What is Computational Social Choice?

Mark C. Wilson
www.cs.auckland.ac.nz/˜mcw/blog/

Department of Computer Science
University of Auckland

UoA CS Seminar, 2010-10-20

References

Computational microeconomics

Social choice

Game theory and mechanism design

Social choice mechanisms

Centre for Mathematical Social Sciences

- Departmental centre in Department of Mathematics, from 2010.

Centre for Mathematical Social Sciences

- Departmental centre in Department of Mathematics, from 2010.
- Members from Maths, CS, Stats, Econ, Philosophy.

Centre for Mathematical Social Sciences

- Departmental centre in Department of Mathematics, from 2010.
- Members from Maths, CS, Stats, Econ, Philosophy.
- Aim to hold an annual summer workshop.

Centre for Mathematical Social Sciences

- Departmental centre in Department of Mathematics, from 2010.
- Members from Maths, CS, Stats, Econ, Philosophy.
- Aim to hold an annual summer workshop.
- http://cmss.auckland.ac.nz.

Good survey articles

Nisa2007 N. Nisan et al. Algorithmic Game Theory (book, 2007).

Good survey articles

Nisa2007 N. Nisan et al. Algorithmic Game Theory (book, 2007).

Shoh2008 Y. Shoham. Computer Science and Game Theory. CACM Aug 2008.

Good survey articles

Nisa2007 N. Nisan et al. Algorithmic Game Theory (book, 2007).

Shoh2008 Y. Shoham. Computer Science and Game Theory. CACM Aug 2008.

DGP2009 C. Daskalaskis et al. The Complexity of Computing a Nash Equilibrium. CACM Feb 2009.

Good survey articles

Nisa2007 N. Nisan et al. Algorithmic Game Theory (book, 2007).

Shoh2008 Y. Shoham. Computer Science and Game Theory. CACM Aug 2008.

DGP2009 C. Daskalaskis et al. The Complexity of Computing a Nash Equilibrium. CACM Feb 2009.

Roug2010 T. Roughgarden. Algorithmic Game Theory. CACM July 2010.

Good survey articles

Nisa2007 N. Nisan et al. Algorithmic Game Theory (book, 2007).

Shoh2008 Y. Shoham. Computer Science and Game Theory. CACM Aug 2008.

DGP2009 C. Daskalaskis et al. The Complexity of Computing a Nash Equilibrium. CACM Feb 2009.

Roug2010 T. Roughgarden. Algorithmic Game Theory. CACM July 2010.

Anth2010 G. Anthes. Mechanism Design Meets Computer Science. CACM Aug 2010.

Good survey articles

Nisa2007 N. Nisan et al. Algorithmic Game Theory (book, 2007).

Shoh2008 Y. Shoham. Computer Science and Game Theory. CACM Aug 2008.

DGP2009 C. Daskalaskis et al. The Complexity of Computing a Nash Equilibrium. CACM Feb 2009.

Roug2010 T. Roughgarden. Algorithmic Game Theory. CACM July 2010.

Anth2010 G. Anthes. Mechanism Design Meets Computer Science. CACM Aug 2010.

Chev2007 Y. Chevaleyre et al. An Introduction to Computational Social Choice. Proceedings SOFSEM 2007.

More specialized papers

CSL2007 Conitzer, Sandholm, Lang. When are elections with few candidates hard to manipulate? JACM 2007.

More specialized papers

CSL2007 Conitzer, Sandholm, Lang. When are elections with few candidates hard to manipulate? JACM 2007.

MPJR2010 R. Meir, M, Polukarov, N. Jennings, J. Rosenschein. Convergence to equilibria in plurality voting. Proc AAAI 2010.

More specialized papers

CSL2007 Conitzer, Sandholm, Lang. When are elections with few candidates hard to manipulate? JACM 2007.

MPJR2010 R. Meir, M, Polukarov, N. Jennings, J. Rosenschein. Convergence to equilibria in plurality voting. Proc AAAI 2010.

RPW2010 R. Reyhani, G. Pritchard, M. Wilson. A new measure of manipulability of voting rules. Submitted, 2010.

A new field is emerging

- In the last decade, computer science and game theory have collided, and a new interdisciplinary field is forming.

A new field is emerging

- In the last decade, computer science and game theory have collided, and a new interdisciplinary field is forming.
- Big philosophical idea: explore the fundamental tension between efficiency (economic or algorithmic) and compatibility with self-interest.

A new field is emerging

- In the last decade, computer science and game theory have collided, and a new interdisciplinary field is forming.
- Big philosophical idea: explore the fundamental tension between efficiency (economic or algorithmic) and compatibility with self-interest.
- "A trend has emerged towards interdisciplinary research involving all of decision theory, game theory, social choice theory, and welfare economics on the one hand, and computer science, artificial intelligence, multiagent systems, operations research, and computational logic on the other."

A new field is emerging

- In the last decade, computer science and game theory have collided, and a new interdisciplinary field is forming.
- Big philosophical idea: explore the fundamental tension between efficiency (economic or algorithmic) and compatibility with self-interest.
- "A trend has emerged towards interdisciplinary research involving all of decision theory, game theory, social choice theory, and welfare economics on the one hand, and computer science, artificial intelligence, multiagent systems, operations research, and computational logic on the other."
- Commercial problems have dominated research on the CS side, but a shift toward a broader viewpoint is evident.

A new field is emerging

- In the last decade, computer science and game theory have collided, and a new interdisciplinary field is forming.
- Big philosophical idea: explore the fundamental tension between efficiency (economic or algorithmic) and compatibility with self-interest.
- "A trend has emerged towards interdisciplinary research involving all of decision theory, game theory, social choice theory, and welfare economics on the one hand, and computer science, artificial intelligence, multiagent systems, operations research, and computational logic on the other."
- Commercial problems have dominated research on the CS side, but a shift toward a broader viewpoint is evident.
- No official name: "computational (micro)economics", "algorithmic game theory", "algorithmic mechanism design"?

Where did it come from?

- The enormous growth in the use of the Internet as a major platform for social and economic interactions.

Where did it come from?

- The enormous growth in the use of the Internet as a major platform for social and economic interactions.
- Strategic behaviour and distributed information aggregation in computer and communications networks.

Where did it come from?

- The enormous growth in the use of the Internet as a major platform for social and economic interactions.
- Strategic behaviour and distributed information aggregation in computer and communications networks.
- Many modern computer science applications involve multiagent systems of autonomous decision-makers (robots, artificial life, bidding agents, . . .).

Where did it come from?

- The enormous growth in the use of the Internet as a major platform for social and economic interactions.
- Strategic behaviour and distributed information aggregation in computer and communications networks.
- Many modern computer science applications involve multiagent systems of autonomous decision-makers (robots, artificial life, bidding agents, . . .).
  - Formerly, they were considered in isolation or as cooperating in a distributed system.

Where did it come from?

- The enormous growth in the use of the Internet as a major platform for social and economic interactions.
- Strategic behaviour and distributed information aggregation in computer and communications networks.
- Many modern computer science applications involve multiagent systems of autonomous decision-makers (robots, artificial life, bidding agents, . . .).
  - Formerly, they were considered in isolation or as cooperating in a distributed system.
  - More recently, there are many situations where they have their own "selfish" preferences, which may conflict with those of other agents.

Where did it come from?

- The enormous growth in the use of the Internet as a major platform for social and economic interactions.
- Strategic behaviour and distributed information aggregation in computer and communications networks.
- Many modern computer science applications involve multiagent systems of autonomous decision-makers (robots, artificial life, bidding agents, . . .).
  - Formerly, they were considered in isolation or as cooperating in a distributed system.
  - More recently, there are many situations where they have their own "selfish" preferences, which may conflict with those of other agents.
  - They cooperate/compete by playing a strategic game.

Internet applications

- Auctions (e.g. Google AdWords). This is the most-studied application and has had the biggest financial impact. Yahoo, Google and Microsoft employ big-name researchers just to study such problems.

Internet applications

- Auctions (e.g. Google AdWords). This is the most-studied application and has had the biggest financial impact. Yahoo, Google and Microsoft employ big-name researchers just to study such problems.
- Recommender systems, collaborative filtering (e.g. Amazon, Netflix , . . .)

Internet applications

- Auctions (e.g. Google AdWords). This is the most-studied application and has had the biggest financial impact. Yahoo, Google and Microsoft employ big-name researchers just to study such problems.
- Recommender systems, collaborative filtering (e.g. Amazon, Netflix , . . .)
- Prediction markets.

Internet applications

- Auctions (e.g. Google AdWords). This is the most-studied application and has had the biggest financial impact. Yahoo, Google and Microsoft employ big-name researchers just to study such problems.
- Recommender systems, collaborative filtering (e.g. Amazon, Netflix , . . .)
- Prediction markets.
- Peer-to-peer networks, network routing.

Internet applications

- Auctions (e.g. Google AdWords). This is the most-studied application and has had the biggest financial impact. Yahoo, Google and Microsoft employ big-name researchers just to study such problems.
- Recommender systems, collaborative filtering (e.g. Amazon, Netflix , . . .)
- Prediction markets.
- Peer-to-peer networks, network routing.
- Social networking sites, reputation.

Internet applications

- Auctions (e.g. Google AdWords). This is the most-studied application and has had the biggest financial impact. Yahoo, Google and Microsoft employ big-name researchers just to study such problems.
- Recommender systems, collaborative filtering (e.g. Amazon, Netflix , . . .)
- Prediction markets.
- Peer-to-peer networks, network routing.
- Social networking sites, reputation.
- Electronic voting?

Some phrases to give the flavour of the field

- ACM Conference on Electronic Commerce, Symposium on Algorithmic Game Theory, Workshop on Computational Social Choice

Some phrases to give the flavour of the field

- ACM Conference on Electronic Commerce, Symposium on Algorithmic Game Theory, Workshop on Computational Social Choice
- Papers: The Complexity of Computing Nash Equilibria, Selfish Routing and the Price of Anarchy, Approximate Mechanism Design without Money, Truthful Fair Division, Combinatorial Auctions

Contributions flow both ways

- Econ $\rightarrow$ CS: distributed computing and networking protocols (such as TCP-IP) have traditionally assumed that components cooperate. However incentives and selfish preferences cannot be ignored. Rational behaviour can lead to suboptimal outcomes if not controlled.

Contributions flow both ways

- Econ $\to$ CS: distributed computing and networking protocols (such as TCP-IP) have traditionally assumed that components cooperate. However incentives and selfish preferences cannot be ignored. Rational behaviour can lead to suboptimal outcomes if not controlled.
- CS $\to$ Econ: traditional models use mathematical existence results such as fixed point theorems. However computational and communication complexity cannot be ignored. Strategies and solutions may not be practically computable.

Basic setup of social choice

- A finite set of m alternatives and n voters. Each voter has a preference over alternatives.

Basic setup of social choice

- A finite set of m alternatives and n voters. Each voter has a preference over alternatives.
- A social choice correspondence aggregates the preferences and outputs a set of alternatives, the winners; a social welfare function outputs a full ranking.

Basic setup of social choice

- A finite set of m alternatives and n voters. Each voter has a preference over alternatives.
- A social choice correspondence aggregates the preferences and outputs a set of alternatives, the winners; a social welfare function outputs a full ranking.
- Used for millenia in human political decision-making (voting, elections, planning, where to build an airport, allocation of objects to people, . . .).

Basic setup of social choice

- A finite set of m alternatives and n voters. Each voter has a preference over alternatives.
- A social choice correspondence aggregates the preferences and outputs a set of alternatives, the winners; a social welfare function outputs a full ranking.
- Used for millenia in human political decision-making (voting, elections, planning, where to build an airport, allocation of objects to people, ...).
- Very often we require only a single winner (social choice function), and tiebreaking procedures are almost always needed. Randomized tiebreaking leads to objects that are not strictly speaking social choice functions.

Some social choice functions

- Scoring rules: fix a vector $1 = w_1 \geq w_2 \geq \cdots \geq w_m = 0$. Voter awards w_1 points to its most preferred alternative, w_2 to second, etc. Highest total score wins. Famous examples: plurality ($w_i = 0$ for $i > 1$); Borda (weights are equally spaced); veto ($w_i = 1$ for $i < m$).

Some social choice functions

- Scoring rules: fix a vector $1 = w_1 \geq w_2 \geq \cdots \geq w_m = 0$. Voter awards w_1 points to its most preferred alternative, w_2 to second, etc. Highest total score wins. Famous examples: plurality ($w_i = 0$ for $i > 1$); Borda (weights are equally spaced); veto ($w_i = 1$ for $i < m$).
- Condorcet rules: if the majority relation has a clear winner, choose it. Otherwise choose something else. Example: Copeland rule: award ± 1 for each pairwise majority victory/defeat, highest total wins.

Some social choice functions

- Scoring rules: fix a vector $1 = w_1 \geq w_2 \geq \cdots \geq w_m = 0$. Voter awards w_1 points to its most preferred alternative, w_2 to second, etc. Highest total score wins. Famous examples: plurality ($w_i = 0$ for $i > 1$); Borda (weights are equally spaced); veto ($w_i = 1$ for $i < m$).
- Condorcet rules: if the majority relation has a clear winner, choose it. Otherwise choose something else. Example: Copeland rule: award ± 1 for each pairwise majority victory/defeat, highest total wins.
- Dictatorship: one voter decides the result, irrespective of the preferences of others.

Classic paradoxes of social choice theory

- Condorcet: the pairwise majority relation can be cyclic.

None is devastating although some may have uncomfortable political implications.

Classic paradoxes of social choice theory

- Condorcet: the pairwise majority relation can be cyclic.
- Arrow: a few simple axioms lead to dictatorship.

None is devastating although some may have uncomfortable political implications.

Classic paradoxes of social choice theory

- Condorcet: the pairwise majority relation can be cyclic.
- Arrow: a few simple axioms lead to dictatorship.
- Simpson: the winner in each of two subgroups of voters may not win in the whole group.

None is devastating although some may have uncomfortable political implications.

Classic paradoxes of social choice theory

- Condorcet: the pairwise majority relation can be cyclic.
- Arrow: a few simple axioms lead to dictatorship.
- Simpson: the winner in each of two subgroups of voters may not win in the whole group.
- Participation: the winner may not remain the winner when extra voters rank it first.

None is devastating although some may have uncomfortable political implications.

(Noncooperative) game theory

- Founded by von Neumann, Nash, et al. in 1940s and 1950s.

(Noncooperative) game theory

- Founded by von Neumann, Nash, et al. in 1940s and 1950s.
- Each player has a finite number of actions; a profile is a choice of one for each player. The utility gained by each player depends only on the profile.

(Noncooperative) game theory

- Founded by von Neumann, Nash, et al. in 1940s and 1950s.
- Each player has a finite number of actions; a profile is a choice of one for each player. The utility gained by each player depends only on the profile.
- Very influential in economics, evolutionary biology, international relations, political sciences, ...

(Noncooperative) game theory

- Founded by von Neumann, Nash, et al. in 1940s and 1950s.
- Each player has a finite number of actions; a profile is a choice of one for each player. The utility gained by each player depends only on the profile.
- Very influential in economics, evolutionary biology, international relations, political sciences, ...
- Classic examples: Chicken, Battle of the Sexes, Prisoners' Dilemma. Suboptimal outcomes can occur because of misalignment of individual incentives, but sometimes don't. It depends on the structure of the game.

Example: load balancing

- We have n players each with one ball, and n bins. Each player must throw its ball into a bin. Moves are simultaneous. The cost to each player is the number of balls in its bin.

Example: load balancing

- We have n players each with one ball, and n bins. Each player must throw its ball into a bin. Moves are simultaneous. The cost to each player is the number of balls in its bin.
- One possible outcome: each ball goes in a unique bin, every player incurs cost 1.

Example: load balancing

- We have n players each with one ball, and n bins. Each player must throw its ball into a bin. Moves are simultaneous. The cost to each player is the number of balls in its bin.
- One possible outcome: each ball goes in a unique bin, every player incurs cost 1.
- The obvious strategy of uniformly randomly choosing a bin has the same expected cost for each player, but the worst-off player has cost of order $\log n / \log \log n$.

Example: load balancing

- We have n players each with one ball, and n bins. Each player must throw its ball into a bin. Moves are simultaneous. The cost to each player is the number of balls in its bin.
- One possible outcome: each ball goes in a unique bin, every player incurs cost 1.
- The obvious strategy of uniformly randomly choosing a bin has the same expected cost for each player, but the worst-off player has cost of order $\log n / \log \log n$.
- Each of these strategy profiles is a Nash equilibrium: given that all other players play the strategy, no player has incentive to deviate. However it is not a dominant strategy equilibrium: if some players deviate, sticking with the strategy may be bad.

Mechanism design

- “Reverse-engineering” in game theory. Sveriges Riksbank (“Nobel”) prize to Hurwicz, Myerson, Maskin in 2007. Applied to raise billions of euros in electromagnetic spectrum auctions.

Mechanism design

- "Reverse-engineering" in game theory. Sveriges Riksbank ("Nobel") prize to Hurwicz, Myerson, Maskin in 2007. Applied to raise billions of euros in electromagnetic spectrum auctions.
- A mechanism is a game with a special player, the designer. The designer's goal is to implement some fixed allocation rule R_1.

Mechanism design

- "Reverse-engineering" in game theory. Sveriges Riksbank ("Nobel") prize to Hurwicz, Myerson, Maskin in 2007. Applied to raise billions of euros in electromagnetic spectrum auctions.
- A mechanism is a game with a special player, the designer. The designer's goal is to implement some fixed allocation rule R_1.
- Each other player has private utility information called its type θ, and must report some type $\hat{\theta}$. Let Θ be the profile of all players types. If designer knew Θ or players always report Θ, the job is easy. However, players can strategically lie, $\hat{\Theta} \neq \Theta$.

The University of Auckland NEW ZEALAND

Mechanism design

- "Reverse-engineering" in game theory. Sveriges Riksbank ("Nobel") prize to Hurwicz, Myerson, Maskin in 2007. Applied to raise billions of euros in electromagnetic spectrum auctions.
- A mechanism is a game with a special player, the designer. The designer's goal is to implement some fixed allocation rule R_1.
- Each other player has private utility information called its type θ, and must report some type $\hat{\theta}$. Let Θ be the profile of all players types. If designer knew Θ or players always report Θ, the job is easy. However, players can strategically lie, $\hat{\Theta} \neq \Theta$.
- The designer announces an allocation rule R_2 (including transfer payments), and uses this on the reported types. Designer aims for $R_2(\hat{\Theta}) = R_1(\Theta)$.

Truthful mechanisms

- Some mechanisms have the property that each player has a dominant strategy to truthfully reveal its type. In other words, there are really no strategic considerations. Each player has a best move no matter what the other players do.

Truthful mechanisms

- Some mechanisms have the property that each player has a dominant strategy to truthfully reveal its type. In other words, there are really no strategic considerations. Each player has a best move no matter what the other players do.
- Classic example: second-price (Vickrey) auction. The winner pays the second-highest bid.

Truthful mechanisms

- Some mechanisms have the property that each player has a dominant strategy to truthfully reveal its type. In other words, there are really no strategic considerations. Each player has a best move no matter what the other players do.
- Classic example: second-price (Vickrey) auction. The winner pays the second-highest bid.
- Classic nonexample: first-price auction. The winner pays its own bid.

Truthful mechanisms

- Some mechanisms have the property that each player has a dominant strategy to truthfully reveal its type. In other words, there are really no strategic considerations. Each player has a best move no matter what the other players do.
- Classic example: second-price (Vickrey) auction. The winner pays the second-highest bid.
- Classic nonexample: first-price auction. The winner pays its own bid.
- Important nonexample: (later) nondictatorial social choice functions.

Example: sealed-bid second price auction

- Player i has a private utility v_i (in common currency) for a fixed object to be auctioned. Players bid simultaneously, once.

Example: sealed-bid second price auction

- Player i has a private utility v_i (in common currency) for a fixed object to be auctioned. Players bid simultaneously, once.
- The allocation rule R_1 is "give the object to the player with highest v_i, and charge him v_i".

Example: sealed-bid second price auction

- Player i has a private utility v_i (in common currency) for a fixed object to be auctioned. Players bid simultaneously, once.
- The allocation rule R_1 is "give the object to the player with highest v_i, and charge him v_i".
- If we announce this then players have an incentive to bid lower than v_i (how much depends on their perception of the bids of other players - the game is complicated).

Example: sealed-bid second price auction

- Player i has a private utility v_i (in common currency) for a fixed object to be auctioned. Players bid simultaneously, once.
- The allocation rule R_1 is "give the object to the player with highest v_i, and charge him v_i".
- If we announce this then players have an incentive to bid lower than v_i (how much depends on their perception of the bids of other players - the game is complicated).
- However, if we announce R_2: "give the object to the highest bidder, and charge him the second-highest bid", there is no incentive to bid untruthfully and players may as well report v_i.

Example: buying a path in a network

- We aim to route a message from node s to node t in a digraph.

Example: buying a path in a network

- We aim to route a message from node s to node t in a digraph.
- Players are arcs of a digraph, and player e incurs cost c_e if the message path uses e. They will be paid.

Example: buying a path in a network

- We aim to route a message from node s to node t in a digraph.
- Players are arcs of a digraph, and player e incurs cost c_e if the message path uses e. They will be paid.
- If players are truthful, standard shortest path algorithms will optimize social welfare (minimize total cost). However, they have clear incentive to report a higher cost than they actually incur.

Example: buying a path in a network

- We aim to route a message from node s to node t in a digraph.
- Players are arcs of a digraph, and player e incurs cost c_e if the message path uses e. They will be paid.
- If players are truthful, standard shortest path algorithms will optimize social welfare (minimize total cost). However, they have clear incentive to report a higher cost than they actually incur.
- The general Vickrey-Clarke-Groves mechanism yields a nice solution. We pay e zero if e is not in the cheapest path, and otherwise pay its reported cost plus a "bonus" equal to its "contribution": the increase in cost of the cheapest path if e were deleted.

More on VCG mechanism

- The payment "internalizes the externality", and reporting the true cost is a dominant strategy for all players, each of whom is guaranteed to cover its cost.

More on VCG mechanism

- The payment “internalizes the externality”, and reporting the true cost is a dominant strategy for all players, each of whom is guaranteed to cover its cost.
- The total of payments may be very much larger than is optimal under truthful reporting. This can be a major difficulty.

More on VCG mechanism

- The payment "internalizes the externality", and reporting the true cost is a dominant strategy for all players, each of whom is guaranteed to cover its cost.
- The total of payments may be very much larger than is optimal under truthful reporting. This can be a major difficulty.
- Another problem: in combinatorial auctions players bid on bundles of goods (such as spectrum licences), and the underlying optimization problem can be NP-hard.

More on VCG mechanism

- The payment "internalizes the externality", and reporting the true cost is a dominant strategy for all players, each of whom is guaranteed to cover its cost.
- The total of payments may be very much larger than is optimal under truthful reporting. This can be a major difficulty.
- Another problem: in combinatorial auctions players bid on bundles of goods (such as spectrum licences), and the underlying optimization problem can be NP-hard.
- VCG only works when we want to maximize the total utility of the players, not for other measures of welfare.

More on VCG mechanism

- The payment "internalizes the externality", and reporting the true cost is a dominant strategy for all players, each of whom is guaranteed to cover its cost.
- The total of payments may be very much larger than is optimal under truthful reporting. This can be a major difficulty.
- Another problem: in combinatorial auctions players bid on bundles of goods (such as spectrum licences), and the underlying optimization problem can be NP-hard.
- VCG only works when we want to maximize the total utility of the players, not for other measures of welfare.
- There is much research on how to get around these difficulties using approximations.

The CS contribution

- Computational complexity: mechanisms may be arbitrarily complex. Strategies, equilibria, ... may be NP-hard (or worse) to compute. In fact they often are.

The CS contribution

- Computational complexity: mechanisms may be arbitrarily complex. Strategies, equilibria, ... may be NP-hard (or worse) to compute. In fact they often are.
- Approximation algorithms: the standard response to hard optimization problems. Concepts such as approximation ratio.

The CS contribution

- Computational complexity: mechanisms may be arbitrarily complex. Strategies, equilibria, . . . may be NP-hard (or worse) to compute. In fact they often are.
- Approximation algorithms: the standard response to hard optimization problems. Concepts such as approximation ratio.
- Worst-case (non-Bayesian) analysis.

Beyond truthful mechanisms

- Perhaps sincerity is overrated: if the designer cares only about the final allocation, and this can be achieved by untruthful behaviour, then why worry about players telling the truth?

Beyond truthful mechanisms

- Perhaps sincerity is overrated: if the designer cares only about the final allocation, and this can be achieved by untruthful behaviour, then why worry about players telling the truth?
- The main problem is that the outcome of the game is easily predicted only when there is a unique dominant strategy (truthtelling) for all players.

Beyond truthful mechanisms

- Perhaps sincerity is overrated: if the designer cares only about the final allocation, and this can be achieved by untruthful behaviour, then why worry about players telling the truth?
- The main problem is that the outcome of the game is easily predicted only when there is a unique dominant strategy (truthtelling) for all players.
- In general there will be many reasonable "predictions" (usually these are Nash equilibria). Problems: in the worst case Nash equilibria are likely not computable in polynomial time [DGP2009]; there are far too many of them.

Beyond truthful mechanisms

- Perhaps sincerity is overrated: if the designer cares only about the final allocation, and this can be achieved by untruthful behaviour, then why worry about players telling the truth?
- The main problem is that the outcome of the game is easily predicted only when there is a unique dominant strategy (truthtelling) for all players.
- In general there will be many reasonable "predictions" (usually these are Nash equilibria). Problems: in the worst case Nash equilibria are likely not computable in polynomial time [DGP2009]; there are far too many of them.
- Which equilibrium do we look at in order to measure the overall welfare? This leads to ideas such as price of anarchy.

Social choice mechanisms

- The type of a player (voter) is its preference order over the alternatives.

Social choice mechanisms

- The type of a player (voter) is its preference order over the alternatives.
- The designer chooses the social choice function R_1 for aggregating the individual preferences, and reports another one R_2.

Social choice mechanisms

- The type of a player (voter) is its preference order over the alternatives.
- The designer chooses the social choice function R_1 for aggregating the individual preferences, and reports another one R_2.
- The strategic action of each voter is to report a preference order (possibly untruthful).

Social choice mechanisms

- The type of a player (voter) is its preference order over the alternatives.
- The designer chooses the social choice function R_1 for aggregating the individual preferences, and reports another one R_2.
- The strategic action of each voter is to report a preference order (possibly untruthful).
- There are no payments.

Social choice mechanisms

- The type of a player (voter) is its preference order over the alternatives.
- The designer chooses the social choice function R_1 for aggregating the individual preferences, and reports another one R_2.
- The strategic action of each voter is to report a preference order (possibly untruthful).
- There are no payments.
- The outcome is a single alternative and this determines the allocation rule (each player receives some "payoff" from that alternative winning).

Impossibility result

- Gibbard and Satterthwaite proved around 1973 that truthful social choice mechanisms are essentially impossible. Long suspected, and widely considered to be devastating.

Impossibility result

- Gibbard and Satterthwaite proved around 1973 that truthful social choice mechanisms are essentially impossible. Long suspected, and widely considered to be devastating.
- Formally, if f is a social choice function, $m \geq 3$, $n \geq 2$, and each alternative can win for some preference profile, then f is a dictatorship or it is sometimes desirable to vote untruthfully.

Impossibility result

- Gibbard and Satterthwaite proved around 1973 that truthful social choice mechanisms are essentially impossible. Long suspected, and widely considered to be devastating.
- Formally, if f is a social choice function, $m \geq 3$, $n \geq 2$, and each alternative can win for some preference profile, then f is a dictatorship or it is sometimes desirable to vote untruthfully.
- The main problem is that in this model we have no way of measuring utility, or of comparing utilities between players. Money is a convenient way of getting past this problem, which is why interesting truthful mechanisms can exist in commercial settings.

Impossibility result

- Gibbard and Satterthwaite proved around 1973 that truthful social choice mechanisms are essentially impossible. Long suspected, and widely considered to be devastating.
- Formally, if f is a social choice function, $m \geq 3$, $n \geq 2$, and each alternative can win for some preference profile, then f is a dictatorship or it is sometimes desirable to vote untruthfully.
- The main problem is that in this model we have no way of measuring utility, or of comparing utilities between players. Money is a convenient way of getting past this problem, which is why interesting truthful mechanisms can exist in commercial settings.
- Manipulation by coalitions is sometimes possible where individual manipulation is not.

Coalitional manipulation example

- Consider a voting situation with 3 alternatives a, b, c and having 4 abc, 3 bca and 2 cab voters. Under the plurality rule, he sincere winner is a.

Coalitional manipulation example

- Consider a voting situation with 3 alternatives a, b, c and having 4 abc, 3 bca and 2 cab voters. Under the plurality rule, he sincere winner is a.
- No coalition can manipulate so that b wins.

Coalitional manipulation example

- Consider a voting situation with 3 alternatives a, b, c and having 4 abc, 3 bca and 2 cab voters. Under the plurality rule, he sincere winner is a.
- No coalition can manipulate so that b wins.
- However, if the bca voters all vote strategically as cba, then c wins.

Coalitional manipulation example

- Consider a voting situation with 3 alternatives a, b, c and having 4 abc, 3 bca and 2 cab voters. Under the plurality rule, he sincere winner is a.
- No coalition can manipulate so that b wins.
- However, if the bca voters all vote strategically as cba, then c wins.
- This is an example of a mechanism that is individually truthful, but not jointly - a group has an incentive to deviate. Voting sincerely is a Nash equilibrium, but not a strong Nash equilibrium.

Computational response to Gibbard-Satterthwaite

- If it is NP-hard to compute a manipulating strategy, perhaps voters will be truthful in practice, even if in theory it is in their interest to deviate.

Computational response to Gibbard-Satterthwaite

- If it is NP-hard to compute a manipulating strategy, perhaps voters will be truthful in practice, even if in theory it is in their interest to deviate.
- Successes: Instant Runoff Voting is NP-hard to manipulate by a single voter [BO1991]; weighted voting rules are almost always NP-hard to manipulate by a coalition, even for a fixed number of alternatives [CSL2007].

Computational response to Gibbard-Satterthwaite

- If it is NP-hard to compute a manipulating strategy, perhaps voters will be truthful in practice, even if in theory it is in their interest to deviate.
- Successes: Instant Runoff Voting is NP-hard to manipulate by a single voter [BO1991]; weighted voting rules are almost always NP-hard to manipulate by a coalition, even for a fixed number of alternatives [CSL2007].
- Problems: NP-hardness is only a worst-case guarantee. Most rules seem easy to manipulate in practice (based on simulation and some analytic results, e.g. [RPW2010]).

Some topics of current interest to me

- Complexity of safe manipulation of voting rules (Egor Ianovski)

Some topics of current interest to me

- Complexity of safe manipulation of voting rules (Egor Ianovski)
- Best-reply dynamics in voting games (Reyhaneh Reyhani)

Some topics of current interest to me

- Complexity of safe manipulation of voting rules (Egor Ianovski)
- Best-reply dynamics in voting games (Reyhaneh Reyhani)
- Convergence to equilibria via polling with incomplete information (Reyhaneh Reyhani)

Some topics of current interest to me

- Complexity of safe manipulation of voting rules (Egor Ianovski)
- Best-reply dynamics in voting games (Reyhaneh Reyhani)
- Convergence to equilibria via polling with incomplete information (Reyhaneh Reyhani)
- Asymptotic probabilistic measures of manipulability (Geoffrey Pritchard)

Some topics of current interest to me

- Complexity of safe manipulation of voting rules (Egor Ianovski)
- Best-reply dynamics in voting games (Reyhaneh Reyhani)
- Convergence to equilibria via polling with incomplete information (Reyhaneh Reyhani)
- Asymptotic probabilistic measures of manipulability (Geoffrey Pritchard)
- Implementation of social choice rules using different solution concepts.

Dynamics in voting games

- Suppose agents can communicate only with a central node.

Dynamics in voting games

- Suppose agents can communicate only with a central node.
- Agents vote sequentially using the plurality rule. After each vote all agents know the current state of the election.

Dynamics in voting games

- Suppose agents can communicate only with a central node.
- Agents vote sequentially using the plurality rule. After each vote all agents know the current state of the election.
- Each tries to obtain its best possible result assuming that its vote will be the last.

Dynamics in voting games

- Suppose agents can communicate only with a central node.
- Agents vote sequentially using the plurality rule. After each vote all agents know the current state of the election.
- Each tries to obtain its best possible result assuming that its vote will be the last.
- A (pure strategy) Nash equilibrium is always reached in $O(m^2n^2)$ iterations.

Dynamics in voting games

- Suppose agents can communicate only with a central node.
- Agents vote sequentially using the plurality rule. After each vote all agents know the current state of the election.
- Each tries to obtain its best possible result assuming that its vote will be the last.
- A (pure strategy) Nash equilibrium is always reached in $O(m^2n^2)$ iterations.
- Small changes to hypotheses lead to a failure to coordinate.

Dynamics in voting games

- Suppose agents can communicate only with a central node.
- Agents vote sequentially using the plurality rule. After each vote all agents know the current state of the election.
- Each tries to obtain its best possible result assuming that its vote will be the last.
- A (pure strategy) Nash equilibrium is always reached in $O(m^2n^2)$ iterations.
- Small changes to hypotheses lead to a failure to coordinate.
- Above results are [MPJR2010]. What happens for other voting rules?

Convergence via polling

- Consider the previous model, but each agent has inertia, a new measure of its risk attitude and available information. Also, instead of sequentially, agents vote simultaneously, and they repeat this procedure. Can interpret as a sequence of opinion polls, and agents strategize based on the incomplete information gleaned from polls.

Convergence via polling

- Consider the previous model, but each agent has inertia, a new measure of its risk attitude and available information. Also, instead of sequentially, agents vote simultaneously, and they repeat this procedure. Can interpret as a sequence of opinion polls, and agents strategize based on the incomplete information gleaned from polls.
- For some inertia distributions, convergence to an equilibrium where only two candidates get votes (Duverger's law). For others, no convergence.

Convergence via polling

- Consider the previous model, but each agent has inertia, a new measure of its risk attitude and available information. Also, instead of sequentially, agents vote simultaneously, and they repeat this procedure. Can interpret as a sequence of opinion polls, and agents strategize based on the incomplete information gleaned from polls.
- For some inertia distributions, convergence to an equilibrium where only two candidates get votes (Duverger's law). For others, no convergence.
- In the zero inertia case, announcing Plurality leads to Instant Runoff.

Convergence via polling

- Consider the previous model, but each agent has inertia, a new measure of its risk attitude and available information. Also, instead of sequentially, agents vote simultaneously, and they repeat this procedure. Can interpret as a sequence of opinion polls, and agents strategize based on the incomplete information gleaned from polls.
- For some inertia distributions, convergence to an equilibrium where only two candidates get votes (Duverger's law). For others, no convergence.
- In the zero inertia case, announcing Plurality leads to Instant Runoff.
- Idea of Reyhaneh Reyhani (PhD student), explored in her thesis work.

Safe manipulation

- Manipulation by a coalition raises hard questions: how do they coordinate?

Safe manipulation

- Manipulation by a coalition raises hard questions: how do they coordinate?
- Slinko and White (2008) introduced safe manipulation. A voter (interpreted as a party leader) issues a call to members to cast a named strategic vote but has no control over how many will follow and how many will remain sincere.

Safe manipulation

- Manipulation by a coalition raises hard questions: how do they coordinate?
- Slinko and White (2008) introduced safe manipulation. A voter (interpreted as a party leader) issues a call to members to cast a named strategic vote but has no control over how many will follow and how many will remain sincere.
- The manipulation is safe if no matter how many follow the call, a worse result is never obtained, and for some number of followers, a better result occurs. A strong incentive to manipulate!

Safe manipulation

- Manipulation by a coalition raises hard questions: how do they coordinate?
- Slinko and White (2008) introduced safe manipulation. A voter (interpreted as a party leader) issues a call to members to cast a named strategic vote but has no control over how many will follow and how many will remain sincere.
- The manipulation is safe if no matter how many follow the call, a worse result is never obtained, and for some number of followers, a better result occurs. A strong incentive to manipulate!
- S & W proved an analogue of Gibbard-Satterthwaite, so we can't avoid safe manipulation.

Safe manipulation

- Manipulation by a coalition raises hard questions: how do they coordinate?
- Slinko and White (2008) introduced safe manipulation. A voter (interpreted as a party leader) issues a call to members to cast a named strategic vote but has no control over how many will follow and how many will remain sincere.
- The manipulation is safe if no matter how many follow the call, a worse result is never obtained, and for some number of followers, a better result occurs. A strong incentive to manipulate!
- S & W proved an analogue of Gibbard-Satterthwaite, so we can't avoid safe manipulation.
- Can complexity help? Can a safe manipulation be found in polynomial time? Egor Ianovski (CS380 project) has solved this open problem for the Borda rule.