Measuring the manipulability of voting rules

Mark C. Wilson
www.cs.auckland.ac.nz/~mcw/blog/
(joint with Geoff Pritchard and Reyhaneh Reyhani)

Department of Computer Science
University of Auckland

Rochester, 2009-10-16

Background on speaker

Recent results in standard framework

Discussion of assumptions

Mechanism design

Aims of this talk

- Publicize my existence, and look for joint work opportunities.

Aims of this talk

- Publicize my existence, and look for joint work opportunities.
- Publicize CMSS at University of Auckland.

Aims of this talk

- Publicize my existence, and look for joint work opportunities.
- Publicize CMSS at University of Auckland.
- Publicize recent results with which I have been involved.

Aims of this talk

- Publicize my existence, and look for joint work opportunities.
- Publicize CMSS at University of Auckland.
- Publicize recent results with which I have been involved.
- Argue that the direction of some recent research in the COMSOC community is misguided.

Centre for Mathematical Social Sciences

- At University of Auckland, with researchers from Mathematics, Computer Science, Statistics, Economics departments.

Centre for Mathematical Social Sciences

- At University of Auckland, with researchers from Mathematics, Computer Science, Statistics, Economics departments.
- www.cs.auckland.ac.nz/~mcw/dmss/

Centre for Mathematical Social Sciences

- At University of Auckland, with researchers from Mathematics, Computer Science, Statistics, Economics departments.
- www.cs.auckland.ac.nz/~mcw/dmss/
- Main interests are in computational social choice, computational game theory, utility theory and decision theory.

Centre for Mathematical Social Sciences

- At University of Auckland, with researchers from Mathematics, Computer Science, Statistics, Economics departments.
- `www.cs.auckland.ac.nz/~mcw/dmss/`
- Main interests are in computational social choice, computational game theory, utility theory and decision theory.
- Running a workshop 18–20 February 2010.

References

PrWi2007 G. Pritchard and M. C. Wilson, Exact results on manipulability of positional voting rules, Social Choice and Welfare 29 (2007), 487–513.

References

PrWi2007 G. Pritchard and M. C. Wilson, Exact results on manipulability of positional voting rules, Social Choice and Welfare 29 (2007), 487–513.

WiPr2007 M. C. Wilson and G. Pritchard, Probability calculations under the IAC hypothesis, Mathematical Social Sciences 54 (2007), 244–256.

References

PrWi2007 G. Pritchard and M. C. Wilson, Exact results on manipulability of positional voting rules, Social Choice and Welfare 29 (2007), 487–513.

WiPr2007 M. C. Wilson and G. Pritchard, Probability calculations under the IAC hypothesis, Mathematical Social Sciences 54 (2007), 244–256.

PrWi2009 G. Pritchard and M. C. Wilson, Asymptotics of the minimum manipulating coalition size for positional rules under IC behaviour, Mathematical Social Sciences 58 (2009), 35–56.

References

PrWi2007 G. Pritchard and M. C. Wilson, Exact results on manipulability of positional voting rules, Social Choice and Welfare 29 (2007), 487–513.

WiPr2007 M. C. Wilson and G. Pritchard, Probability calculations under the IAC hypothesis, Mathematical Social Sciences 54 (2007), 244–256.

PrWi2009 G. Pritchard and M. C. Wilson, Asymptotics of the minimum manipulating coalition size for positional rules under IC behaviour, Mathematical Social Sciences 58 (2009), 35–56.

RPW201x R. Reyhani, G. Pritchard and M. C. Wilson, A new measure of the difficulty of manipulation of voting rules, preprint 2009.

Basic setup

- We consider a voting rule (social choice function) that aggregates preferences of a set V of n voters.

Basic setup

- We consider a voting rule (social choice function) that aggregates preferences of a set V of n voters.
- Each voter submits a total order of the m candidates.

Basic setup

- We consider a voting rule (social choice function) that aggregates preferences of a set V of n voters.
- Each voter submits a total order of the m candidates.
- We break ties symmetrically: choose a tied winner uniformly at random. This is not an essential assumption but makes computation somewhat easier.

Basic setup

- We consider a voting rule (social choice function) that aggregates preferences of a set V of n voters.
- Each voter submits a total order of the m candidates.
- We break ties symmetrically: choose a tied winner uniformly at random. This is not an essential assumption but makes computation somewhat easier.
- We can describe the individual votes by a profile, an ordered list of the individual votes. There are $(m!)^n$ of these.

Basic setup

- We consider a voting rule (social choice function) that aggregates preferences of a set V of n voters.
- Each voter submits a total order of the m candidates.
- We break ties symmetrically: choose a tied winner uniformly at random. This is not an essential assumption but makes computation somewhat easier.
- We can describe the individual votes by a profile, an ordered list of the individual votes. There are $(m!)^n$ of these.
- For anonymous voting rules there is symmetry between candidates, so we need only the succinct input (voting situation). List the $m!$ possible votes in some way, and then list the number n_i of voters with ordering i. There are $\binom{n+m!-1}{n}$ of these.

Manipulation

- Suppose that some subset S of voters each submit a vote that differs from their sincere preference, and thereby achieve a better result than if they had voted sincerely (assuming all other voters vote sincerely each time). We say that the profile is manipulable by S.

Manipulation

- Suppose that some subset S of voters each submit a vote that differs from their sincere preference, and thereby achieve a better result than if they had voted sincerely (assuming all other voters vote sincerely each time). We say that the profile is manipulable by S.
- A profile is manipulable by some S if and only if the sincere strategy does not give a strong Nash equilibrium of the associated game.

Manipulation

- Suppose that some subset S of voters each submit a vote that differs from their sincere preference, and thereby achieve a better result than if they had voted sincerely (assuming all other voters vote sincerely each time). We say that the profile is manipulable by S.
- A profile is manipulable by some S if and only if the sincere strategy does not give a strong Nash equilibrium of the associated game.
- Gibbard-Satterthwaite theorem says that if $m \geq 3$ and the rule is fair to voters and candidates, then it is manipulable in some situation.

Manipulation

- Suppose that some subset S of voters each submit a vote that differs from their sincere preference, and thereby achieve a better result than if they had voted sincerely (assuming all other voters vote sincerely each time). We say that the profile is manipulable by S.
- A profile is manipulable by some S if and only if the sincere strategy does not give a strong Nash equilibrium of the associated game.
- Gibbard-Satterthwaite theorem says that if $m \geq 3$ and the rule is fair to voters and candidates, then it is manipulable in some situation.
- Since manipulation is essentially unavoidable, how can we minimize its impact?

Manipulation example

- Candidates a, b, c and Borda rule (each voter gives his 1st choice 2 points, 2nd gets 1 point, 3rd gets 0).

Manipulation example

- Candidates a, b, c and Borda rule (each voter gives his 1st choice 2 points, 2nd gets 1 point, 3rd gets 0).
- Voting situation: 2 bac, 2 abc, 2 acb, 3 cba. Sincere scores are $10, 9, 8$ for a, b, c.

Manipulation example

- Candidates a, b, c and Borda rule (each voter gives his 1st choice 2 points, 2nd gets 1 point, 3rd gets 0).
- Voting situation: 2 bac, 2 abc, 2 acb, 3 cba. Sincere scores are $10, 9, 8$ for a, b, c.
- If a cba changes vote to bca, then a and b tie, while if 2 cba change in that way, b wins clearly.

Manipulation example

- Candidates a, b, c and Borda rule (each voter gives his 1st choice 2 points, 2nd gets 1 point, 3rd gets 0).
- Voting situation: 2 bac, 2 abc, 2 acb, 3 cba. Sincere scores are $10, 9, 8$ for a, b, c.
- If a cba changes vote to bca, then a and b tie, while if 2 cba change in that way, b wins clearly.
- If both bac changed to bca, then c would win, and this goes against the preferences of such voters, so they can't manipulate by so voting (nor in any other way).

Manipulation example

- Candidates a, b, c and Borda rule (each voter gives his 1st choice 2 points, 2nd gets 1 point, 3rd gets 0).
- Voting situation: 2 bac, 2 abc, 2 acb, 3 cba. Sincere scores are $10, 9, 8$ for a, b, c.
- If a cba changes vote to bca, then a and b tie, while if 2 cba change in that way, b wins clearly.
- If both bac changed to bca, then c would win, and this goes against the preferences of such voters, so they can't manipulate by so voting (nor in any other way).
- Manipulability can be described by systems of integer linear (in)equalities for most commonly used rules, including all scoring rules, Copeland's rule, etc.

Linear system example: Borda, $m = 3$

- Suppose that the sincere election result is $|a| > |b| \geq |c|$, and we want to manipulate in favour of b.

Linear system example: Borda, $m = 3$

- Suppose that the sincere election result is $|a| > |b| \geq |c|$, and we want to manipulate in favour of b.
- Assuming other voters vote naively, an optimal strategy is for some bac voters (say y) and all cba to vote bca.

Linear system example: Borda, $m = 3$

- Suppose that the sincere election result is $|a| > |b| \geq |c|$, and we want to manipulate in favour of b.
- Assuming other voters vote naively, an optimal strategy is for some bac voters (say y) and all cba to vote bca.
- Let $|a|'$ denote a's score after a strategic attempt as above. Then the attempt is successful if and only if $|b|' > |a|', |c|'$. We can express $|a|'$ as a linear combination of the n_i and y, and also eliminate y. This yields $n_i \geq 0$, $\sum_i n_i = n$, and

$$
\begin{aligned}
0 &\leq n_1 + n_2 - n_3 - n_4 \\
0 &\leq n_3 + n_4 - n_5 - n_6 \\
0 &\leq -n_1 - n_2 + n_3 + n_4 + n_6 \\
0 &\leq -n_1 - n_2 + 2n_3 + 2n_4 - n_5 + 2n_2.
\end{aligned}
$$

Linear system example: Borda, $m = 3$

- Suppose that the sincere election result is $|a| > |b| \geq |c|$, and we want to manipulate in favour of b.
- Assuming other voters vote naively, an optimal strategy is for some bac voters (say y) and all cba to vote bca.
- Let $|a|'$ denote a's score after a strategic attempt as above. Then the attempt is successful if and only if $|b|' > |a|', |c|'$. We can express $|a|'$ as a linear combination of the n_i and y, and also eliminate y. This yields $n_i \geq 0$, $\sum_i n_i = n$, and

$$
\begin{aligned}
0 &\leq n_1 + n_2 - n_3 - n_4 \\
0 &\leq n_3 + n_4 - n_5 - n_6 \\
0 &\leq -n_1 - n_2 + n_3 + n_4 + n_6 \\
0 &\leq -n_1 - n_2 + 2n_3 + 2n_4 - n_5 + 2n_2.
\end{aligned}
$$

- Can do this for any coalition X (above is $X = V$).

IAC culture computations

- Each voting situation is equally likely. A special case of a Polya-Eggenberger distribution, used to model contagion.

IAC culture computations

- Each voting situation is equally likely. A special case of a Polya-Eggenberger distribution, used to model contagion.
- Computations reduce to counting lattice points in polytopes, since each point has the same probability.

IAC culture computations

- Each voting situation is equally likely. A special case of a Polya-Eggenberger distribution, used to model contagion.
- Computations reduce to counting lattice points in polytopes, since each point has the same probability.
- There is a large theory of such counting (e.g. Ehrhart polynomials), with nice algorithms e.g. (Barvinok), and software available (e.g. LattE).

IAC culture computations

- Each voting situation is equally likely. A special case of a Polya-Eggenberger distribution, used to model contagion.
- Computations reduce to counting lattice points in polytopes, since each point has the same probability.
- There is a large theory of such counting (e.g. Ehrhart polynomials), with nice algorithms e.g. (Barvinok), and software available (e.g. LattE).
- Sample result: let $P(n)$ be the probability that manipulation by some coalition is possible, for Borda with $m = 3$ under IAC. Then $P(n)$ is the ration of quasipolynomials in n with leading coefficient $132953/264600 \approx 0.5024678760$.

IAC culture computations

- Each voting situation is equally likely. A special case of a Polya-Eggenberger distribution, used to model contagion.
- Computations reduce to counting lattice points in polytopes, since each point has the same probability.
- There is a large theory of such counting (e.g. Ehrhart polynomials), with nice algorithms e.g. (Barvinok), and software available (e.g. LattE).
- Sample result: let $P(n)$ be the probability that manipulation by some coalition is possible, for Borda with $m = 3$ under IAC. Then $P(n)$ is the ration of quasipolynomials in n with leading coefficient $132953/264600 \approx 0.5024678760$.
- PrWi2007 explains this kind of result in detail.

IC culture computations

- Each voter chooses a profile uniformly at random.

IC culture computations

- Each voter chooses a profile uniformly at random.
- Computations generally reduce to central limit approximations.

IC culture computations

- Each voter chooses a profile uniformly at random.
- Computations generally reduce to central limit approximations.
- Sample results: for fixed m and every positional rule except antiplurality, a random profile is manipulable with probability that approaches 1 exponentially fast as $n \to \infty$.

IC culture computations

- Each voter chooses a profile uniformly at random.
- Computations generally reduce to central limit approximations.
- Sample results: for fixed m and every positional rule except antiplurality, a random profile is manipulable with probability that approaches 1 exponentially fast as $n \to \infty$.
- Such results are "classical".

Some measures of manipulability

- Indicator: is the situation manipulable by some coalition, of unspecified size? Used in the vast majority of papers until the last few years. Does not discriminate between large and small coalitions.

Some measures of manipulability

- Indicator: is the situation manipulable by some coalition, of unspecified size? Used in the vast majority of papers until the last few years. Does not discriminate between large and small coalitions.
- Size of smallest manipulating coalition for the given situation. Used in our papers and introduced by Chamberlin (1985). Better than previous one, but doesn't take into account the prevalence of coalitions of small size.

Some measures of manipulability

- Indicator: is the situation manipulable by some coalition, of unspecified size? Used in the vast majority of papers until the last few years. Does not discriminate between large and small coalitions.
- Size of smallest manipulating coalition for the given situation. Used in our papers and introduced by Chamberlin (1985). Better than previous one, but doesn't take into account the prevalence of coalitions of small size.
- Is the situation manipulable by k or fewer voters?

Some measures of manipulability

- Indicator: is the situation manipulable by some coalition, of unspecified size? Used in the vast majority of papers until the last few years. Does not discriminate between large and small coalitions.
- Size of smallest manipulating coalition for the given situation. Used in our papers and introduced by Chamberlin (1985). Better than previous one, but doesn't take into account the prevalence of coalitions of small size.
- Is the situation manipulable by k or fewer voters?
- Averages of these over all situations according to the culture.

A new measure of manipulability (RPW201x)

- For a fixed situation, choose voters randomly without replacement, until our set first contains a manipulating coalition.

A new measure of manipulability (RPW201x)

- For a fixed situation, choose voters randomly without replacement, until our set first contains a manipulating coalition.
- We can think of an instigator of manipulation, who does not know the preferences of the voters, but wants to disrupt the election by manipulating. This agent interviews voters one by one until a coalition is found.

A new measure of manipulability (RPW201x)

- For a fixed situation, choose voters randomly without replacement, until our set first contains a manipulating coalition.
- We can think of an instigator of manipulation, who does not know the preferences of the voters, but wants to disrupt the election by manipulating. This agent interviews voters one by one until a coalition is found.
- Let Q be the (defective) random variable thus obtained.

A new measure of manipulability (RPW201x)

- For a fixed situation, choose voters randomly without replacement, until our set first contains a manipulating coalition.
- We can think of an instigator of manipulation, who does not know the preferences of the voters, but wants to disrupt the election by manipulating. This agent interviews voters one by one until a coalition is found.
- Let Q be the (defective) random variable thus obtained.
- It is easy to show that $\Pr(Q \leq k)$ equals the probability that a randomly chosen k-subset of V contains a manipulating coalition.

A new measure of manipulability (RPW201x)

- For a fixed situation, choose voters randomly without replacement, until our set first contains a manipulating coalition.
- We can think of an instigator of manipulation, who does not know the preferences of the voters, but wants to disrupt the election by manipulating. This agent interviews voters one by one until a coalition is found.
- Let Q be the (defective) random variable thus obtained.
- It is easy to show that $\Pr(Q \leq k)$ equals the probability that a randomly chosen k-subset of V contains a manipulating coalition.
- Thus Q measures both the size and prevalence of manipulating coalitions. It contains more information than many other measures.

Results for scoring rules

- Exact computation of distribution function of M (resp. Q) for $m = 3$ up to $n = 150$ (resp. $n = 25$) for 6 rules, under 2 probability distributions. (PrWi2007, RPW201x)

Results for scoring rules

- Exact computation of distribution function of M (resp. Q) for $m = 3$ up to $n = 150$ (resp. $n = 25$) for 6 rules, under 2 probability distributions. (PrWi2007, RPW201x)
- Under uniform distribution (IC culture), an analytic description of asymptotic (in n) size of M for any fixed m. (PrWi2009)

Results for scoring rules

- Exact computation of distribution function of M (resp. Q) for $m = 3$ up to $n = 150$ (resp. $n = 25$) for 6 rules, under 2 probability distributions. (PrWi2007, RPW201x)
- Under uniform distribution (IC culture), an analytic description of asymptotic (in n) size of M for any fixed m. (PrWi2009)
- A heuristic argument as to why we should have $Q \leq CM$ with high probability (under IC), where C depends on the rule and m. (Future work by PhD student)

$\Pr(Q \leq k)$ for $n = 25$ under IAC, 3 rules

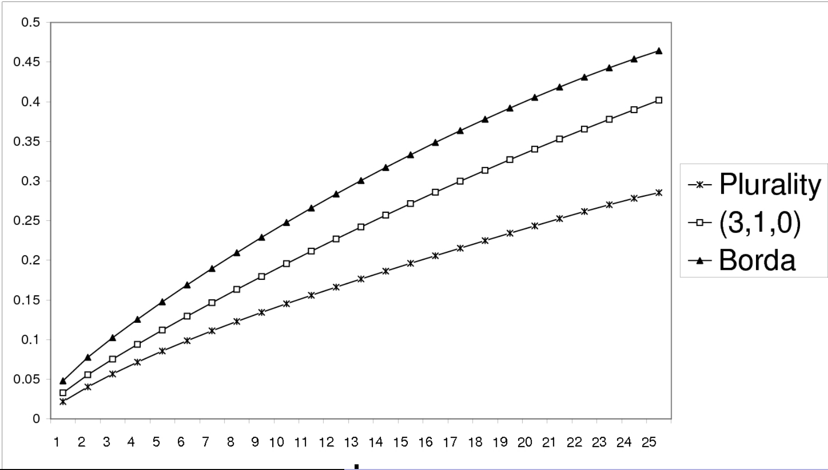

$\Pr(Q \leq k)$ for $n = 25$ under IC, 3 rules

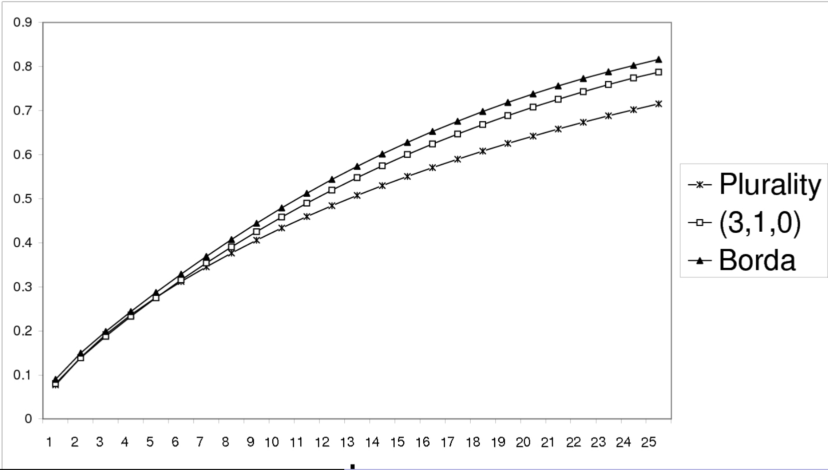

Asymptotics of M under IC for positional rules: results

- For fixed m, and fixed weight vector w, $M/\sqrt{n}$ converges in distribution to a random variable A_w.

Asymptotics of M under IC for positional rules: results

- For fixed m, and fixed weight vector w, $M/\sqrt{n}$ converges in distribution to a random variable A_w.
- The form of A_w is $C(m, w)B$ for some absolute random variable B (except for antiplurality).

Asymptotics of M under IC for positional rules: results

- For fixed m, and fixed weight vector w, $M/\sqrt{n}$ converges in distribution to a random variable A_w.
- The form of A_w is $C(m, w)B$ for some absolute random variable B (except for antiplurality).
- Hence we can compare rules according to the constants C.

Asymptotics of M under IC for positional rules: results

- For fixed m, and fixed weight vector w, $M/\sqrt{n}$ converges in distribution to a random variable A_w.
- The form of A_w is $C(m, w)B$ for some absolute random variable B (except for antiplurality).
- Hence we can compare rules according to the constants C.
- The plurality rule is always dominated.

Asymptotics of M under IC for positional rules: results

- For fixed m, and fixed weight vector w, $M/\sqrt{n}$ converges in distribution to a random variable A_w.
- The form of A_w is $C(m,w)B$ for some absolute random variable B (except for antiplurality).
- Hence we can compare rules according to the constants C.
- The plurality rule is always dominated.
- The Borda rule is the most resistant to very small coalitions.

Asymptotics of M under IC for positional rules: results

- For fixed m, and fixed weight vector w, $M/\sqrt{n}$ converges in distribution to a random variable A_w.
- The form of A_w is $C(m,w)B$ for some absolute random variable B (except for antiplurality).
- Hence we can compare rules according to the constants C.
- The plurality rule is always dominated.
- The Borda rule is the most resistant to very small coalitions.
- For most m, the $m/2$-approval rule dominates the others.

Asymptotics of M under IC for positional rules: results

- For fixed m, and fixed weight vector w, $M/\sqrt{n}$ converges in distribution to a random variable A_w.
- The form of A_w is $C(m,w)B$ for some absolute random variable B (except for antiplurality).
- Hence we can compare rules according to the constants C.
- The plurality rule is always dominated.
- The Borda rule is the most resistant to very small coalitions.
- For most m, the $m/2$-approval rule dominates the others.
- There is approximate symmetry between plurality and antiplurality for $m \geq 6$, contrary to the situation for $m = 3$.

Asymptotics of M under IC for positional rules: proof sketch

- The scores obey a multivariate central limit theorem, with mean and covariances of order n.

Asymptotics of M under IC for positional rules: proof sketch

- The scores obey a multivariate central limit theorem, with mean and covariances of order n.
- Manipulation is relatively easy because much weight is placed on situations where scores are nearly tied.

Asymptotics of M under IC for positional rules: proof sketch

- The scores obey a multivariate central limit theorem, with mean and covariances of order n.
- Manipulation is relatively easy because much weight is placed on situations where scores are nearly tied.
- The integer linear program describing manipulability can be relaxed by removing integrality.

Asymptotics of M under IC for positional rules: proof sketch

- The scores obey a multivariate central limit theorem, with mean and covariances of order n.
- Manipulation is relatively easy because much weight is placed on situations where scores are nearly tied.
- The integer linear program describing manipulability can be relaxed by removing integrality.
- We may simplify much more by restricting to manipulations where b overtakes a, and ignoring other candidates.

Asymptotics of M under IC for positional rules: proof sketch

- The scores obey a multivariate central limit theorem, with mean and covariances of order n.
- Manipulation is relatively easy because much weight is placed on situations where scores are nearly tied.
- The integer linear program describing manipulability can be relaxed by removing integrality.
- We may simplify much more by restricting to manipulations where b overtakes a, and ignoring other candidates.
- Linear programming duality leads to a problem we can solve explicitly.

Extensions

- Numerical results indicate that Copeland's rule M and Q are (much?) larger than for positional rules. What are the asymptotics, for example under IC?

Extensions

- Numerical results indicate that Copeland's rule M and Q are (much?) larger than for positional rules. What are the asymptotics, for example under IC?
- Prove or disprove conjecture on $Q \leq CM$.

Extensions

- Numerical results indicate that Copeland's rule M and Q are (much?) larger than for positional rules. What are the asymptotics, for example under IC?
- Prove or disprove conjecture on $Q \leq CM$.
- Extend results on positional rules to the class of Conitzer and Xia.

Standard assumptions in the COMSOC literature

- Computational complexity theory provides a good measure of manipulability.

Standard assumptions in the COMSOC literature

- Computational complexity theory provides a good measure of manipulability.
- Manipulation is something to be minimized.

Standard assumptions in the COMSOC literature

- Computational complexity theory provides a good measure of manipulability.
- Manipulation is something to be minimized.

Standard assumptions in the COMSOC literature

- Computational complexity theory provides a good measure of manipulability.
- Manipulation is something to be minimized.

In my opinion, both of these are wrong, especially the second.

Computational complexity: pro and con

- Under weak assumptions, (unweighted) voting rules seem to be easy to manipulate on average (Conitzer-Sandholm, Friedgut-Kalai-Nisan, Dobzinski-Procaccia, Conitzer-Xia).

Computational complexity: pro and con

- Under weak assumptions, (unweighted) voting rules seem to be easy to manipulate on average (Conitzer-Sandholm, Friedgut-Kalai-Nisan, Dobzinski-Procaccia, Conitzer-Xia).
- When m is small, as in many human applications, then almost all rules are manipulable in polynomial time.

Computational complexity: pro and con

- Under weak assumptions, (unweighted) voting rules seem to be easy to manipulate on average (Conitzer-Sandholm, Friedgut-Kalai-Nisan, Dobzinski-Procaccia, Conitzer-Xia).
- When m is small, as in many human applications, then almost all rules are manipulable in polynomial time.
- When m is large, as for search engines, then some rules are NP-hard to manipulate even for small n.

Computational complexity: pro and con

- Under weak assumptions, (unweighted) voting rules seem to be easy to manipulate on average (Conitzer-Sandholm, Friedgut-Kalai-Nisan, Dobzinski-Procaccia, Conitzer-Xia).
- When m is small, as in many human applications, then almost all rules are manipulable in polynomial time.
- When m is large, as for search engines, then some rules are NP-hard to manipulate even for small n.
- Results can be quite crude. For example, for fixed m under IC, there is a threshold around $k = \sqrt{n}$ where manipulability switches from almost impossible to almost inevitable, but complexity results say nothing about this.

Implementing voting rules

- Assume that all voters have common knowledge and are rational with no limit on computational power. Try to design a voting rule such that everyone has incentive to vote sincerely.

Implementing voting rules

- Assume that all voters have common knowledge and are rational with no limit on computational power. Try to design a voting rule such that everyone has incentive to vote sincerely.
- By Gibbard-Satterthwaite this is a vain hope. One idea is to choose a voting rule, then hope to design a mechanism (a "game form") such that the equilibrium outcomes of the game are the same as the sincere outcome of the given voting rule.

Implementing voting rules

- Assume that all voters have common knowledge and are rational with no limit on computational power. Try to design a voting rule such that everyone has incentive to vote sincerely.
- By Gibbard-Satterthwaite this is a vain hope. One idea is to choose a voting rule, then hope to design a mechanism (a "game form") such that the equilibrium outcomes of the game are the same as the sincere outcome of the given voting rule.
- If we use Nash equilibrium, this again implies dictatorship (Maskin). However other solution concepts exist that can be implemented in this way. They may require enormous computational power.

Implementing voting rules

- Assume that all voters have common knowledge and are rational with no limit on computational power. Try to design a voting rule such that everyone has incentive to vote sincerely.
- By Gibbard-Satterthwaite this is a vain hope. One idea is to choose a voting rule, then hope to design a mechanism (a "game form") such that the equilibrium outcomes of the game are the same as the sincere outcome of the given voting rule.
- If we use Nash equilibrium, this again implies dictatorship (Maskin). However other solution concepts exist that can be implemented in this way. They may require enormous computational power.
- The mechanism announced to players must be just a voting rule. Using one to implement another still leaves the question: which one are we trying to implement, and why?

Manipulability: why it can be a good thing

- A dictatorship is not manipulable, and Gibbard-Satterthwaite says essentially nothing else is. So minimizing manipulability has major welfare consequences and should not be the only goal.

Manipulability: why it can be a good thing

- A dictatorship is not manipulable, and Gibbard-Satterthwaite says essentially nothing else is. So minimizing manipulability has major welfare consequences and should not be the only goal.
- Allowing manipulation can give voters more expressivity by restoring information lost in the voting rule (for example, full preference order, intensity of preference). Lehtinen (Public Choice, 2007; European J. Political Economy, 2008) argues via simulations that strategic voting can improve overall social welfare.

Manipulability: why it can be a good thing

- A dictatorship is not manipulable, and Gibbard-Satterthwaite says essentially nothing else is. So minimizing manipulability has major welfare consequences and should not be the only goal.
- Allowing manipulation can give voters more expressivity by restoring information lost in the voting rule (for example, full preference order, intensity of preference). Lehtinen (Public Choice, 2007; European J. Political Economy, 2008) argues via simulations that strategic voting can improve overall social welfare.
- Dowding and van Hees (British J. Politics, 2008) argue that encouraging strategic voting has many benefits for democracy. Buchanan and Yeo (Public Choice, 2006) argue that in fact all voting is strategic.

The main problem

- Find a measure of social welfare.

The main problem

- Find a measure of social welfare.
- Design a voting mechanism to maximize this measure when all voters are strategic.

The main problem

- Find a measure of social welfare.
- Design a voting mechanism to maximize this measure when all voters are strategic.
- There are only m outcomes so for a fixed voting situation each measure won't differentiate much between rules, but averaged over all voting situations it might.

The main problem

- Find a measure of social welfare.
- Design a voting mechanism to maximize this measure when all voters are strategic.
- There are only m outcomes so for a fixed voting situation each measure won't differentiate much between rules, but averaged over all voting situations it might.
- This requires some way of dealing with multiple equilibria, since there are very many in such situations.

The main problem

- Find a measure of social welfare.
- Design a voting mechanism to maximize this measure when all voters are strategic.
- There are only m outcomes so for a fixed voting situation each measure won't differentiate much between rules, but averaged over all voting situations it might.
- This requires some way of dealing with multiple equilibria, since there are very many in such situations.
- One candidate: price of anarchy — ratio of best case welfare by central planner to worst case welfare in an equilibrium.