# Dynamics of voting games: preliminary report

Mark C. Wilson
www.cs.auckland.ac.nz/˜mcw/blog/
(joint with Reyhaneh Reyhani)

Department of Computer Science
University of Auckland

CMSS Summer Workshop, Auckland, 2010-12-20

# Basic setup

- We have a set $\mathcal{C}$ of alternatives (candidates) and set $\mathcal{V}$ of voters, with $m := |\mathcal{C}|, n := |\mathcal{V}|$.
- Each voter $v$ submits a permutation $L(v)$ of the candidates. This defines the set $\mathcal{T}$ of types, and $|\mathcal{T}| = m!$.
- A profile is a function $\mathcal{V} \to \mathcal{T}$. A voting situation is a multiset from $\mathcal{T}$ with total weight $n$.
- The scoring rule determined by a vector $w$ with $w_1 \geq w_2 \geq \cdots \geq w_{m-1} \geq w_m$ assigns the score

$$|c| := \sum_{t \in T} |\{v \in \mathcal{V} \mid L(v) = t\}| w_{L(v)^{-1}(c)}.$$

- Special cases:
  - plurality: $w = (1, 0, 0, \ldots, 0)$;
  - antiplurality (veto): $w = (1, 1, \ldots, 1, 0)$;
  - Borda: $w = (m-1, m-2, \ldots, 1, 0)$.

# Basic setup

- We have a set $\mathcal{C}$ of alternatives (candidates) and set $\mathcal{V}$ of voters, with $m := |\mathcal{C}|, n := |\mathcal{V}|$.
- Each voter $v$ submits a permutation $L(v)$ of the candidates. This defines the set $\mathcal{T}$ of types, and $|\mathcal{T}| = m!$.
- A profile is a function $\mathcal{V} \to \mathcal{T}$. A voting situation is a multiset from $\mathcal{T}$ with total weight $n$.
- The scoring rule determined by a vector $w$ with $w_1 \geq w_2 \geq \cdots \geq w_{m-1} \geq w_m$ assigns the score

$$|c| := \sum_{t \in T} |\{v \in \mathcal{V} \mid L(v) = t\}| w_{L(v)^{-1}(c)}.$$

- Special cases:
  - plurality: $w = (1, 0, 0, \ldots, 0)$;
  - antiplurality (veto): $w = (1, 1, \ldots, 1, 0)$;
  - Borda: $w = (m-1, m-2, \ldots, 1, 0)$.

# Basic setup

- We have a set $\mathcal{C}$ of alternatives (candidates) and set $\mathcal{V}$ of voters, with $m := |\mathcal{C}|, n := |\mathcal{V}|$.
- Each voter $v$ submits a permutation $L(v)$ of the candidates. This defines the set $\mathcal{T}$ of types, and $|\mathcal{T}| = m!$.
- A profile is a function $\mathcal{V} \to \mathcal{T}$. A voting situation is a multiset from $\mathcal{T}$ with total weight $n$.
- The scoring rule determined by a vector $w$ with $w_1 \geq w_2 \geq \cdots \geq w_{m-1} \geq w_m$ assigns the score

$$|c| := \sum_{t \in T} |\{v \in \mathcal{V} \mid L(v) = t\}| w_{L(v)^{-1}(c)}.$$

- Special cases:
  - plurality: $w = (1, 0, 0, \ldots, 0)$;
  - antiplurality (veto): $w = (1, 1, \ldots, 1, 0)$;
  - Borda: $w = (m-1, m-2, \ldots, 1, 0)$.

# Basic setup

- We have a set $\mathcal{C}$ of alternatives (candidates) and set $\mathcal{V}$ of voters, with $m := |\mathcal{C}|, n := |\mathcal{V}|$.
- Each voter $v$ submits a permutation $L(v)$ of the candidates. This defines the set $\mathcal{T}$ of types, and $|\mathcal{T}| = m!$.
- A profile is a function $\mathcal{V} \to \mathcal{T}$. A voting situation is a multiset from $\mathcal{T}$ with total weight $n$.
- The scoring rule determined by a vector $w$ with $w_1 \geq w_2 \geq \cdots \geq w_{m-1} \geq w_m$ assigns the score

$$|c| := \sum_{t \in T} |\{v \in \mathcal{V} \mid L(v) = t\}| w_{L(v)^{-1}(c)}.$$

- Special cases:
  - plurality: $w = (1, 0, 0, \dots, 0)$;
  - antiplurality (veto): $w = (1, 1, \dots, 1, 0)$;
  - Borda: $w = (m-1, m-2, \dots, 1, 0)$.

# Basic setup

- We have a set $\mathcal{C}$ of alternatives (candidates) and set $\mathcal{V}$ of voters, with $m := |\mathcal{C}|, n := |\mathcal{V}|$.
- Each voter $v$ submits a permutation $L(v)$ of the candidates. This defines the set $\mathcal{T}$ of types, and $|\mathcal{T}| = m!$.
- A profile is a function $\mathcal{V} \to \mathcal{T}$. A voting situation is a multiset from $\mathcal{T}$ with total weight $n$.
- The scoring rule determined by a vector $w$ with $w_1 \geq w_2 \geq \cdots \geq w_{m-1} \geq w_m$ assigns the score

$$|c| := \sum_{t \in T} |\{v \in \mathcal{V} \mid L(v) = t\}| w_{L(v)^{-1}(c)}.$$

- Special cases:
  - plurality: $w = (1, 0, 0, \ldots, 0)$;
  - antiplurality (veto): $w = (1, 1, \ldots, 1, 0)$;
  - Borda: $w = (m-1, m-2, \ldots, 1, 0)$.

# Basic setup

- We have a set $\mathcal{C}$ of alternatives (candidates) and set $\mathcal{V}$ of voters, with $m := |\mathcal{C}|, n := |\mathcal{V}|$.
- Each voter $v$ submits a permutation $L(v)$ of the candidates. This defines the set $\mathcal{T}$ of types, and $|\mathcal{T}| = m!$.
- A profile is a function $\mathcal{V} \to \mathcal{T}$. A voting situation is a multiset from $\mathcal{T}$ with total weight $n$.
- The scoring rule determined by a vector $w$ with $w_1 \geq w_2 \geq \cdots \geq w_{m-1} \geq w_m$ assigns the score

$$|c| := \sum_{t \in T} |\{v \in \mathcal{V} \mid L(v) = t\}| w_{L(v)^{-1}(c)}.$$

- Special cases:
  - plurality: $w = (1, 0, 0, \ldots, 0)$;
  - antiplurality (veto): $w = (1, 1, \ldots, 1, 0)$;
  - Borda: $w = (m-1, m-2, \ldots, 1, 0)$.

# Basic setup

- We have a set $\mathcal{C}$ of alternatives (candidates) and set $\mathcal{V}$ of voters, with $m := |\mathcal{C}|, n := |\mathcal{V}|$.
- Each voter $v$ submits a permutation $L(v)$ of the candidates. This defines the set $\mathcal{T}$ of types, and $|\mathcal{T}| = m!$.
- A profile is a function $\mathcal{V} \to \mathcal{T}$. A voting situation is a multiset from $\mathcal{T}$ with total weight $n$.
- The scoring rule determined by a vector $w$ with $w_1 \geq w_2 \geq \cdots \geq w_{m-1} \geq w_m$ assigns the score

$$|c| := \sum_{t \in T} |\{v \in \mathcal{V} \mid L(v) = t\}| w_{L(v)^{-1}(c)}.$$

- Special cases:
  - plurality: $w = (1, 0, 0, \ldots, 0)$;
  - antiplurality (veto): $w = (1, 1, \ldots, 1, 0)$;
  - Borda: $w = (m-1, m-2, \ldots, 1, 0)$.

# Basic setup

- We have a set $\mathcal{C}$ of alternatives (candidates) and set $\mathcal{V}$ of voters, with $m := |\mathcal{C}|, n := |\mathcal{V}|$.
- Each voter $v$ submits a permutation $L(v)$ of the candidates. This defines the set $\mathcal{T}$ of types, and $|\mathcal{T}| = m!$.
- A profile is a function $\mathcal{V} \to \mathcal{T}$. A voting situation is a multiset from $\mathcal{T}$ with total weight $n$.
- The scoring rule determined by a vector $w$ with $w_1 \geq w_2 \geq \cdots \geq w_{m-1} \geq w_m$ assigns the score

$$|c| := \sum_{t \in T} |\{v \in \mathcal{V} \mid L(v) = t\}| w_{L(v)^{-1}(c)}.$$

- Special cases:
    - plurality: $w = (1, 0, 0, \ldots, 0)$;
    - antiplurality (veto): $w = (1, 1, \ldots, 1, 0)$;
    - Borda: $w = (m-1, m-2, \ldots, 1, 0)$.

# Voting games

- Player (voter) action is to submit an expressed vote (possibly different from its sincere preference).
- Gibbard-Satterthwaite and other theorems show that dominant strategies don't always exist.
- Far too many Nash equilibria exist for this to be a useful concept, so refinements are probably needed.
- Meir, Polukarov, Jennings, Rosenschein (AAAI 2010) studied best-reply dynamics (BRD).
- If BRD converges, it necessarily finds a Nash equilibrium.

# Voting games

- Player (voter) action is to submit an expressed vote (possibly different from its sincere preference).
- Gibbard-Satterthwaite and other theorems show that dominant strategies don't always exist.
- Far too many Nash equilibria exist for this to be a useful concept, so refinements are probably needed.
- Meir, Polukarov, Jennings, Rosenschein (AAAI 2010) studied best-reply dynamics (BRD).
- If BRD converges, it necessarily finds a Nash equilibrium.

# Voting games

- Player (voter) action is to submit an expressed vote (possibly different from its sincere preference).
- Gibbard-Satterthwaite and other theorems show that dominant strategies don't always exist.
- Far too many Nash equilibria exist for this to be a useful concept, so refinements are probably needed.
- Meir, Polukarov, Jennings, Rosenschein (AAAI 2010) studied best-reply dynamics (BRD).
- If BRD converges, it necessarily finds a Nash equilibrium.

# Voting games

- Player (voter) action is to submit an expressed vote (possibly different from its sincere preference).
- Gibbard-Satterthwaite and other theorems show that dominant strategies don't always exist.
- Far too many Nash equilibria exist for this to be a useful concept, so refinements are probably needed.
- Meir, Polukarov, Jennings, Rosenschein (AAAI 2010) studied best-reply dynamics (BRD).
- If BRD converges, it necessarily finds a Nash equilibrium.

# Voting games

- Player (voter) action is to submit an expressed vote (possibly different from its sincere preference).
- Gibbard-Satterthwaite and other theorems show that dominant strategies don't always exist.
- Far too many Nash equilibria exist for this to be a useful concept, so refinements are probably needed.
- Meir, Polukarov, Jennings, Rosenschein (AAAI 2010) studied best-reply dynamics (BRD).
- If BRD converges, it necessarily finds a Nash equilibrium.

# Summary of our recent activity in this area

- Repeated polling (Reyhaneh Reyhani: well developed, relations to STV, Duverger's law - not today).
- Best reply dynamics of plurality games (Reyhaneh Reyhani: preliminary, today).
- Sequential/strategic entry (Ben Skudder: very preliminary, not today).

The University of Auckland
NEW ZEALAND

# Summary of our recent activity in this area

- Repeated polling (Reyhaneh Reyhani: well developed, relations to STV, Duverger's law - not today).
- Best reply dynamics of plurality games (Reyhaneh Reyhani: preliminary, today).
- Sequential/strategic entry (Ben Skudder: very preliminary, not today).

# Summary of our recent activity in this area

- Repeated polling (Reyhaneh Reyhani: well developed, relations to STV, Duverger's law - not today).
- Best reply dynamics of plurality games (Reyhaneh Reyhani: preliminary, today).
- Sequential/strategic entry (Ben Skudder: very preliminary, not today).

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

The University of Auckland
NEW ZEALAND

# Summary of Meir et al

- Assumptions:
  - Fixed assumptions: myopic, no communication between players and zero knowledge of others.
  - Other assumptions:
    - Behaviour : best reply at each step or arbitrary improvement step.
    - Indifference: keep last move or report sincere preference.
    - Initial state: sincere profile or arbitrary profile.
    - Tiebreaking: deterministic or uniform random.
    - Voters: unweighted or weighted.
- Results:
  - Convergence for plurality under red hypotheses in at most $m^2n^2$ steps. Also, deterministic tiebreaking from an arbitrary initial state converges for unweighted voters. The winner is the sincere winner or a candidate at most 1 point behind initially.

In each case, changing each red hypothesis and keeping the others yields examples of non-convergence.

# Our aims

- Consider best-reply dynamics for more general voting rules (scoring rules). Restrict to sincere initial state, unweighted voters, keep last move if no improvement possible.
- Conjecture and prove general positive results, where possible. Otherwise clearly explain the negative results.
- Use best-reply dynamics to refine equilibria (better predictive value) and measure manipulability.
- Have done preliminary simulation results for several rules including Borda, $2$-approval, antiplurality.

# A basic issue

- Best reply is not unique, because several preference orders may yield the same result.
- Traditional game-theoretic idea is to randomize, and use mixed strategies.
- However we often avoid random tiebreaking because we want a deterministic voting rule, so we may wish to restrict to pure strategies.
- One option is to use a fixed enumeration of the preference orders, and choose the one with smallest index among all actions giving the best result.
- Example: $abc, acb, bac, bca, cab, cba$, standard lexicographic order.

# A basic issue

- Best reply is not unique, because several preference orders may yield the same result.
- Traditional game-theoretic idea is to randomize, and use mixed strategies.
- However we often avoid random tiebreaking because we want a deterministic voting rule, so we may wish to restrict to pure strategies.
- One option is to use a fixed enumeration of the preference orders, and choose the one with smallest index among all actions giving the best result.
- Example: $abc, acb, bac, bca, cab, cba$, standard lexicographic order.

# A basic issue

- Best reply is not unique, because several preference orders may yield the same result.
- Traditional game-theoretic idea is to randomize, and use mixed strategies.
- However we often avoid random tiebreaking because we want a deterministic voting rule, so we may wish to restrict to pure strategies.
- One option is to use a fixed enumeration of the preference orders, and choose the one with smallest index among all actions giving the best result.
- Example: $abc, acb, bac, bca, cab, cba$, standard lexicographic order.

# A basic issue

- Best reply is not unique, because several preference orders may yield the same result.
- Traditional game-theoretic idea is to randomize, and use mixed strategies.
- However we often avoid random tiebreaking because we want a deterministic voting rule, so we may wish to restrict to pure strategies.
- One option is to use a fixed enumeration of the preference orders, and choose the one with smallest index among all actions giving the best result.
- Example: $abc, acb, bac, bca, cab, cba$, standard lexicographic order.

The University of Auckland
NEW ZEALAND

# A basic issue

- Best reply is not unique, because several preference orders may yield the same result.
- Traditional game-theoretic idea is to randomize, and use mixed strategies.
- However we often avoid random tiebreaking because we want a deterministic voting rule, so we may wish to restrict to pure strategies.
- One option is to use a fixed enumeration of the preference orders, and choose the one with smallest index among all actions giving the best result.
- Example: $abc, acb, bac, bca, cab, cba$, standard lexicographic order.

# Cycles

## Example

Consider P=$(abc, bca)$ and voting rule Borda, so sincere scoreboard is (2, 3, 1) and sincere winner $b$. Best reply of voter $1$ is $acb$, giving $3$-way tie. Under alphabetical tie breaking, the winner is $a$. Player $2$ changes to $cba$ and the winner switches from $a$ to $c$.
We list the current state $P$ after each number of stages (stage $i$ is the list of $i$th moves of all players).

| | | |
|---|---|---|
| 0 | $P = (abc, bac)$ | $S(P) = (2, 3, 1)$ |
| 1 | $P = (acb, cba)$ | $S(P) = (2, 1, 3)$ |
| 2 | $P = (abc, bac)$ | $S(P) = (2, 3, 1)$ |

# Order of players matters

## Example

Consider $P = (acb, acb, cab, cba)$ under Borda.

0 $P = (acb, acb, cab, cba)$ $S(P) = (5, 1, 6)$

1 $P = (abc, acb, cba, cba)$ $S(P) = (4, 3, 5)$

2 $P = (abc, abc, cba, bac)$ $S(P) = (4, 5, 3)$

3 $P = (acb, abc, cba, cba)$ $S(P) = (4, 3, 5)$

4 $P = (abc, abc, cba, bac)$ $S(P) = (4, 5, 3)$

A cycle of length $2$ has been reached.

Consider another profile for the same voting situation, $P' = (acb, acb, cab, cba)$. We obtain:

0 $P = (acb, acb, cba, cab)$ $S(P) = (5, 1, 6)$

1 $P = (abc, acb, cba, cba)$ $S(P) = (4, 3, 5)$

2 $P = (abc, abc, bca, cab)$ $S(P) = (5, 4, 3)$

3 $P = (abc, abc, bca, cab)$ $S(P) = (5, 4, 3)$

In this case, convergence has occurred.

# Best reply is not unique, and choices matter

## Example

Consider the scoring rule $(3, 2, 0)$ and sincere profile $P = (acb, acb, bca)$, using mixed strategies.

| | | |
|---|---|---|
| 0 | $P = (acb, acb, bca)$ | $S(P) = (6, 3, 6)$ |
| 1 | $P = (acb, acb, cba)$ | $S(P) = (6, 2, 7)$ |
| 2 | $P = (bac, acb, bac)$ | $S(P) = (5, 6, 4)$ |
| 3 | $P = (acb, acb, cba)$ | $S(P) = (6, 2, 7)$ |
| 4 | $P = (bac, acb, bac)$ | $S(P) = (5, 6, 4)$ |
| 5 | $P = (acb, acb, cba)$ | $S(P) = (6, 2, 7)$ |
| 6 | $P = (abc, acb, cba)$ | $S(P) = (6, 4, 5)$ |
| 7 | $P = (abc, acb, cba)$ | $S(P) = (6, 4, 5)$ |

However, we reach the same equilibrium after 3 iterations using our pure strategy algorithm, omitting stages 2 to 5.

# Pure vs mixed strategies

- It may take longer to converge when using mixed strategies (infinitely often, as previous example shows). This is intuitively clear.
- Conversely, mixed strategies sometimes allow quicker convergence than pure ones. This seems less obvious.
- Example: $m = n = 4$, Borda, initial state $(abcd, abcd, bacd, bdca)$ does not converge using our pure strategy setup, but does using mixed strategies.

The University of Auckland
NEW ZEALAND

# Pure vs mixed strategies

- It may take longer to converge when using mixed strategies (infinitely often, as previous example shows). This is intuitively clear.
- Conversely, mixed strategies sometimes allow quicker convergence than pure ones. This seems less obvious.
- Example: $m = n = 4$, Borda, initial state $(abcd, abcd, bacd, bdca)$ does not converge using our pure strategy setup, but does using mixed strategies.

The University of Auckland
NEW ZEALAND

# Pure vs mixed strategies

- It may take longer to converge when using mixed strategies (infinitely often, as previous example shows). This is intuitively clear.
- Conversely, mixed strategies sometimes allow quicker convergence than pure ones. This seems less obvious.
- Example: $m = n = 4$, Borda, initial state $(abcd, abcd, bacd, bdca)$ does not converge using our pure strategy setup, but does using mixed strategies.

# Other topics

- It seems more realistic in some interpretations to start from a zero initial state, rather than assuming all voters first vote once. This doesn't seem to affect results much but needs more exploration, because the first player has more influence.
- The set of possible winners is larger for Borda - a candidate can be more than 1 point behind initially and still win.
- If only $k$ voters play this game and the other $n - k$ always vote sincerely, similar results appear.

# Other topics

- It seems more realistic in some interpretations to start from a zero initial state, rather than assuming all voters first vote once. This doesn't seem to affect results much but needs more exploration, because the first player has more influence.
- The set of possible winners is larger for Borda - a candidate can be more than 1 point behind initially and still win.
- If only $k$ voters play this game and the other $n-k$ always vote sincerely, similar results appear.

# Other topics

- It seems more realistic in some interpretations to start from a zero initial state, rather than assuming all voters first vote once. This doesn't seem to affect results much but needs more exploration, because the first player has more influence.
- The set of possible winners is larger for Borda - a candidate can be more than 1 point behind initially and still win.
- If only $k$ voters play this game and the other $n - k$ always vote sincerely, similar results appear.

# Conjectures and facts

- BRD always converges for antiplurality. The upper bound on number of iterations is more than for plurality.
- BRD always converges for $m = n = 3$, for some class of rules including Borda, but not all rules. For other values of $m$ and $n$, the class does not include Borda.
- BRD converges in $(m-1)n$ steps from the sincere initial state, for plurality.
- Given an initial voting situation, for all representing profiles for which BRD converges, the equilibrium strategies are the same (up to permutation of voters), hence the winner is the same. However, the rate of convergence depends on the profile.

# Conjectures and facts

- BRD always converges for antiplurality. The upper bound on number of iterations is more than for plurality.
- BRD always converges for $m = n = 3$, for some class of rules including Borda, but not all rules. For other values of $m$ and $n$, the class does not include Borda.
- BRD converges in $(m-1)n$ steps from the sincere initial state, for plurality.
- Given an initial voting situation, for all representing profiles for which BRD converges, the equilibrium strategies are the same (up to permutation of voters), hence the winner is the same. However, the rate of convergence depends on the profile.

# Conjectures and facts

- BRD always converges for antiplurality. The upper bound on number of iterations is more than for plurality.
- BRD always converges for $m = n = 3$, for some class of rules including Borda, but not all rules. For other values of $m$ and $n$, the class does not include Borda.
- BRD converges in $(m-1)n$ steps from the sincere initial state, for plurality.
- Given an initial voting situation, for all representing profiles for which BRD converges, the equilibrium strategies are the same (up to permutation of voters), hence the winner is the same. However, the rate of convergence depends on the profile.

# Conjectures and facts

- BRD always converges for antiplurality. The upper bound on number of iterations is more than for plurality.
- BRD always converges for $m = n = 3$, for some class of rules including Borda, but not all rules. For other values of $m$ and $n$, the class does not include Borda.
- BRD converges in $(m-1)n$ steps from the sincere initial state, for plurality.
- Given an initial voting situation, for all representing profiles for which BRD converges, the equilibrium strategies are the same (up to permutation of voters), hence the winner is the same. However, the rate of convergence depends on the profile.

# Speculations

- How does convergence of BRD relate to difficulty of manipulation of a voting rule?
- For example, suppose the sincere initial state is an equilibrium. Then convergence is immediate, and by usual definitions the rule is not (individually) manipulable at that profile.
- On the other hand, if convergence is slow or cycling occurs, it seems reasonable to assume that manipulation is computationally hard.
- The concepts of counterthreat and reaction of Pattanaik may be given a clearer motivation via this model.
- Each Nash equilibrium occurs as a limit of best-reply dynamics, but some may be more stable than others, and have larger basins of attraction. Does this lead to a probability distribution on the equilibria, which may be useful for prediction or welfare comparisons?

# Speculations

- How does convergence of BRD relate to difficulty of manipulation of a voting rule?
- For example, suppose the sincere initial state is an equilibrium. Then convergence is immediate, and by usual definitions the rule is not (individually) manipulable at that profile.
- On the other hand, if convergence is slow or cycling occurs, it seems reasonable to assume that manipulation is computationally hard.
- The concepts of counterthreat and reaction of Pattanaik may be given a clearer motivation via this model.
- Each Nash equilibrium occurs as a limit of best-reply dynamics, but some may be more stable than others, and have larger basins of attraction. Does this lead to a probability distribution on the equilibria, which may be useful for prediction or welfare comparisons?

# Speculations

- How does convergence of BRD relate to difficulty of manipulation of a voting rule?
- For example, suppose the sincere initial state is an equilibrium. Then convergence is immediate, and by usual definitions the rule is not (individually) manipulable at that profile.
- On the other hand, if convergence is slow or cycling occurs, it seems reasonable to assume that manipulation is computationally hard.
- The concepts of counterthreat and reaction of Pattanaik may be given a clearer motivation via this model.
- Each Nash equilibrium occurs as a limit of best-reply dynamics, but some may be more stable than others, and have larger basins of attraction. Does this lead to a probability distribution on the equilibria, which may be useful for prediction or welfare comparisons?

# Speculations

- How does convergence of BRD relate to difficulty of manipulation of a voting rule?
- For example, suppose the sincere initial state is an equilibrium. Then convergence is immediate, and by usual definitions the rule is not (individually) manipulable at that profile.
- On the other hand, if convergence is slow or cycling occurs, it seems reasonable to assume that manipulation is computationally hard.
- The concepts of counterthreat and reaction of Pattanaik may be given a clearer motivation via this model.
- Each Nash equilibrium occurs as a limit of best-reply dynamics, but some may be more stable than others, and have larger basins of attraction. Does this lead to a probability distribution on the equilibria, which may be useful for prediction or welfare comparisons?

# Speculations

- How does convergence of BRD relate to difficulty of manipulation of a voting rule?
- For example, suppose the sincere initial state is an equilibrium. Then convergence is immediate, and by usual definitions the rule is not (individually) manipulable at that profile.
- On the other hand, if convergence is slow or cycling occurs, it seems reasonable to assume that manipulation is computationally hard.
- The concepts of counterthreat and reaction of Pattanaik may be given a clearer motivation via this model.
- Each Nash equilibrium occurs as a limit of best-reply dynamics, but some may be more stable than others, and have larger basins of attraction. Does this lead to a probability distribution on the equilibria, which may be useful for prediction or welfare comparisons?