Iterated regret minimization for voting games

Mark C. Wilson
(joint with Miranda Emery, University of Auckland)

Department of Computer Science
University of Auckland
www.cs.auckland.ac.nz/~mcw/

COMSOC 2014, Pittsburgh, 2014-06-25

Background

- Voting games have (too) many Nash equilibria. Some of these are socially very undesirable.

Background

- Voting games have (too) many Nash equilibria. Some of these are socially very undesirable.
- Example: under common voting rules, if all voters have the same preference order but decide to submit the reversed one, this is a Nash equilibrium.

Background

- Voting games have (too) many Nash equilibria. Some of these are socially very undesirable.
- Example: under common voting rules, if all voters have the same preference order but decide to submit the reversed one, this is a Nash equilibrium.
- Several other solution concepts have been studied in the context of voting games. We consider simultaneous voting (many recent papers consider sequential voting).

Background

- Voting games have (too) many Nash equilibria. Some of these are socially very undesirable.
- Example: under common voting rules, if all voters have the same preference order but decide to submit the reversed one, this is a Nash equilibrium.
- Several other solution concepts have been studied in the context of voting games. We consider simultaneous voting (many recent papers consider sequential voting).
- We investigate the new solution concept iterated regret minimization (Halpern & Pass, GEB2012, 184–207) in this context.

Regret

images/regret_trooper-490x392.jpg

Regret minimization

- Regret minimization is an important concept in decision theory, at least since Savage (1951).

Regret minimization

- Regret minimization is an important concept in decision theory, at least since Savage (1951).
- There is substantial evidence that it describes human behaviour better than maximizing expected utility does.

Regret minimization

- Regret minimization is an important concept in decision theory, at least since Savage (1951).
- There is substantial evidence that it describes human behaviour better than maximizing expected utility does.
- The regret for an agent of action α is the difference in payoff between the payoff of the best action and the payoff of α.

Regret minimization

- Regret minimization is an important concept in decision theory, at least since Savage (1951).
- There is substantial evidence that it describes human behaviour better than maximizing expected utility does.
- The regret for an agent of action α is the difference in payoff between the payoff of the best action and the payoff of α.
- In a simultaneous game, the regret of α is computed with respect to a partial profile giving all other players' actions.

Regret minimization

- Regret minimization is an important concept in decision theory, at least since Savage (1951).
- There is substantial evidence that it describes human behaviour better than maximizing expected utility does.
- The regret for an agent of action α is the difference in payoff between the payoff of the best action and the payoff of α.
- In a simultaneous game, the regret of α is computed with respect to a partial profile giving all other players' actions.
- We can then maximize over all possible partial profiles, to get the maximum regret (worst-case regret) of α.

Regret minimization

- Regret minimization is an important concept in decision theory, at least since Savage (1951).
- There is substantial evidence that it describes human behaviour better than maximizing expected utility does.
- The regret for an agent of action α is the difference in payoff between the payoff of the best action and the payoff of α.
- In a simultaneous game, the regret of α is computed with respect to a partial profile giving all other players' actions.
- We can then maximize over all possible partial profiles, to get the maximum regret (worst-case regret) of α.
- Now choose an action α that minimizes maximum regret. This is a rather risk-averse strategy.

Example: regret vs utility

Suppose an agent can make actions $A1, A2$ or $A3$, and opponent can make $O1, O2, O3$.

	O1		O2		O3		Min utility	Max regret
A1	-4	7	4	0	12	0	-4	7
A2	-2	5	3	1	8	4	-2	5
A3	3	0	2	2	1	11	1	11

Table: Payoffs in black, regrets in blue

The maximin utility solution is to play A3, but the minimax regret solution is to play A2. Note that $A2$ is strictly dominated by the mixed strategy $0.65A1 + 0.35A3$.

Basic voting setup

- A finite set V of voters, $n := |V|$.

Basic voting setup

- A finite set V of voters, $n := |V|$.
- A finite set A of alternatives, $m := |A|$.

Basic voting setup

- A finite set V of voters, $n := |V|$.
- A finite set A of alternatives, $m := |A|$.
- $L(A) :=$ set of linear orders on A, representing the possible strict preference orders.

Basic voting setup

- A finite set V of voters, $n := |V|$.
- A finite set A of alternatives, $m := |A|$.
- $L(A) :=$ set of linear orders on A, representing the possible strict preference orders.
- A profile is a map $V \to L(A)$ giving each voter's preference order.

Basic voting setup

- A finite set V of voters, $n := |V|$.
- A finite set A of alternatives, $m := |A|$.
- $L(A) :=$ set of linear orders on A, representing the possible strict preference orders.
- A profile is a map $V \to L(A)$ giving each voter's preference order.
- A voting rule is a mapping taking each profile to a nonempty subset of A (the winners).

Basic voting setup

- A finite set V of voters, $n := |V|$.
- A finite set A of alternatives, $m := |A|$.
- $L(A) :=$ set of linear orders on A, representing the possible strict preference orders.
- A profile is a map $V \to L(A)$ giving each voter's preference order.
- A voting rule is a mapping taking each profile to a nonempty subset of A (the winners).
- A social choice function is a voting rule with no ties (a tiebreaking method has been specified).

Voting games

- A given voting rule and set of players and candidates specifies a game form or mechanism.

Voting games

- A given voting rule and set of players and candidates specifies a game form or mechanism.
- Once payoffs are described, this gives a game. We consider the simultaneous case. This is a typical strategic form game with ordinal payoffs.

Voting games

- A given voting rule and set of players and candidates specifies a game form or mechanism.
- Once payoffs are described, this gives a game. We consider the simultaneous case. This is a typical strategic form game with ordinal payoffs.
- Payoffs in voting games are usually described ordinally. Our desired solution concept requires more information. We need to assign utilities of candidates to voters. One way is to choose utilities uniformly and independently from $[0, 1]$ (leads to Impartial Culture).

Voting games

- A given voting rule and set of players and candidates specifies a game form or mechanism.
- Once payoffs are described, this gives a game. We consider the simultaneous case. This is a typical strategic form game with ordinal payoffs.
- Payoffs in voting games are usually described ordinally. Our desired solution concept requires more information. We need to assign utilities of candidates to voters. One way is to choose utilities uniformly and independently from $[0, 1]$ (leads to Impartial Culture).
- The possible actions of a voter are to submit any of the $m!$ preference orders.

Iterated regret minimization

- We iteratively eliminate all strategies that do not minimize maximum regret.

Iterated regret minimization

- We iteratively eliminate all strategies that do not minimize maximum regret.
- At each iteration, we imagine that each other player is as “rational” as we were in the last iteration.

Iterated regret minimization

- We iteratively eliminate all strategies that do not minimize maximum regret.
- At each iteration, we imagine that each other player is as "rational" as we were in the last iteration.
- This is analogous to admissibility or dominance solvability, where players iteratively eliminate weakly dominated strategies.

Iterated regret minimization

- We iteratively eliminate all strategies that do not minimize maximum regret.
- At each iteration, we imagine that each other player is as "rational" as we were in the last iteration.
- This is analogous to admissibility or dominance solvability, where players iteratively eliminate weakly dominated strategies.
- Some such iterated deletion solution concepts have been controversial. Halpern and Pass give detailed justification for this as a solution concept, and state "It seems particularly appealing when considering inexperienced but intelligent players that play a game for the first time."

Relation to dominance solvability

- IRM is analogous to dominance solvability, in that we iteratively eliminate strategies, increasing our expectations of rationality of the other players.

Relation to dominance solvability

- IRM is analogous to dominance solvability, in that we iteratively eliminate strategies, increasing our expectations of rationality of the other players.
- The difference is that with DS, we eliminate weakly dominated strategies.

Relation to dominance solvability

- IRM is analogous to dominance solvability, in that we iteratively eliminate strategies, increasing our expectations of rationality of the other players.
- The difference is that with DS, we eliminate weakly dominated strategies.
- For voting games, DS seems to be much less decisive than IRM. For plurality rule, the profiles where DS gives a unique winner have been characterized (Dhillon-Lockwood 2004) and are relatively few.

Relation to dominance solvability

- IRM is analogous to dominance solvability, in that we iteratively eliminate strategies, increasing our expectations of rationality of the other players.
- The difference is that with DS, we eliminate weakly dominated strategies.
- For voting games, DS seems to be much less decisive than IRM. For plurality rule, the profiles where DS gives a unique winner have been characterized (Dhillon-Lockwood 2004) and are relatively few.
- In contrast, IRM always outputs the sincere outcome.

Theoretical results

- The sincere vote can have maximum regret (for example, for IRV).

Theoretical results

- The sincere vote can have maximum regret (for example, for IRV).
- For a wide class of rules including positional scoring rules but not including IRV, in the first iteration each player places her top candidate first and bottom candidate last. Hence for $m = 3$, they vote sincerely.

Theoretical results

- The sincere vote can have maximum regret (for example, for IRV).
- For a wide class of rules including positional scoring rules but not including IRV, in the first iteration each player places her top candidate first and bottom candidate last. Hence for $m = 3$, they vote sincerely.
- For plurality and veto, the IRM outcome always equals the sincere outcome.

Theoretical results

- The sincere vote can have maximum regret (for example, for IRV).
- For a wide class of rules including positional scoring rules but not including IRV, in the first iteration each player places her top candidate first and bottom candidate last. Hence for $m = 3$, they vote sincerely.
- For plurality and veto, the IRM outcome always equals the sincere outcome.
- These results hold for all tiebreaking methods and all utilities consistent with the voters' preference order.

Example: IRM outcome

Voter	Preference order
v_1	ACDB
v_2	BCAD
v_3	BDCA
v_4	CDAB
v_5	DABC

Table: A profile where IRM elects the sincere loser.

Sincere social Borda ranking is $C \succ D \succ A \succ B$. Under IRM assuming Borda utilities, B, the sincere loser, wins.

Welfare considerations

- How good is the IRM outcome, socially?

Welfare considerations

- How good is the IRM outcome, socially?
- We can measure the utilitarian welfare (the sum of voters' utilities) of a given outcome, or the egalitarian welfare (minimum utility of a voter).

Welfare considerations

- How good is the IRM outcome, socially?
- We can measure the utilitarian welfare (the sum of voters' utilities) of a given outcome, or the egalitarian welfare (minimum utility of a voter).
- Another measure is net satisfaction, the difference between the number of voters preferring the IRM outcome and the number preferring the sincere outcome.

Welfare considerations

- How good is the IRM outcome, socially?
- We can measure the utilitarian welfare (the sum of voters' utilities) of a given outcome, or the egalitarian welfare (minimum utility of a voter).
- Another measure is net satisfaction, the difference between the number of voters preferring the IRM outcome and the number preferring the sincere outcome.
- One could also measure the change in Condorcet efficiency (how often the Condorcet winner is elected, when it exists).

Experimental results I

Table: Summary statistics with $m = 4$: Borda (B) and 2-approval (2A) rule and utilities. % S gives percentage of simulated elections where IRM winner is sincere winner.

Util/rule	n	Mean N	Mean $\overline{U}$	Mean $\overline{E}$	% S
B/B	5	0.010	-0.004	-0.016	90.4
B/B	30	0.0003	-0.0002	0.0	99
2A/B	5	-0.014	-0.027	-0.020	88.9
2A/B	30	-0.00001	-0.001	-0.00001	100
2A/2A	5	-0.013	-0.013	0.0	100
2A/2A	30	0.0	0.0	0.0	100

Experimental results II

Figure: Rank of IRM winner: Borda utilities, IRV, $m = 3, n = 5$. Left bars, U; right, E.

images/ranksIRV.png

Welfare: the broader view

- Many (computational) social choice theorists have focused on measuring/reducing manipulability of voting rules.

Welfare: the broader view

- Many (computational) social choice theorists have focused on measuring/reducing manipulability of voting rules.
- Gibbard-Satterthwaite: under mild hypotheses, the minimum is attained only at the dictatorial rule. How should we trade off dictatorship with manipulability?

Welfare: the broader view

- Many (computational) social choice theorists have focused on measuring/reducing manipulability of voting rules.
- Gibbard-Satterthwaite: under mild hypotheses, the minimum is attained only at the dictatorial rule. How should we trade off dictatorship with manipulability?
- My opinion: who cares about manipulation? Surely overall system performance (welfare) is more important.

Welfare: the broader view

- Many (computational) social choice theorists have focused on measuring/reducing manipulability of voting rules.
- Gibbard-Satterthwaite: under mild hypotheses, the minimum is attained only at the dictatorial rule. How should we trade off dictatorship with manipulability?
- My opinion: who cares about manipulation? Surely overall system performance (welfare) is more important.
- Under the assumption of sincere behaviour, we can compare rules with respect to a fixed welfare measure. For example, Borda maximizes expected utilitarian welfare under Impartial Culture (Boutilier *et al*, EC 2012).

Welfare: the broader view

- Many (computational) social choice theorists have focused on measuring/reducing manipulability of voting rules.
- Gibbard-Satterthwaite: under mild hypotheses, the minimum is attained only at the dictatorial rule. How should we trade off dictatorship with manipulability?
- My opinion: who cares about manipulation? Surely overall system performance (welfare) is more important.
- Under the assumption of sincere behaviour, we can compare rules with respect to a fixed welfare measure. For example, Borda maximizes expected utilitarian welfare under Impartial Culture (Boutilier *et al*, EC 2012).
- We can also compare rules under various other models of voter behaviour. So far, many authors have found (sometimes to their surprise) that welfare is badly affected in the worst case, but on average it is often increased by strategic behaviour.

Recent experimental work on welfare and voting games

- Lehtinen (Euro. J. Pol. Econ 2008, Public Choice 2007): for plurality/Borda/approval, a kind of myopic best response increases utilitarian welfare.

Recent experimental work on welfare and voting games

- Lehtinen (Euro. J. Pol. Econ 2008, Public Choice 2007): for plurality/Borda/approval, a kind of myopic best response increases utilitarian welfare.
- Xia & Conitzer (AAAI 2010): sequential voting under plurality/veto results in positive net satisfaction.

Recent experimental work on welfare and voting games

- Lehtinen (Euro. J. Pol. Econ 2008, Public Choice 2007): for plurality/Borda/approval, a kind of myopic best response increases utilitarian welfare.
- Xia & Conitzer (AAAI 2010): sequential voting under plurality/veto results in positive net satisfaction.
- Thompson, Leyton-Brown, Lev & Rosenschein (AAMAS 2013): for plurality, truth-biased Nash equilibrium increases Condorcet efficiency.

Recent experimental work on welfare and voting games

- Lehtinen (Euro. J. Pol. Econ 2008, Public Choice 2007): for plurality/Borda/approval, a kind of myopic best response increases utilitarian welfare.
- Xia & Conitzer (AAAI 2010): sequential voting under plurality/veto results in positive net satisfaction.
- Thompson, Leyton-Brown, Lev & Rosenschein (AAMAS 2013): for plurality, truth-biased Nash equilibrium increases Condorcet efficiency.
- Meir, Lev & Rosenschein (COMSOC 2014): for plurality, under a wide class of heuristics, many welfare measures increase.

Recent experimental work on welfare and voting games

- Lehtinen (Euro. J. Pol. Econ 2008, Public Choice 2007): for plurality/Borda/approval, a kind of myopic best response increases utilitarian welfare.
- Xia & Conitzer (AAAI 2010): sequential voting under plurality/veto results in positive net satisfaction.
- Thompson, Leyton-Brown, Lev & Rosenschein (AAMAS 2013): for plurality, truth-biased Nash equilibrium increases Condorcet efficiency.
- Meir, Lev & Rosenschein (COMSOC 2014): for plurality, under a wide class of heuristics, many welfare measures increase.
- Wilson and various students (unpublished): some generalizations and overlaps with above, similar results.

Recent experimental work on welfare and voting games

- Lehtinen (Euro. J. Pol. Econ 2008, Public Choice 2007): for plurality/Borda/approval, a kind of myopic best response increases utilitarian welfare.
- Xia & Conitzer (AAAI 2010): sequential voting under plurality/veto results in positive net satisfaction.
- Thompson, Leyton-Brown, Lev & Rosenschein (AAMAS 2013): for plurality, truth-biased Nash equilibrium increases Condorcet efficiency.
- Meir, Lev & Rosenschein (COMSOC 2014): for plurality, under a wide class of heuristics, many welfare measures increase.
- Wilson and various students (unpublished): some generalizations and overlaps with above, similar results.
- This paper: for Borda, IRM has modestly positive net satisfaction, and trivially negative net utilitarian welfare.

Conclusion

- IRM has reasonably good justification as a solution concept, and is more decisive than dominance solvability.

Conclusion

- IRM has reasonably good justification as a solution concept, and is more decisive than dominance solvability.
- IRM has a strong bias toward sincere outcomes.

Conclusion

- IRM has reasonably good justification as a solution concept, and is more decisive than dominance solvability.
- IRM has a strong bias toward sincere outcomes.
- IRM has small (sometimes positive, sometimes negative) effects on overall welfare.

Invitation

- Centre for Mathematical Social Sciences at University of Auckland is holding 6th annual summer workshop.
- This year: during week 8-12 December 2014. Theme is Social Networks, but all areas welcome.
- Invited speakers: Matthew Jackson (Stanford); Damon Centola (U. Penn.).
- Previous visitors from this community: J. Röthe, E. Elkind, P. Faliszewski, J. Lang, T. Walsh, G. Erdelyi, D. Baumeister, M. Brill, N. Betzler, C. Puppe, W. Zwicker, Y. Zick, ...
- No fee to participate, some possible (small) funding support.