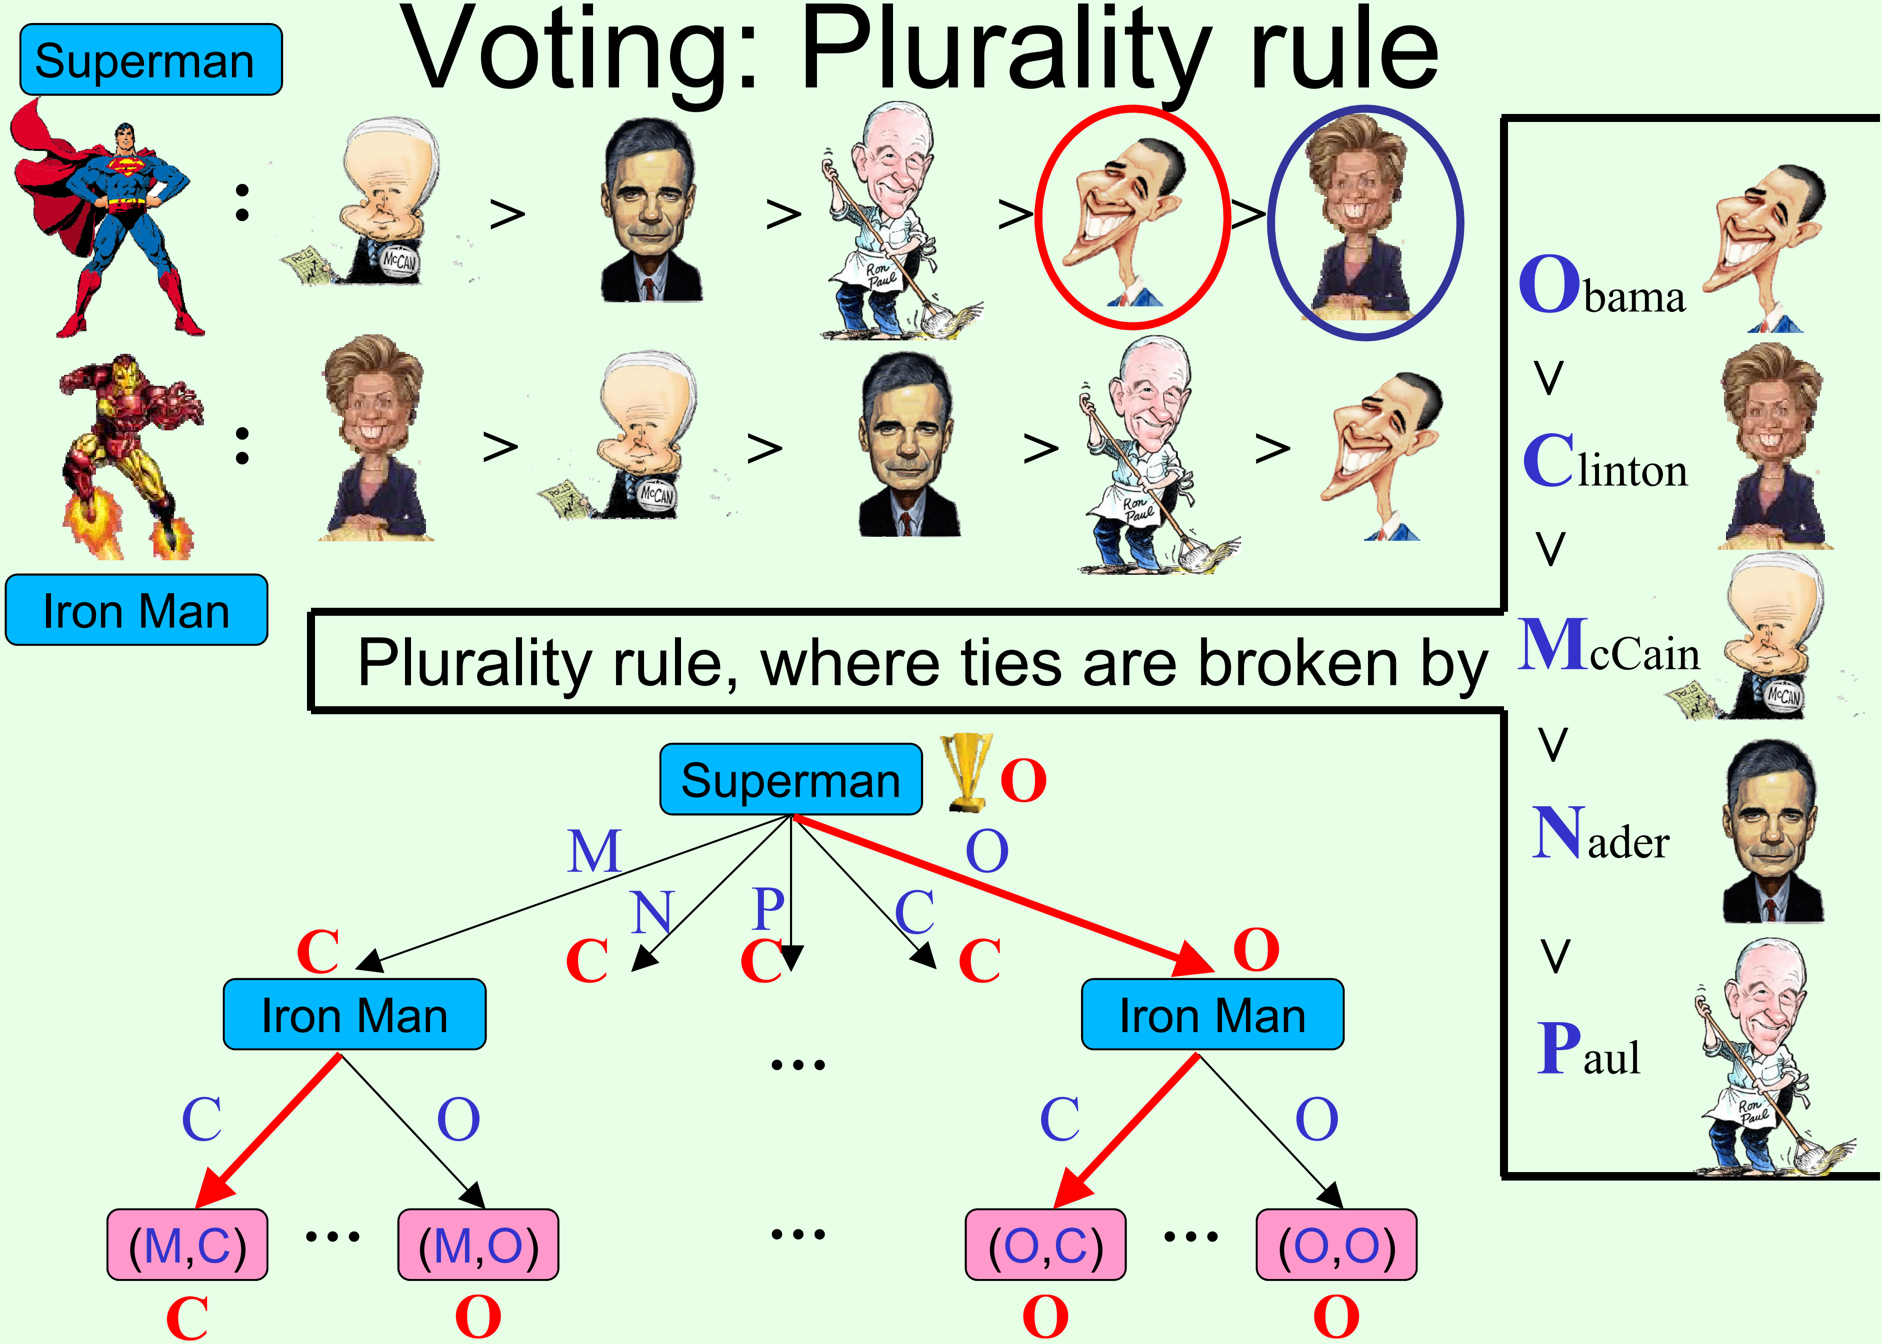

Literature

- Voting games where voters cast votes one after another
  - [Sloth GEB-93, Dekel and Piccione JPE-00, Battaglini GEB-05, Desmedt & Elkind EC-10]

Key questions

- How can we compute the backward-induction winner efficiently (for general voting rules)?
- How good/bad is the backward-induction winner?

Computing $SG_r(P)$

- Backward induction:
  - A state in stage i corresponds to a profile for voters 1, …, i-1
  - For each state (starting from the terminal states), we compute the winner if we reach that point
- Making the computation more efficient:
  - depending on r, some states are equivalent
  - can merge these into a single state
  - drastically speeds up computation

An equivalence relationship between profiles

- The plurality rule
- 160 voters have cast their votes, 20 voters remaining

50 votes $x>y>z$		31 votes $x>y>z$
30 votes $x>z>y$	=	21 votes $y>z>x$
70 votes $y>x>z$		0 votes $z>y>x$
10 votes $z>x>y$		
(80, 70, 10)		(31, 21, 0)
↑ ↑ ↑		↑ ↑ ↑
x y z		x y z

- This equivalence relationship is captured in a concept called *compilation complexity* [Chevaleyre et al. IJCAI-09, Xia & C. AAAI-10]

Paradoxes

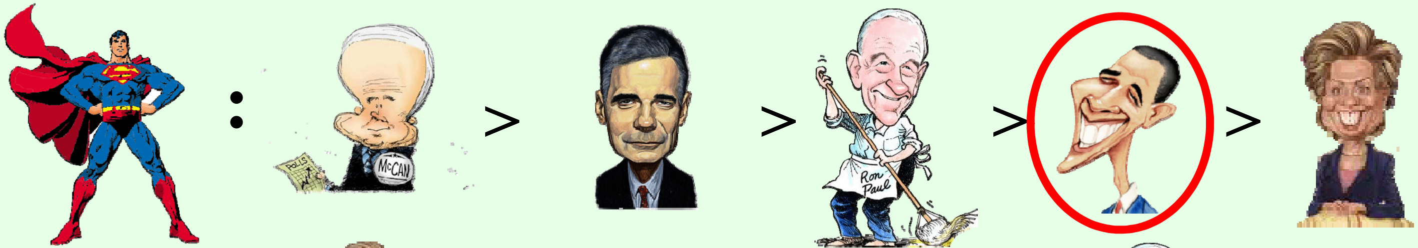

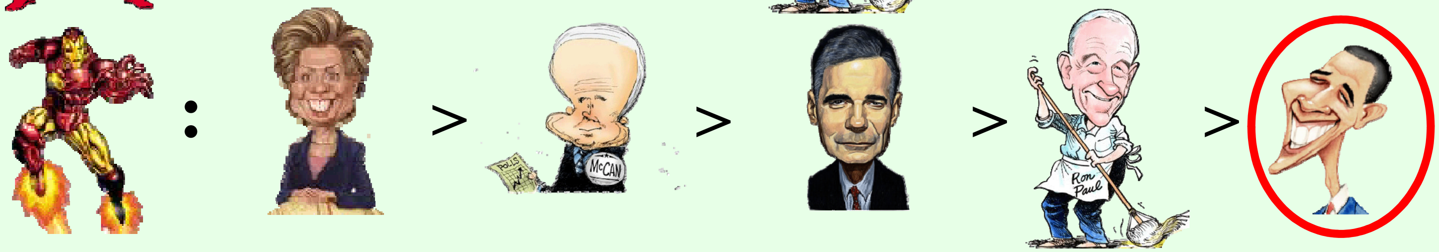

- Plurality rule, where ties are broken according to

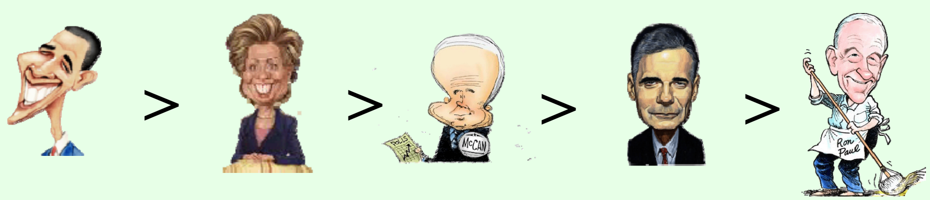

- The SG_{Plu} winner is 
- Paradox: the SG_{Plu} winner is ranked almost in the bottom position in all voters' true preferences

What causes the paradox?

- Q: Is it due to defects in the plurality rule / tiebreaking scheme, or it is because of the strategic behavior?
- A: The strategic behavior!
  - by showing a ubiquitous paradox

Domination index

- For any voting rule r, the **domination index** of r when there are n voters, denoted by $\mathrm{DI}_r(n)$, is:
- the smallest number k such that for any alternative c, any coalition of $n/2+k$ voters can guarantee that c wins.
  - The DI of any **majority consistent** rule r is 1, including any Condorcet-consistent rule, plurality, plurality with runoff, Bucklin, and STV
  - The DI of any positional scoring rule is no more than $n/2-n/m$
  - Defined for a voting rule (not for the voting game using the rule)
  - Closely related to the **anonymous veto function** [Moulin 91]

Main theorem (ubiquity of paradox)

- **Theorem**: For any voting rule r and any n, there exists an n-profile P such that:
  - (*many voters are miserable*) $SG_r(P)$ is ranked somewhere in the bottom two positions in the true preferences of $n\text{-}2\cdot \mathrm{DI}_r(n)$ voters
  - (*almost Condorcet loser*) if $\mathrm{DI}_r(n) < n/4$, then $SG_r(P)$ loses to all but one alternative in pairwise elections.

Proof

- **Lemma:** Let P be a profile. An alternative d is **not** the winner $SG_r(P)$ if there exists another alternative c and a subprofile $P_k = (V_{i_1}, \ldots, V_{i_k})$ of P that satisfies the following conditions: (1) $k \geq \lfloor n/2 \rfloor + \mathrm{DI}_r(n)$, (2) $c > d$ in each vote in P_k, (3) for any $1 \leq x < y \leq k$, $\mathrm{Up}(V_{i_x}, c) \supseteq \mathrm{Up}(V_{i_y}, c)$, where $\mathrm{Up}(V_{i_x}, c)$ is the set of alternatives ranked higher than c in V_{i_x}

$$V_1 = \ldots = V_{\lfloor n/2 \rfloor - \mathrm{DI}_r(n)} = [c_3 > \ldots > c_m > c_1 > c_2]$$

$$V_{\lfloor n/2 \rfloor - \mathrm{DI}_r(n)+1} = \ldots = V_{\lfloor n/2 \rfloor + \mathrm{DI}_r(n)} = [c_1 > c_2 > c_3 > \ldots > c_m]$$

$$V_{\lfloor n/2 \rfloor + \mathrm{DI}_r(n)+1} = \ldots = V_n = [c_2 > c_3 > \ldots > c_m > c_1]$$

- c_2 is not a winner (letting $c = c_1$ and $d = c_2$ in the lemma)
- For any $i \geq 3$, c_i is not a winner (letting $c = c_2$ and $d = c_i$ in the lemma)

What do these paradoxes mean?

- These paradoxes state that for any rule r that has a low domination index, *sometimes* the backward-induction outcome of the Stackelberg voting game is undesirable
  - the DI of any majority consistent rule is 1
- Worst-case result
- Surprisingly, on average (by simulation)
  - # { voters who prefer the SG_r winner to the truthful r winner}

    > # { voters who prefer the truthful r winner to the SG_r winner}

Simulation results

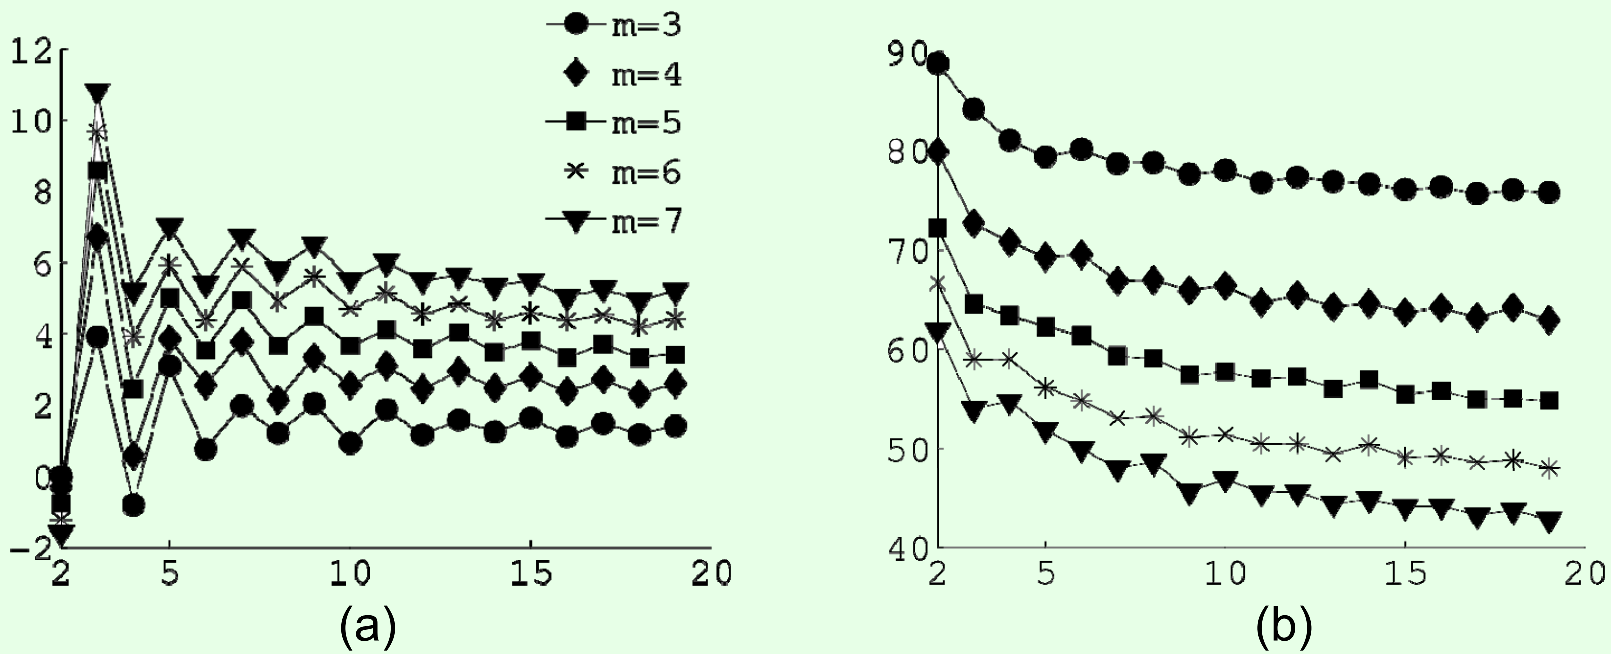

- Simulations for the plurality rule (25000 profiles uniformly at random)
  - x-axis is #voters, y-axis is the percentage of voters
  - (a) percentage of voters where $SG_r(P) > r(P)$ minus percentage of voters where $r(P) > SG_r(P)$
  - (b) percentage of profiles where the $SG_r(P) = r(P)$
- SG_r winner is preferred to the truthful r winner by more voters than vice versa
  - Whether this means that SG_r is "better" is debatable

Ph.D. applicants may be substitutes or complements…

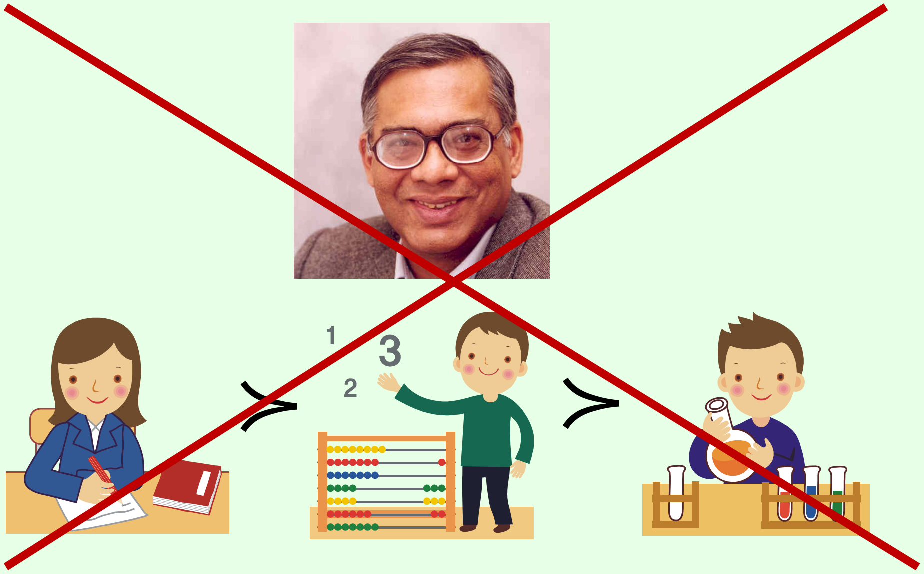

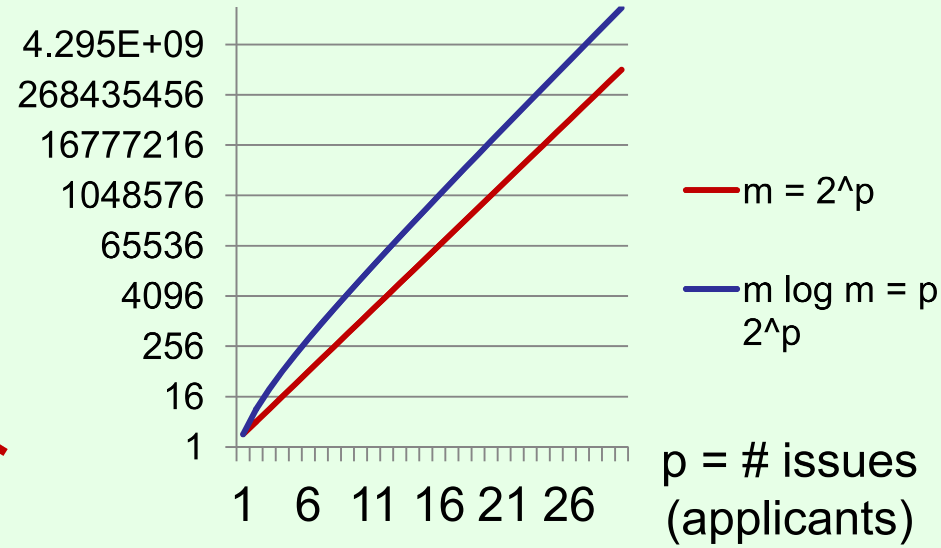

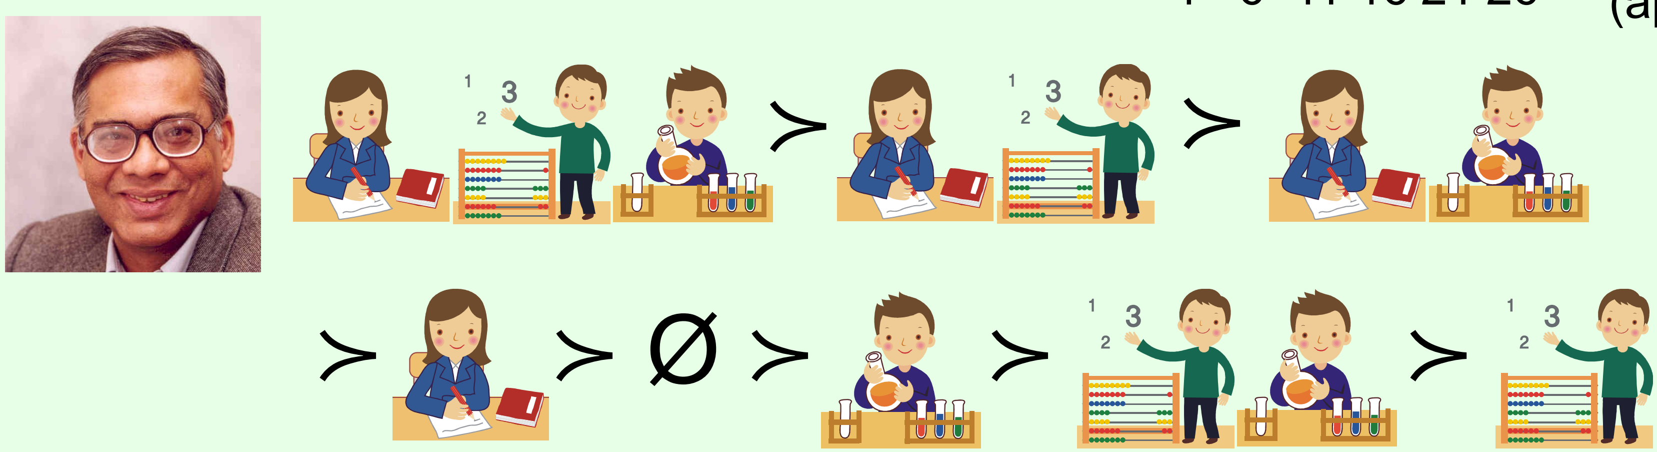

Sequential voting

see Lang & Xia [2009]

- Issues: main dish, wine
- Order: main dish > wine
- Local rules are majority rules
- V_1:  $\succ$  ,  :  $\succ$  ,  :  $\succ$ 
- V_2:  $\succ$  ,  :  $\succ$  ,  :  $\succ$ 
- V_3:  $\succ$  ,  :  $\succ$  ,  :  $\succ$ 
- **Step 1:** 
- **Step 2: given**  **,**  **is the winner for wine**
- **Winner: (**  **,**  **)**

- Xia, C., Lang [2008, 2010, 2011] study rules that do not require preferences to have this structure

Sequential voting and strategic voting

S **T**

$V_1 : st > \bar{s}t > s\bar{t} > \bar{s}\bar{t}$

$V_2 : s\bar{t} > st > \bar{s}t > \bar{s}\bar{t}$

$V_3 : \bar{s}t > \bar{s}\bar{t} > s\bar{t} > st$

- In the first stage, the voters vote simultaneously to determine **S**; then, in the second stage, the voters vote simultaneously to determine **T**
- If **S** is built, then in the second step $t > \bar{t}$, $\bar{t} > t$, $\bar{t} > t$ so the winner is $s\bar{t}$
- If **S** is **not** built, then in the 2nd step $t > \bar{t}$, $t > \bar{t}$, $t > \bar{t}$ so the winner is $\bar{s}t$
- In the first step, the voters are effectively comparing $s\bar{t}$ and $\bar{s}t$, so the votes are $\bar{s} > s$, $s > \bar{s}$, $\bar{s} > s$, and the final winner is $\bar{s}t$

[Xia, C., Lang 2011; see also Farquharson 1969, McKelvey & Niemi 1978, Moulin 1979, Gretlein 1983, Dutta & Sen 1993]

Strategic sequential voting (SSP)

- Binary issues (two possible values each)
- Voters vote simultaneously on issues, one issue after another according to O
- For each issue, the majority rule is used to determine the value of that issue
- Game-theoretic aspects:
  - A complete-information extensive-form game
  - The winner is unique

Voting tree

- The winner is the same as the (truthful) winner of the following *voting tree*

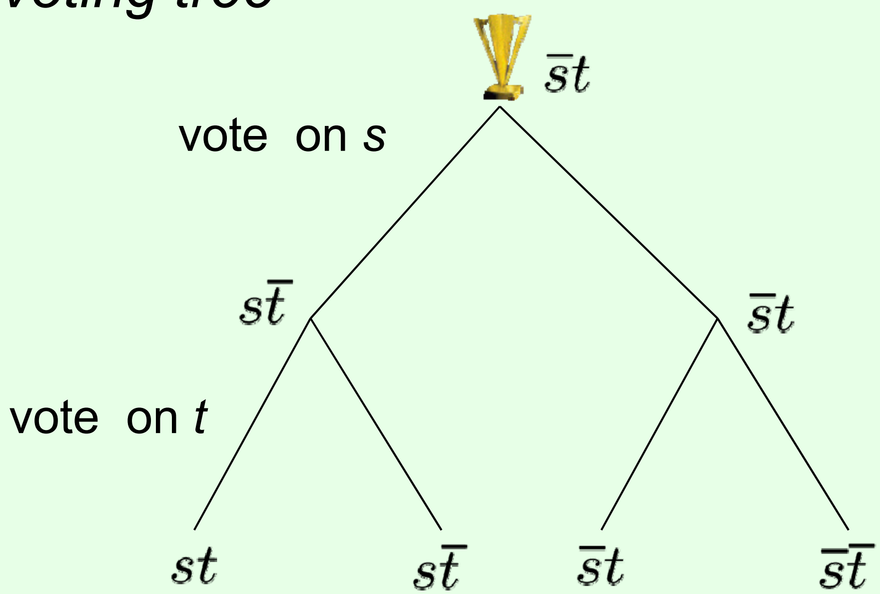

- "Within-state-dominant-strategy-backward-induction"
- Similar relationships between backward induction and voting trees have been observed previously [McKelvey&Niemi JET 78], [Moulin Econometrica 79], [Gretlein IJGT 83], [Dutta & Sen SCW 93]

Paradoxes [Xia, C., Lang EC 2011]

- Strong paradoxes for strategic sequential voting (SSP)
- Slightly weaker paradoxes for SSP that hold for any O (the order in which issues are voted on)
- Restricting voters' preferences to escape paradoxes
- Other multiple-election paradoxes:

  [Brams, Kilgour & Zwicker SCW 98], [Scarsini SCW 98], [Lacy & Niou JTP 00], [Saari & Sieberg 01 APSR], [Lang & Xia MSS 09]

Multiple-election paradoxes for SSP

- **Main theorem** (informally). For any $p \geq 2$ and any $n \geq 2p^2 + 1$, there exists an n-profile such that the SSP winner is
  – Pareto dominated by almost every other candidate
  – ranked almost at the bottom (exponentially low positions) in every vote
  – an almost Condorcet loser

Is there any better choice of the order O?

- **Theorem** (informally). For any $p \geq 2$ and $n \geq 2^{p+1}$, there exists an n-profile such that for **any** order O over $\{x_1, \ldots, x_p\}$, the SSP_O winner is ranked somewhere in the bottom $p+2$ positions.
  - The winner is ranked almost at the bottom in every vote
  - The winner is still an almost Condorcet loser
  - I.e., at least some of the paradoxes cannot be avoided by a better choice of O

Getting rid of the paradoxes

- **Theorem(s)** (informally)
  - Restricting the preferences to be separable or lexicographic gets rid of the paradoxes
  - Restricting the preferences to be O-legal does not get rid of the paradoxes

Agenda control

- **Theorem.** For any $p \geq 4$, there exists a profile P such that any alternative can be made to win under this profile by changing the order O over issues
  - The chair has full power over the outcome by agenda control (for this profile)

Crowdsourcing societal tradeoffs

[C., Brill, Freeman AAMAS'15 Blue Sky track; C., Freeman, Brill, Li AAAI'16]

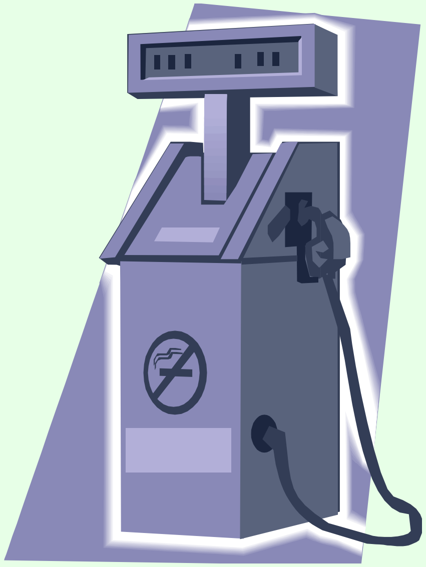

1 bag of landfill trash *is as bad as* using *x* gallons of gasoline

How to determine x?

- Other examples: clearing an acre of forest, fishing a ton of bluefin tuna, causing the average person to sit in front of a screen for another 5 minutes a day, …

A challenge

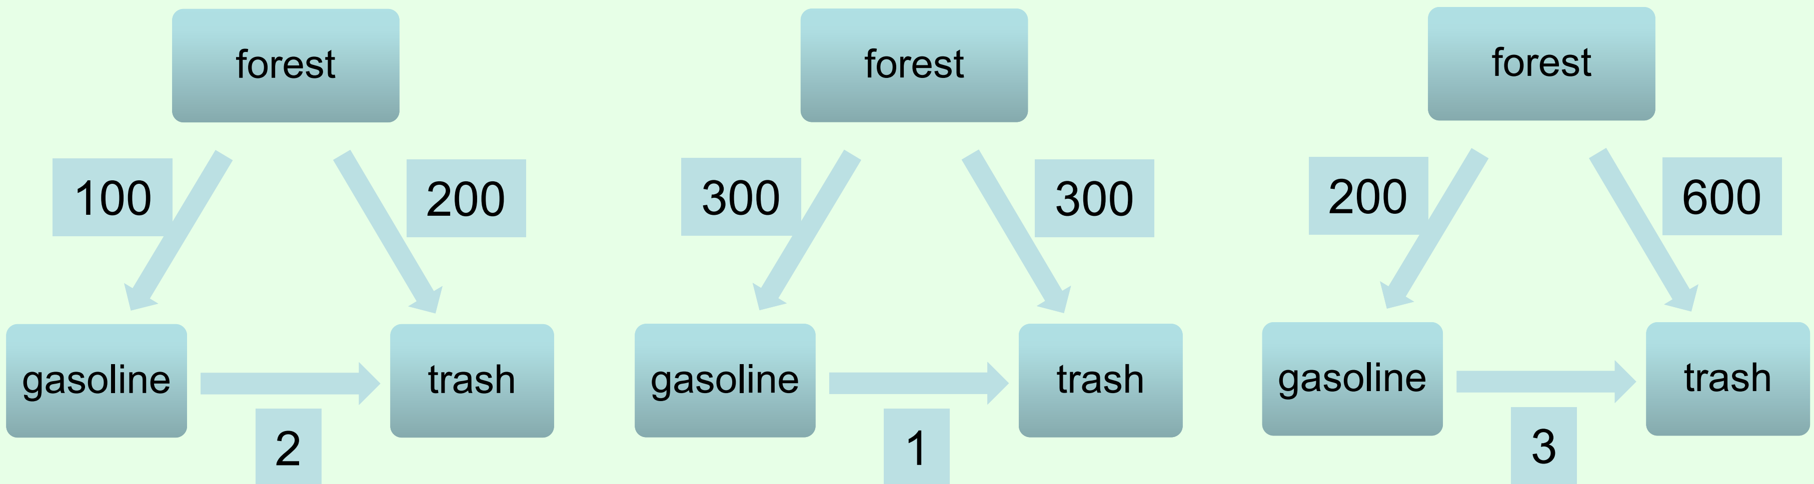

Just taking medians pairwise results in inconsistency

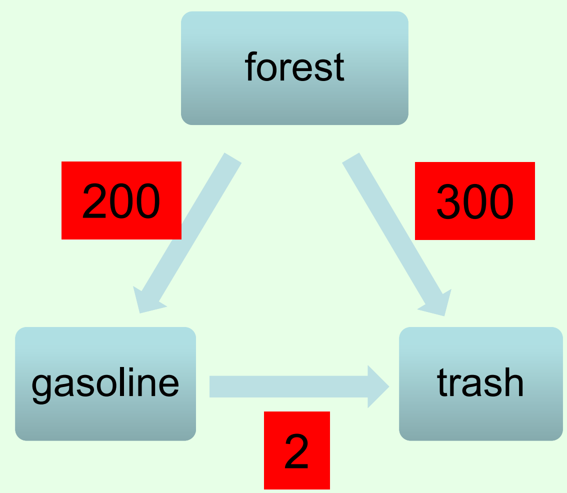

Conclusion

- Game-theoretic analysis of voting can appear hopeless
  - Impossibility results, multiplicity of equilibria, highly combinatorial domain
- Some variants still allow clean analysis
- Other variants provide a good challenge for computer scientists
  - Worst case analysis, algorithms, complexity, dynamics / learning, …

Thank you for your attention!