Some probabilistic questions in social choice

Mark C. Wilson
www.cs.auckland.ac.nz/~mcw/

Department of Computer Science
University of Auckland

Probability Seminar, UC Berkeley, 2013-04-03

Motivation

- I am not a probabilist, but am interested in probabilistic analysis of combinatorial situations.

Motivation

- I am not a probabilist, but am interested in probabilistic analysis of combinatorial situations.
- This year I am co-authoring a monograph with Robin Pemantle, *Analytic Combinatorics in Several Variables*. It contains several interesting probabilistic topics.

Motivation

- I am not a probabilist, but am interested in probabilistic analysis of combinatorial situations.
- This year I am co-authoring a monograph with Robin Pemantle, *Analytic Combinatorics in Several Variables*. It contains several interesting probabilistic topics.
- Today, will talk about something different. I am visiting UCB (Elchanan Mossel) and focusing on social choice, "the theory of collective decision-making without money".

Motivation

- I am not a probabilist, but am interested in probabilistic analysis of combinatorial situations.
- This year I am co-authoring a monograph with Robin Pemantle, *Analytic Combinatorics in Several Variables*. It contains several interesting probabilistic topics.
- Today, will talk about something different. I am visiting UCB (Elchanan Mossel) and focusing on social choice, "the theory of collective decision-making without money".
- Social choice theory gives rise to many interesting questions. There is still much scope for probabilists, in my opinion.

Motivation

- I am not a probabilist, but am interested in probabilistic analysis of combinatorial situations.
- This year I am co-authoring a monograph with Robin Pemantle, *Analytic Combinatorics in Several Variables*. It contains several interesting probabilistic topics.
- Today, will talk about something different. I am visiting UCB (Elchanan Mossel) and focusing on social choice, "the theory of collective decision-making without money".
- Social choice theory gives rise to many interesting questions. There is still much scope for probabilists, in my opinion.
- "Computational" social choice is very active, and is related to algorithmic mechanism design and multiagent systems. Complexity-theoretic results dominate the recent literature, and probability techniques used so far are fairly basic.

Basic social choice setup

- A finite set V of voters, $n := |V|$. Without loss of generality, $V = [n] := \{1, \ldots, n\}$.

Basic social choice setup

- A finite set V of voters, $n := |V|$. Without loss of generality, $V = [n] := \{1, \ldots, n\}$.
- A finite set A of alternatives, $m := |A|$. Wlog, $A = [m]$.

Basic social choice setup

- A finite set V of voters, $n := |V|$. Without loss of generality, $V = [n] := \{1, \ldots, n\}$.
- A finite set A of alternatives, $m := |A|$. Wlog, $A = [m]$.
- $L(A) :=$ set of linear orders on A, representing the possible strict preference orders. Identified with the set S_m of permutations on $[m]$.

Basic social choice setup

- A finite set V of voters, $n := |V|$. Without loss of generality, $V = [n] := \{1, \ldots, n\}$.
- A finite set A of alternatives, $m := |A|$. Wlog, $A = [m]$.
- $L(A) :=$ set of linear orders on A, representing the possible strict preference orders. Identified with the set S_m of permutations on $[m]$.
- Profile: map $V \to L(A)$ giving each voter's preference order. The set of profiles is identified with $(S_m)^n$.

Basic social choice setup

- A finite set V of voters, $n := |V|$. Without loss of generality, $V = [n] := \{1, \ldots, n\}$.
- A finite set A of alternatives, $m := |A|$. Wlog, $A = [m]$.
- $L(A) :=$ set of linear orders on A, representing the possible strict preference orders. Identified with the set S_m of permutations on $[m]$.
- Profile: map $V \to L(A)$ giving each voter's preference order. The set of profiles is identified with $(S_m)^n$.
- Voting situation: an orbit on profile space of the symmetric group on the voters. Can be represented as a composition (ordered partition) of n with $m!$ parts. We usually normalize by dividing by n, yielding the preference simplex.

Basic social choice setup

- A finite set V of voters, $n := |V|$. Without loss of generality, $V = [n] := \{1, \ldots, n\}$.
- A finite set A of alternatives, $m := |A|$. Wlog, $A = [m]$.
- $L(A) :=$ set of linear orders on A, representing the possible strict preference orders. Identified with the set S_m of permutations on $[m]$.
- Profile: map $V \to L(A)$ giving each voter's preference order. The set of profiles is identified with $(S_m)^n$.
- Voting situation: an orbit on profile space of the symmetric group on the voters. Can be represented as a composition (ordered partition) of n with $m!$ parts. We usually normalize by dividing by n, yielding the preference simplex.
- There are $m!^n$ profiles and $\binom{n+m!-1}{n}$ voting situations.

IC distribution

- Fix m and n. The IC distribution is the uniform distribution on S_m^n (generated by choosing each component independently and uniformly from S_m).

IC distribution

- Fix m and n. The IC distribution is the uniform distribution on S_m^n (generated by choosing each component independently and uniformly from S_m).
- This is a very special $(\alpha = 0)$ case of the Polyà-Eggenberger distribution.

IC distribution

- Fix m and n. The IC distribution is the uniform distribution on S_m^n (generated by choosing each component independently and uniformly from S_m).
- This is a very special ($\alpha = 0$) case of the Polyà-Eggenberger distribution.
- Not surprisingly, central limit theorems figure prominently in IC analyses. Other distributions in which each voter picks a preference order IID are similar in many ways.

IC distribution

- Fix m and n. The IC distribution is the uniform distribution on S_m^n (generated by choosing each component independently and uniformly from S_m).
- This is a very special ($\alpha = 0$) case of the Polyà-Eggenberger distribution.
- Not surprisingly, central limit theorems figure prominently in IC analyses. Other distributions in which each voter picks a preference order IID are similar in many ways.
- This is by far the most common distribution used in social choice. It has a relatively high probability of a very close election. Not usually a realistic model, but a useful extreme case, and fairly tractable.

Some other distributions

- The Polya-Eggenberger distribution with $\alpha = 1$ is the IAC distribution. This turns out to be the uniform distribution on voting situations: each configuration of anonymous voters is equally likely. Computations involve Ehrhart theory, volumes of polytopes, etc.

Some other distributions

- The Polya-Eggenberger distribution with $\alpha = 1$ is the IAC distribution. This turns out to be the uniform distribution on voting situations: each configuration of anonymous voters is equally likely. Computations involve Ehrhart theory, volumes of polytopes, etc.
- Mallows model: probability of a permutation π is proportional to $q^{i(\pi)}$ where $0 < q < 1$ and i is the number of inversions.

Some other distributions

- The Polya-Eggenberger distribution with $\alpha = 1$ is the IAC distribution. This turns out to be the uniform distribution on voting situations: each configuration of anonymous voters is equally likely. Computations involve Ehrhart theory, volumes of polytopes, etc.
- Mallows model: probability of a permutation π is proportional to $q^{i(\pi)}$ where $0 < q < 1$ and i is the number of inversions.
- Spatial model: each voter has an ideal point in some Euclidean space, and prefers candidates in inverse relation to their distance from this point. In dimension 1, this gives single-peaked preferences.

Some social choice rules

- The scoring rule defined by a fixed weight vector $(w_1, \ldots, w_m)$ gives w_i points to each candidate for each voter listing it in ith position. Highest total score wins. Special cases: plurality $(1, 0, 0, \ldots, 0)$; Borda $w_i = (m - i)/(m - 1)$, veto $(1, 1, \ldots, 1, 0)$.

Some social choice rules

- The scoring rule defined by a fixed weight vector $(w_1, \ldots, w_m)$ gives w_i points to each candidate for each voter listing it in ith position. Highest total score wins. Special cases: plurality $(1, 0, 0, \ldots, 0)$; Borda $w_i = (m - i)/(m - 1)$, veto $(1, 1, \ldots, 1, 0)$.
- Condorcet rules: if there is a Condorcet winner (preferred by some majority of voters to each other alternative), choose it. Otherwise choose something else. Most famous is Copeland rule (chess scoring).

Some social choice rules

- The scoring rule defined by a fixed weight vector $(w_1, \ldots, w_m)$ gives w_i points to each candidate for each voter listing it in ith position. Highest total score wins. Special cases: plurality $(1, 0, 0, \ldots, 0)$; Borda $w_i = (m - i)/(m - 1)$, veto $(1, 1, \ldots, 1, 0)$.
- Condorcet rules: if there is a Condorcet winner (preferred by some majority of voters to each other alternative), choose it. Otherwise choose something else. Most famous is Copeland rule (chess scoring).
- A huge generalization, including all rules used in practice, is hyperplane rules (generalized scoring rules). These are all anonymous, and the winner is constant on each chamber of the preference simplex, the chambers being defined by a finite set of hyperplanes.

Some key problems

- Fix a profile π, let $S \subseteq V$, let a be the winner under sincere voting and let $c \in A \setminus \{a\}$. Suppose that the members of S can change the winner to c by changing their votes.

Some key problems

- Fix a profile π, let $S \subseteq V$, let a be the winner under sincere voting and let $c \in A \setminus \{a\}$. Suppose that the members of S can change the winner to c by changing their votes.
- If all members of S prefer c to a, this is coalitional manipulation; otherwise it is bribery. When c is specified, the problem is constructive, otherwise destructive.

Some key problems

- Fix a profile π, let $S \subseteq V$, let a be the winner under sincere voting and let $c \in A \setminus \{a\}$. Suppose that the members of S can change the winner to c by changing their votes.
- If all members of S prefer c to a, this is coalitional manipulation; otherwise it is bribery. When c is specified, the problem is constructive, otherwise destructive.
- When S is specified, we have a decision problem: can they do it? When S is not specified, we have an optimization problem: minimize $|S|$ so that they can. We call S a winning coalition in general.

Some key problems

- Fix a profile π, let $S \subseteq V$, let a be the winner under sincere voting and let $c \in A \setminus \{a\}$. Suppose that the members of S can change the winner to c by changing their votes.
- If all members of S prefer c to a, this is coalitional manipulation; otherwise it is bribery. When c is specified, the problem is constructive, otherwise destructive.
- When S is specified, we have a decision problem: can they do it? When S is not specified, we have an optimization problem: minimize $|S|$ so that they can. We call S a winning coalition in general.
- Many old results compute the (asymptotic) probability that a winning coalition exists, under IC or IAC, at least for $m = 3$. Other distributions could be tried.

Unifying result for optimization problems

- Xia (2012) defines a class of optimization problems which includes several shown below.

Unifying result for optimization problems

- Xia (2012) defines a class of optimization problems which includes several shown below.
- He proves that in the IID case, for hyperplane rules, the optimal value is always one of: 0, ∞, $\Theta(\sqrt{n})$, $\Theta(n)$.

Unifying result for optimization problems

- Xia (2012) defines a class of optimization problems which includes several shown below.
- He proves that in the IID case, for hyperplane rules, the optimal value is always one of: 0, ∞, $\Theta(\sqrt{n})$, $\Theta(n)$.
- For smaller classes of rules and for IC, more specific results can be derived.

Margin of victory

- The Hamming metric d_H on S_m^n measures how many of the n components differ. Induced by the discrete metric on S_m.

Margin of victory

- The Hamming metric d_H on S_m^n measures how many of the n components differ. Induced by the discrete metric on S_m.
- For each profile π, the minimum Hamming distance to a profile with a different winner we denote by $\Delta_H(\pi)$. Xia (2012) calls this the margin of victory.

Margin of victory

- The Hamming metric d_H on S_m^n measures how many of the n components differ. Induced by the discrete metric on S_m.
- For each profile π, the minimum Hamming distance to a profile with a different winner we denote by $\Delta_H(\pi)$. Xia (2012) calls this the margin of victory.
- This is the optimization version of destructive bribery. Application: if $\Delta_H(\pi)$ is large, an election recount is likely not needed.

Margin of victory

- The Hamming metric d_H on S_m^n measures how many of the n components differ. Induced by the discrete metric on S_m.
- For each profile π, the minimum Hamming distance to a profile with a different winner we denote by $\Delta_H(\pi)$. Xia (2012) calls this the margin of victory.
- This is the optimization version of destructive bribery. Application: if $\Delta_H(\pi)$ is large, an election recount is likely not needed.
- Xia (2012): In the IID case, Δ_H has order $\sqrt{n}$ or n. Pritchard & Wilson (2009): For scoring rules under IC, it is $\sqrt{n}$ and very explicit formulae exist in terms of the weight vector. Mossel, Procaccia, & Racz (2013): for hyperplane rules under IC, the phase transition near $\sqrt{n}$ is smooth.

Swap robustness

- The Kemeny metric d_K on S_m^n measures the minimum number of single-voter adjacent transpositions needed to convert one permutation to the other. Induced by the bubblesort metric (Kendall's tau) on S_m.

Swap robustness

- The Kemeny metric d_K on S_m^n measures the minimum number of single-voter adjacent transpositions needed to convert one permutation to the other. Induced by the bubblesort metric (Kendall's tau) on S_m.
- For each profile π, define $\Delta_K(\pi)$ as before, with K replacing H. Shiryaev, Lu, & Elkind (2013) call this the robustness radius at π.

Swap robustness

- The Kemeny metric d_K on S_m^n measures the minimum number of single-voter adjacent transpositions needed to convert one permutation to the other. Induced by the bubblesort metric (Kendall's tau) on S_m.
- For each profile π, define $\Delta_K(\pi)$ as before, with K replacing H. Shiryaev, Lu, & Elkind (2013) call this the robustness radius at π.
- This is the optimization version of destructive swap bribery. Application: if this is large, then small errors by voters are unlikely to change the outcome.

Swap robustness

- The Kemeny metric d_K on S_m^n measures the minimum number of single-voter adjacent transpositions needed to convert one permutation to the other. Induced by the bubblesort metric (Kendall's tau) on S_m.
- For each profile π, define $\Delta_K(\pi)$ as before, with K replacing H. Shiryaev, Lu, & Elkind (2013) call this the robustness radius at π.
- This is the optimization version of destructive swap bribery. Application: if this is large, then small errors by voters are unlikely to change the outcome.
- Note that $\min_\pi \Delta_H(\pi) = 1 = \min_\pi \Delta_K(\pi)$ for most common rules. Shiryaev et al. compute $\max_\pi \Delta_K(\pi)$ for several common rules.

Manipulability

- So far we have looked at sensitivity of a voting rule to changes in the input (bribery model). We can consider the same results under the manipulation model.

Manipulability

- So far we have looked at sensitivity of a voting rule to changes in the input (bribery model). We can consider the same results under the manipulation model.
- Denote the analogue of Δ in this model by Δ^*.

Manipulability

- So far we have looked at sensitivity of a voting rule to changes in the input (bribery model). We can consider the same results under the manipulation model.
- Denote the analogue of Δ in this model by Δ^*.
- The Gibbard-Satterthwaite theorem says that unless the rule is degenerate, when $m \geq 3$ and $n \geq 2$ then $\min_\pi \Delta^*_H(\pi) = 1$. That is, a single voter can manipulate the outcome.

Minimum manipulating coalition size

- Another name for $\Delta^*_H(\pi)$ is the minimum manipulating coalition size.

Minimum manipulating coalition size

- Another name for $\Delta^*_H(\pi)$ is the minimum manipulating coalition size.
- Pritchard & Wilson (2009): for each scoring rule under IC, $\Delta^*_H/\sqrt{n}$ converges in distribution to an explicit distribution depending on the weight vector,

$$\Pr(n^{-1/2}\Delta^*_H \leq v) \to \Pr(V_w \leq w).$$

Minimum manipulating coalition size

- Another name for $\Delta_H^*(\pi)$ is the minimum manipulating coalition size.
- Pritchard & Wilson (2009): for each scoring rule under IC, $\Delta_H^*/\sqrt{n}$ converges in distribution to an explicit distribution depending on the weight vector,

$$\Pr(n^{-1/2}\Delta_H^* \leq v) \to \Pr(V_w \leq w).$$

- The limiting distribution can be explicitly computed as the minimum of a linear program in dimension 2 defined by m constraints, where the objective function involves Gaussian random variables.

Minimum manipulating coalition size

- Another name for $\Delta^*_H(\pi)$ is the minimum manipulating coalition size.
- Pritchard & Wilson (2009): for each scoring rule under IC, $\Delta^*_H/\sqrt{n}$ converges in distribution to an explicit distribution depending on the weight vector,

$$\Pr(n^{-1/2}\Delta^*_H \leq v) \to \Pr(V_w \leq w).$$

- The limiting distribution can be explicitly computed as the minimum of a linear program in dimension 2 defined by m constraints, where the objective function involves Gaussian random variables.
- For example, the asymptotically least manipulable scoring rule is "vote for your top half of the list" ($m/2$-approval), for $m \geq 5$.

Minimum manipulating coalition size

- Another name for $\Delta^*_H(\pi)$ is the minimum manipulating coalition size.
- Pritchard & Wilson (2009): for each scoring rule under IC, $\Delta^*_H/\sqrt{n}$ converges in distribution to an explicit distribution depending on the weight vector,

$$\Pr(n^{-1/2}\Delta^*_H \leq v) \to \Pr(V_w \leq w).$$

- The limiting distribution can be explicitly computed as the minimum of a linear program in dimension 2 defined by m constraints, where the objective function involves Gaussian random variables.
- For example, the asymptotically least manipulable scoring rule is "vote for your top half of the list" ($m/2$-approval), for $m \geq 5$.
- The phase transition near $\sqrt{n}$ is smooth.

A note on "coalitional manipulability"

- Computational social choice literature is dominated by results on the computational complexity of the "coalitional manipulation" decision problem.

A note on "coalitional manipulability"

- Computational social choice literature is dominated by results on the computational complexity of the "coalitional manipulation" decision problem.
- Strangely, the term "coalitional manipulation" is used for a problem where $S \not\subseteq V$ and members of S have no preferences. The correct term is "control by adding voters". Denote the optimal value by Γ.

A note on "coalitional manipulability"

- Computational social choice literature is dominated by results on the computational complexity of the "coalitional manipulation" decision problem.
- Strangely, the term "coalitional manipulation" is used for a problem where $S \not\subseteq V$ and members of S have no preferences. The correct term is "control by adding voters". Denote the optimal value by Γ.
- Xia (2012): For a large class of rules, Γ and Δ_H are of the same asymptotic order (any distribution, fixed m), so we will ignore Γ today. However, there may be some point in studying it probabilistically.

The probability of making a change

- For each metric (Hamming or Kemeny), each model (bribery or manipulation), and each $k \geq 1$, we can consider the probability P that a uniformly randomly chosen profile at distance k makes the appropriate change.

The probability of making a change

- For each metric (Hamming or Kemeny), each model (bribery or manipulation), and each $k \geq 1$, we can consider the probability P that a uniformly randomly chosen profile at distance k makes the appropriate change.
- Procaccia & Rosenschein (2007): compute $\max_\pi P_{K;1}(\pi)$ for several common rules. Application: if small, we can successfully aggregate preferences in noisy environments.

The probability of making a change

- For each metric (Hamming or Kemeny), each model (bribery or manipulation), and each $k \geq 1$, we can consider the probability P that a uniformly randomly chosen profile at distance k makes the appropriate change.
- Procaccia & Rosenschein (2007): compute $\max_\pi P_{K;1}(\pi)$ for several common rules. Application: if small, we can successfully aggregate preferences in noisy environments.
- The Gibbard-Satterthwaite theorem says that unless the rule is very degenerate, $\max_\pi P^*_{H;1}(\pi) > 0$ when $m \geq 3$ and $n \geq 2$.

The probability of making a change

- For each metric (Hamming or Kemeny), each model (bribery or manipulation), and each $k \geq 1$, we can consider the probability P that a uniformly randomly chosen profile at distance k makes the appropriate change.
- Procaccia & Rosenschein (2007): compute $\max_\pi P_{K;1}(\pi)$ for several common rules. Application: if small, we can successfully aggregate preferences in noisy environments.
- The Gibbard-Satterthwaite theorem says that unless the rule is very degenerate, $\max_\pi P^*_{H;1}(\pi) > 0$ when $m \geq 3$ and $n \geq 2$.
- Mossel & Racz (2012): a general lower bound for $\Pr_\pi P^*_{H;1}(\pi)$ as a polynomial in $n^{-1}, m^{-1}, \varepsilon$ where ε is the distance to the set of degenerate rules.

Table of results

	min	max	IC/IID	other dist'ns
Δ_H	easy		PW2009, MPR2013	
Δ_K	easy	SLE2013		
Δ_H^*	G1973,S1975		PW2009, X2012	
Δ_K^*	MR2012			
$P_{H;1}$				
$P_{K;1}$		PR2007		
$P_{H;1}^*$	G1973,S1975		MR2012	
$P_{K;1}^*$				

A query model for manipulation

- Query voters in turn by sampling uniformly without replacement.

A query model for manipulation

- Query voters in turn by sampling uniformly without replacement.
- Let Q be the number of queries required before we have found a manipulating coalition. For anonymous rules, this has the flavour of a coupon-collecting problem, where any of several minimal sub-multisets must be collected. Note that $Q \geq \Delta_H^*$ and Q is random even when π is given.

A query model for manipulation

- Query voters in turn by sampling uniformly without replacement.
- Let Q be the number of queries required before we have found a manipulating coalition. For anonymous rules, this has the flavour of a coupon-collecting problem, where any of several minimal sub-multisets must be collected. Note that $Q \geq \Delta_H^*$ and Q is random even when π is given.
- Can show that $\Pr(Q \leq k)$ is the probability of choosing a winning coalition if we first choose size uniformly, then a coalition uniformly.

A query model for manipulation

- Query voters in turn by sampling uniformly without replacement.
- Let Q be the number of queries required before we have found a manipulating coalition. For anonymous rules, this has the flavour of a coupon-collecting problem, where any of several minimal sub-multisets must be collected. Note that $Q \geq \Delta_H^*$ and Q is random even when π is given.
- Can show that $\Pr(Q \leq k)$ is the probability of choosing a winning coalition if we first choose size uniformly, then a coalition uniformly.
- Conjecture: for some class of rules including scoring rules, for some distributions including IC, there is a positive constant K such that $Q \leq K\Delta_H^*$ with high probability (as $n \to \infty$).

A query model for manipulation

- Query voters in turn by sampling uniformly without replacement.
- Let Q be the number of queries required before we have found a manipulating coalition. For anonymous rules, this has the flavour of a coupon-collecting problem, where any of several minimal sub-multisets must be collected. Note that $Q \geq \Delta_H^*$ and Q is random even when π is given.
- Can show that $\Pr(Q \leq k)$ is the probability of choosing a winning coalition if we first choose size uniformly, then a coalition uniformly.
- Conjecture: for some class of rules including scoring rules, for some distributions including IC, there is a positive constant K such that $Q \leq K\Delta_H^*$ with high probability (as $n \to \infty$).
- The analogous problem for other models also makes sense.

Aside: who cares about manipulation?

- Manipulability is often attacked on grounds of fairness, informational efficiency, etc.

Aside: who cares about manipulation?

- Manipulability is often attacked on grounds of fairness, informational efficiency, etc.
- From a consequentialist viewpoint, it is hard to understand why manipulation should be considered harmful. My opinion: the obsession with measuring (and minimizing) success of strategic voting has distracted from the main issue, welfare.

Aside: who cares about manipulation?

- Manipulability is often attacked on grounds of fairness, informational efficiency, etc.
- From a consequentialist viewpoint, it is hard to understand why manipulation should be considered harmful. My opinion: the obsession with measuring (and minimizing) success of strategic voting has distracted from the main issue, welfare.
- For reasonable rules, in order for a manipulation by few voters to succeed, c must be close to a in overall support. Thus changing the winner to c might not be particularly bad for overall social welfare. I don't know of any quantitative work on this.

Social welfare measures

- Social choice theory usually deals only with rankings, not cardinal utilities.

Social welfare measures

- Social choice theory usually deals only with rankings, not cardinal utilities.
- For scoring rules, can use the score of the winner as the measure of social welfare.

Social welfare measures

- Social choice theory usually deals only with rankings, not cardinal utilities.
- For scoring rules, can use the score of the winner as the measure of social welfare.
- Suppose that we *do* know the utility for each voter of each candidate being chosen by the rule, $u_v(c)$. We can then look at utilitarian: $U(c) := \sum_v u_v(c)$ or egalitarian: $E(c) := \min_v u_v(c)$ welfare measures.

Social welfare measures

- Social choice theory usually deals only with rankings, not cardinal utilities.
- For scoring rules, can use the score of the winner as the measure of social welfare.
- Suppose that we *do* know the utility for each voter of each candidate being chosen by the rule, $u_v(c)$. We can then look at utilitarian: $U(c) := \sum_v u_v(c)$ or egalitarian: $E(c) := \min_v u_v(c)$ welfare measures.
- Distortion (Price of anarchy): ratio of social welfare of optimal winner to social welfare of decentralized winner.

Socially optimal rules - sincere voting

- Consider randomized voting rules, where a profile is mapped to a probability distribution over A (a point in the m-simplex).

Socially optimal rules - sincere voting

- Consider randomized voting rules, where a profile is mapped to a probability distribution over A (a point in the m-simplex).
- A rule is average-case optimal if it maximizes expected social welfare on every profile.

Socially optimal rules - sincere voting

- Consider randomized voting rules, where a profile is mapped to a probability distribution over A (a point in the m-simplex).
- A rule is average-case optimal if it maximizes expected social welfare on every profile.
- Boutilier et al. (2012): if utilities are drawn IID, from a distribution symmetric on alternatives, and we use U, then the average-case optimal rule is a scoring rule. In particular, the Borda rule arises from the uniform distribution on $[0, 1]$.

Socially optimal rules - sincere voting

- Consider randomized voting rules, where a profile is mapped to a probability distribution over A (a point in the m-simplex).
- A rule is average-case optimal if it maximizes expected social welfare on every profile.
- Boutilier et al. (2012): if utilities are drawn IID, from a distribution symmetric on alternatives, and we use U, then the average-case optimal rule is a scoring rule. In particular, the Borda rule arises from the uniform distribution on $[0, 1]$.
- They also study worst-case optimality and show that randomized rules are qualitatively better than deterministic ones.

Socially optimal rules - sincere voting

- Consider randomized voting rules, where a profile is mapped to a probability distribution over A (a point in the m-simplex).
- A rule is average-case optimal if it maximizes expected social welfare on every profile.
- Boutilier et al. (2012): if utilities are drawn IID, from a distribution symmetric on alternatives, and we use U, then the average-case optimal rule is a scoring rule. In particular, the Borda rule arises from the uniform distribution on $[0, 1]$.
- They also study worst-case optimality and show that randomized rules are qualitatively better than deterministic ones.
- Problem: what can be said about other utility distributions? other welfare measures? commonly used rules?

Choosing a parliament

- We want to allocate seats to parties in a "proportional representation" system. Ignoring roundoff, this can be thought of as a randomized voting rule, where probability corresponds to fraction of seats. This gives more justification for randomized voting rules.

Choosing a parliament

- We want to allocate seats to parties in a "proportional representation" system. Ignoring roundoff, this can be thought of as a randomized voting rule, where probability corresponds to fraction of seats. This gives more justification for randomized voting rules.
- Rules used in practice only look at first preferences (plurality). However, many more general methods are possible, e.g. optimal voting rules as in Boutilier et al.

Choosing a parliament

- We want to allocate seats to parties in a "proportional representation" system. Ignoring roundoff, this can be thought of as a randomized voting rule, where probability corresponds to fraction of seats. This gives more justification for randomized voting rules.
- Rules used in practice only look at first preferences (plurality). However, many more general methods are possible, e.g. optimal voting rules as in Boutilier et al.
- Voters likely have preferences over parliaments, not just parties. Expected utility doesn't seem right in this case.

Choosing a parliament

- We want to allocate seats to parties in a "proportional representation" system. Ignoring roundoff, this can be thought of as a randomized voting rule, where probability corresponds to fraction of seats. This gives more justification for randomized voting rules.
- Rules used in practice only look at first preferences (plurality). However, many more general methods are possible, e.g. optimal voting rules as in Boutilier et al.
- Voters likely have preferences over parliaments, not just parties. Expected utility doesn't seem right in this case.
- There are general methods for "optimal proportional representation" by Monroe (1995) and Chamberlin & Courant (1983). These use Borda score as a welfare measure and are a slightly different model. I know of no analytic results about their average-case behaviour, nor other welfare measures.

Socially optimal rules - strategic voting

- Once we consider manipulation in this context, we enter the realm of game theory. In full generality this is very hard, because there are (too) many equilibria under strategic voting, some of which have very bad social welfare.

Socially optimal rules - strategic voting

- Once we consider manipulation in this context, we enter the realm of game theory. In full generality this is very hard, because there are (too) many equilibria under strategic voting, some of which have very bad social welfare.
- Restrictions on allowed strategic behaviour are commonly made (sequential voting, asynchronous best-reply voting, . . .). Thompson, Lev, Leyton-Brown & Rosenschein (2012) have a nice equilibrium refinement.

Socially optimal rules - strategic voting

- Once we consider manipulation in this context, we enter the realm of game theory. In full generality this is very hard, because there are (too) many equilibria under strategic voting, some of which have very bad social welfare.
- Restrictions on allowed strategic behaviour are commonly made (sequential voting, asynchronous best-reply voting, . . .). Thompson, Lev, Leyton-Brown & Rosenschein (2012) have a nice equilibrium refinement.
- Simulation results (Lehtinen 2006, Xia & Conitzer 2010, some of my students, Thompson et al.) so far show that widespread strategic voting generally increases overall social welfare.

Socially optimal rules - strategic voting

- Once we consider manipulation in this context, we enter the realm of game theory. In full generality this is very hard, because there are (too) many equilibria under strategic voting, some of which have very bad social welfare.
- Restrictions on allowed strategic behaviour are commonly made (sequential voting, asynchronous best-reply voting, . . .). Thompson, Lev, Leyton-Brown & Rosenschein (2012) have a nice equilibrium refinement.
- Simulation results (Lehtinen 2006, Xia & Conitzer 2010, some of my students, Thompson et al.) so far show that widespread strategic voting generally increases overall social welfare.
- It would be very desirable to have some analytic results even for IC. I don't see to make progress. Any ideas?

Other topics of interest (to me)

- Efficient aggregation of information: Condorcet Jury Theorem, wisdom of crowds, voting rules as maximum likelihood estimators.

Other topics of interest (to me)

- Efficient aggregation of information: Condorcet Jury Theorem, wisdom of crowds, voting rules as maximum likelihood estimators.
- Extension of the above to social networks (e.g. Mossel, Tamuz & Neeman 2012).

Other topics of interest (to me)

- Efficient aggregation of information: Condorcet Jury Theorem, wisdom of crowds, voting rules as maximum likelihood estimators.
- Extension of the above to social networks (e.g. Mossel, Tamuz & Neeman 2012).
- Asymptotic probabilities of interesting events: Condorcet efficiency, no-show paradox, etc.

Other topics of interest (to me)

- Efficient aggregation of information: Condorcet Jury Theorem, wisdom of crowds, voting rules as maximum likelihood estimators.
- Extension of the above to social networks (e.g. Mossel, Tamuz & Neeman 2012).
- Asymptotic probabilities of interesting events: Condorcet efficiency, no-show paradox, etc.
- Fair division algorithms, divisible and indivisible case. Average-case analysis and randomized algorithms.

Other topics of interest (to me)

- Efficient aggregation of information: Condorcet Jury Theorem, wisdom of crowds, voting rules as maximum likelihood estimators.
- Extension of the above to social networks (e.g. Mossel, Tamuz & Neeman 2012).
- Asymptotic probabilities of interesting events: Condorcet efficiency, no-show paradox, etc.
- Fair division algorithms, divisible and indivisible case. Average-case analysis and randomized algorithms.
- Disaggregation of a global probability distribution into local distributions - models of changes in party support.

Other topics of interest (to me)

- Efficient aggregation of information: Condorcet Jury Theorem, wisdom of crowds, voting rules as maximum likelihood estimators.
- Extension of the above to social networks (e.g. Mossel, Tamuz & Neeman 2012).
- Asymptotic probabilities of interesting events: Condorcet efficiency, no-show paradox, etc.
- Fair division algorithms, divisible and indivisible case. Average-case analysis and randomized algorithms.
- Disaggregation of a global probability distribution into local distributions - models of changes in party support.
- Relation of Q to semivalues, probabilistic models of coalition formation, power indices.