A Natural Adaptive Process for Collective Decision-Making

Florian Brandl
(joint work with Felix Brandt)

Social Choice Theory: Standard Assumptions

- Fixed set of voters
- Fixed set of alternatives
- Voters are able to rank-order all alternatives
- Fixed preferences (single election)
- Central authority
  - collects reported preferences
  - computes election outcome
  - convinces voters of outcome's correctness

1	**1**	**1**
a	*a*	*b*
b	*c*	*c*
c	*b*	*a*

The Goal

- Our goal is to devise an ongoing adaptive voting process with desirable axiomatic properties.
  - Voters may arrive, leave, and change their preferences over time.
  - Voters are never asked for their complete preferences.
  - Get rid of the central authority.
- The only device we want to use is an urn filled with balls of various types that allows for two primitive operations:

  (i) randomly sampling a ball

  (ii) replacing a sampled ball of one type by a ball of another type

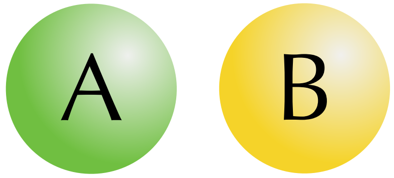
A
B

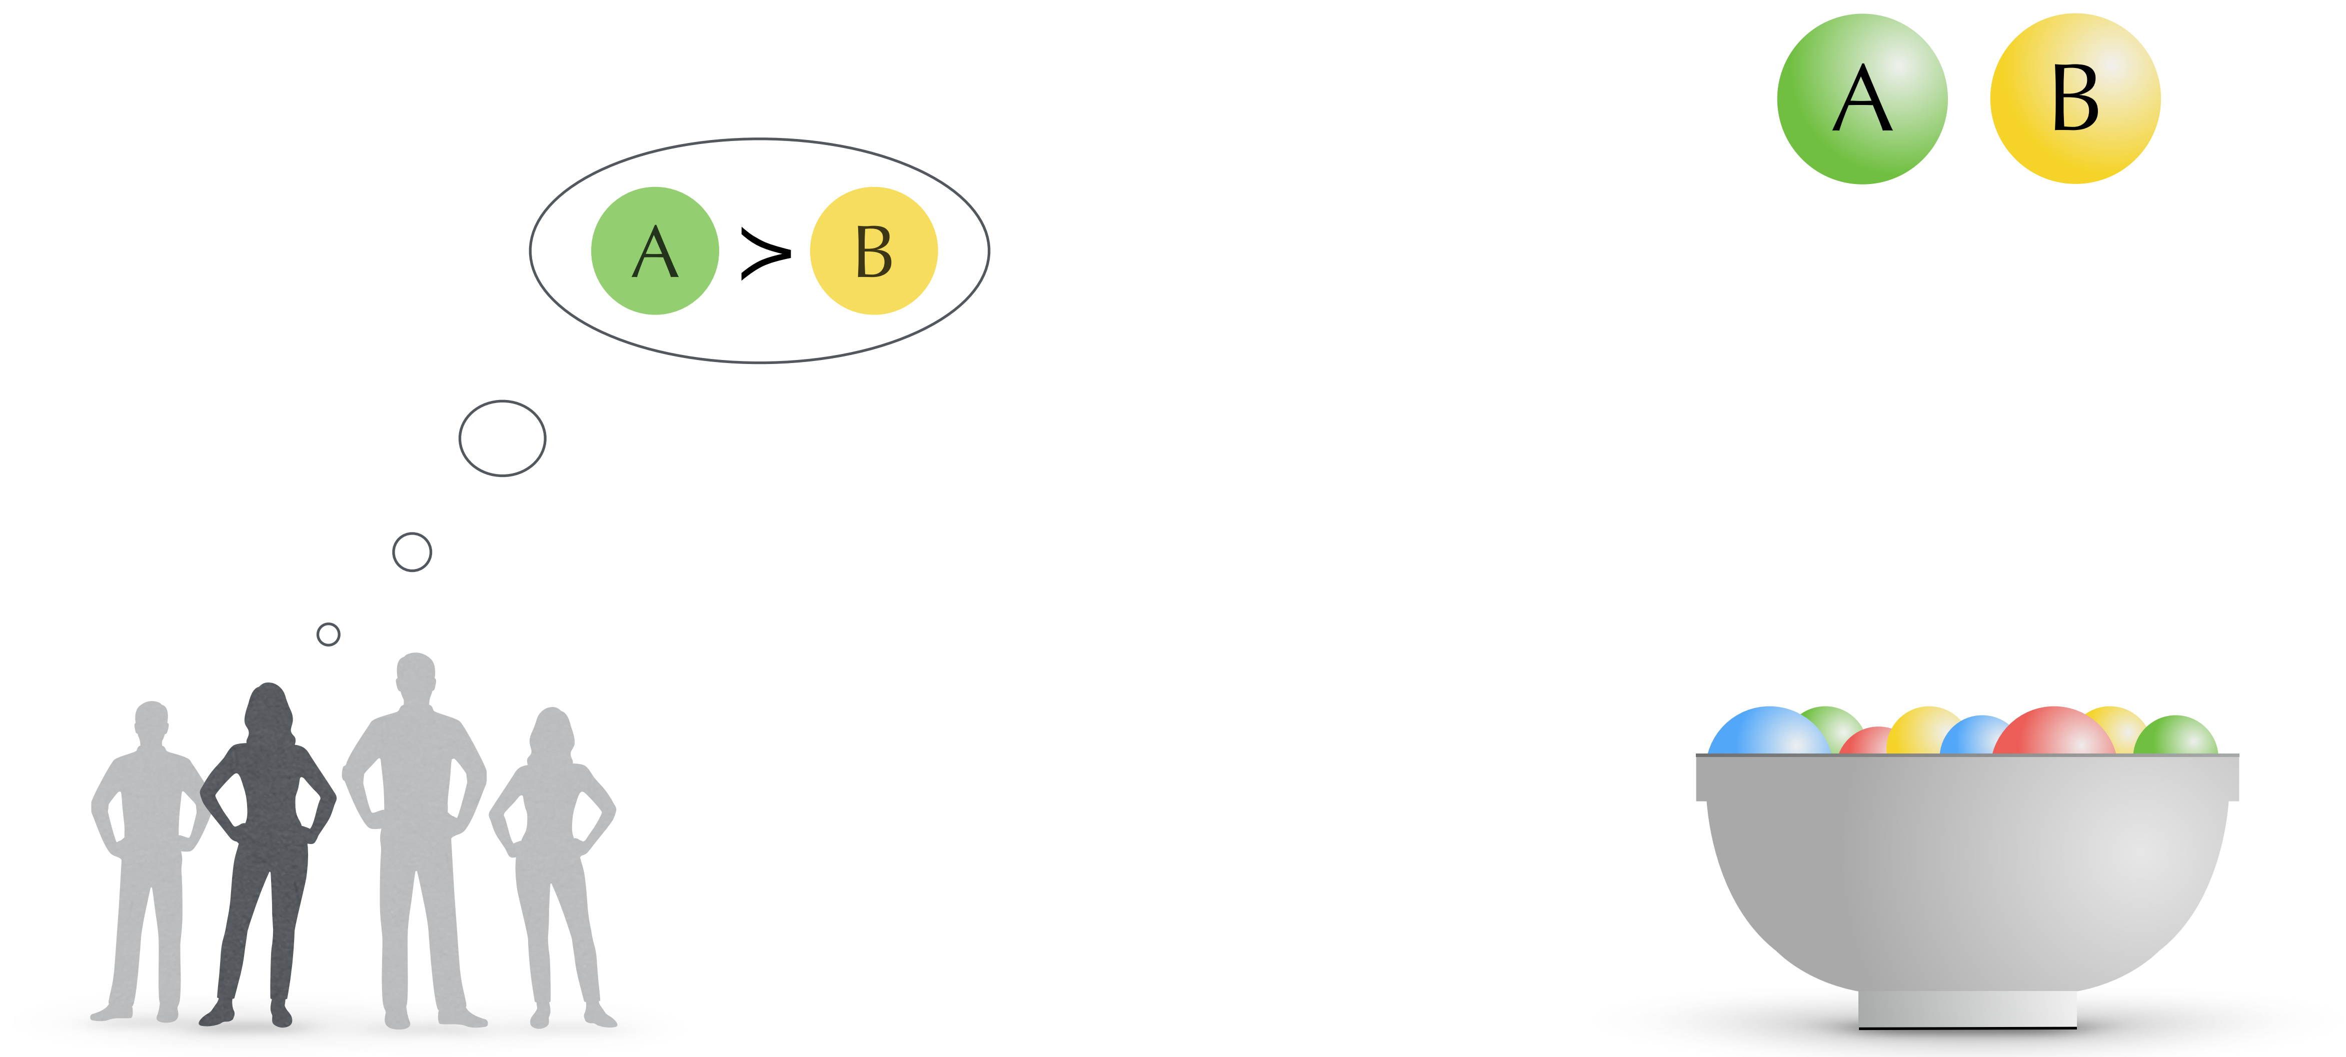
A > B
A
B

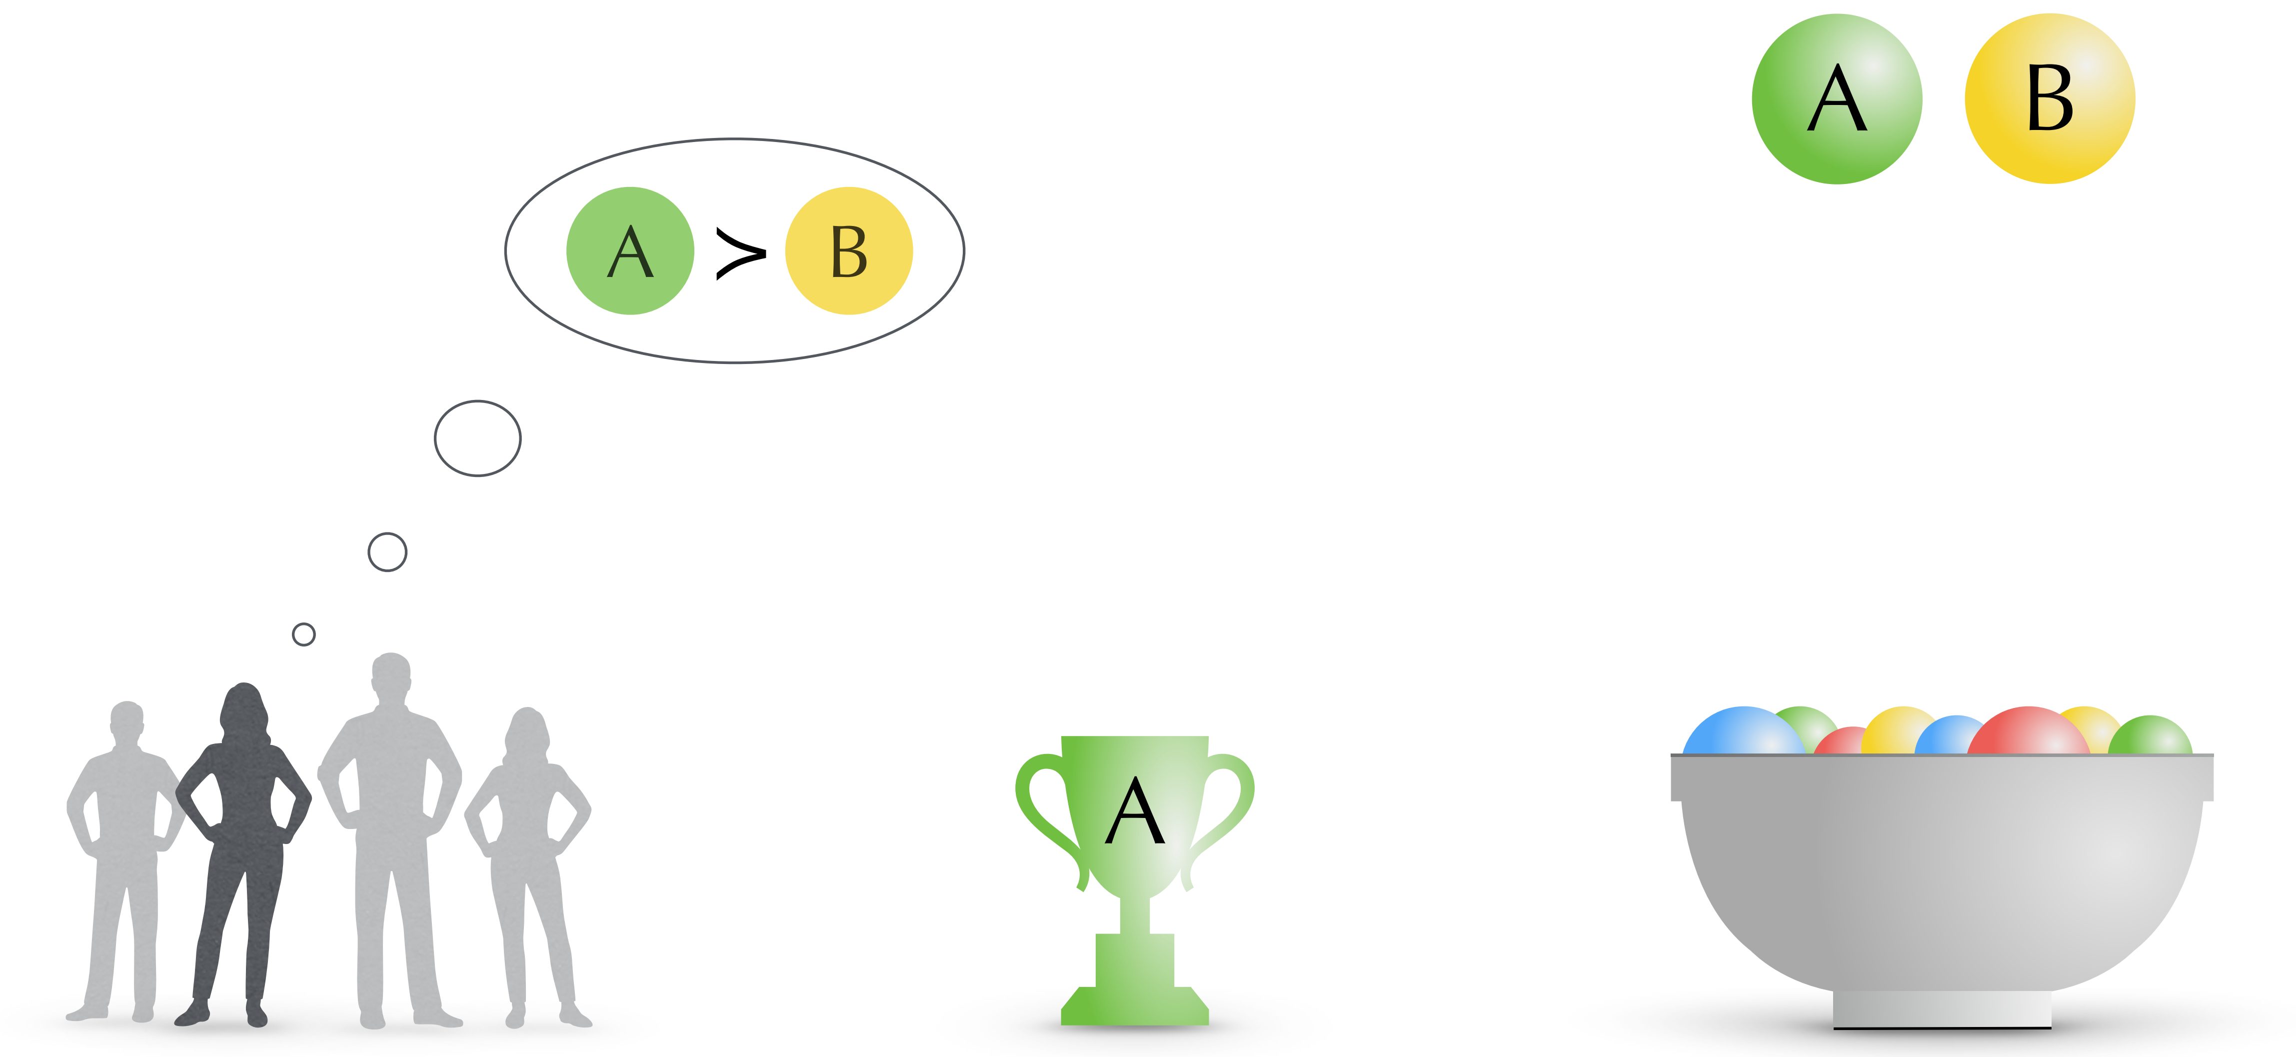
A > B
A
A
B


A


A


A
C
C
B
D
A
D
A
D


C
A
C
C
B
D
A
D
A
D


A C C B D A D A D C D D A D

- What can be said about the distribution of winners?
    - More generally, what about the distribution of balls in the urn?

Urn-Based Voting Process

- Urn filled with N balls, each carrying the label of an alternative.
  - Initial distribution of balls in urn is irrelevant.
- Repeat for each round:

1. A randomly selected voter i draws two balls from urn.
   - Assume the labels of these balls are x and y and $x \succ_i y$.
2. x is declared the winner of this round.
3. Voter i will change the label of the second ball to x and put both balls (now carrying the same label) back into the urn.
4. With some small probability r (called mutation rate), a randomly drawn ball is re-labelled with a random alternative.

100	**100**	**100**
a	*a*	*b*
b	*c*	*c*
c	*b*	*a*

- N=50 balls
- Mutation rate r=0.02
- 1000 rounds

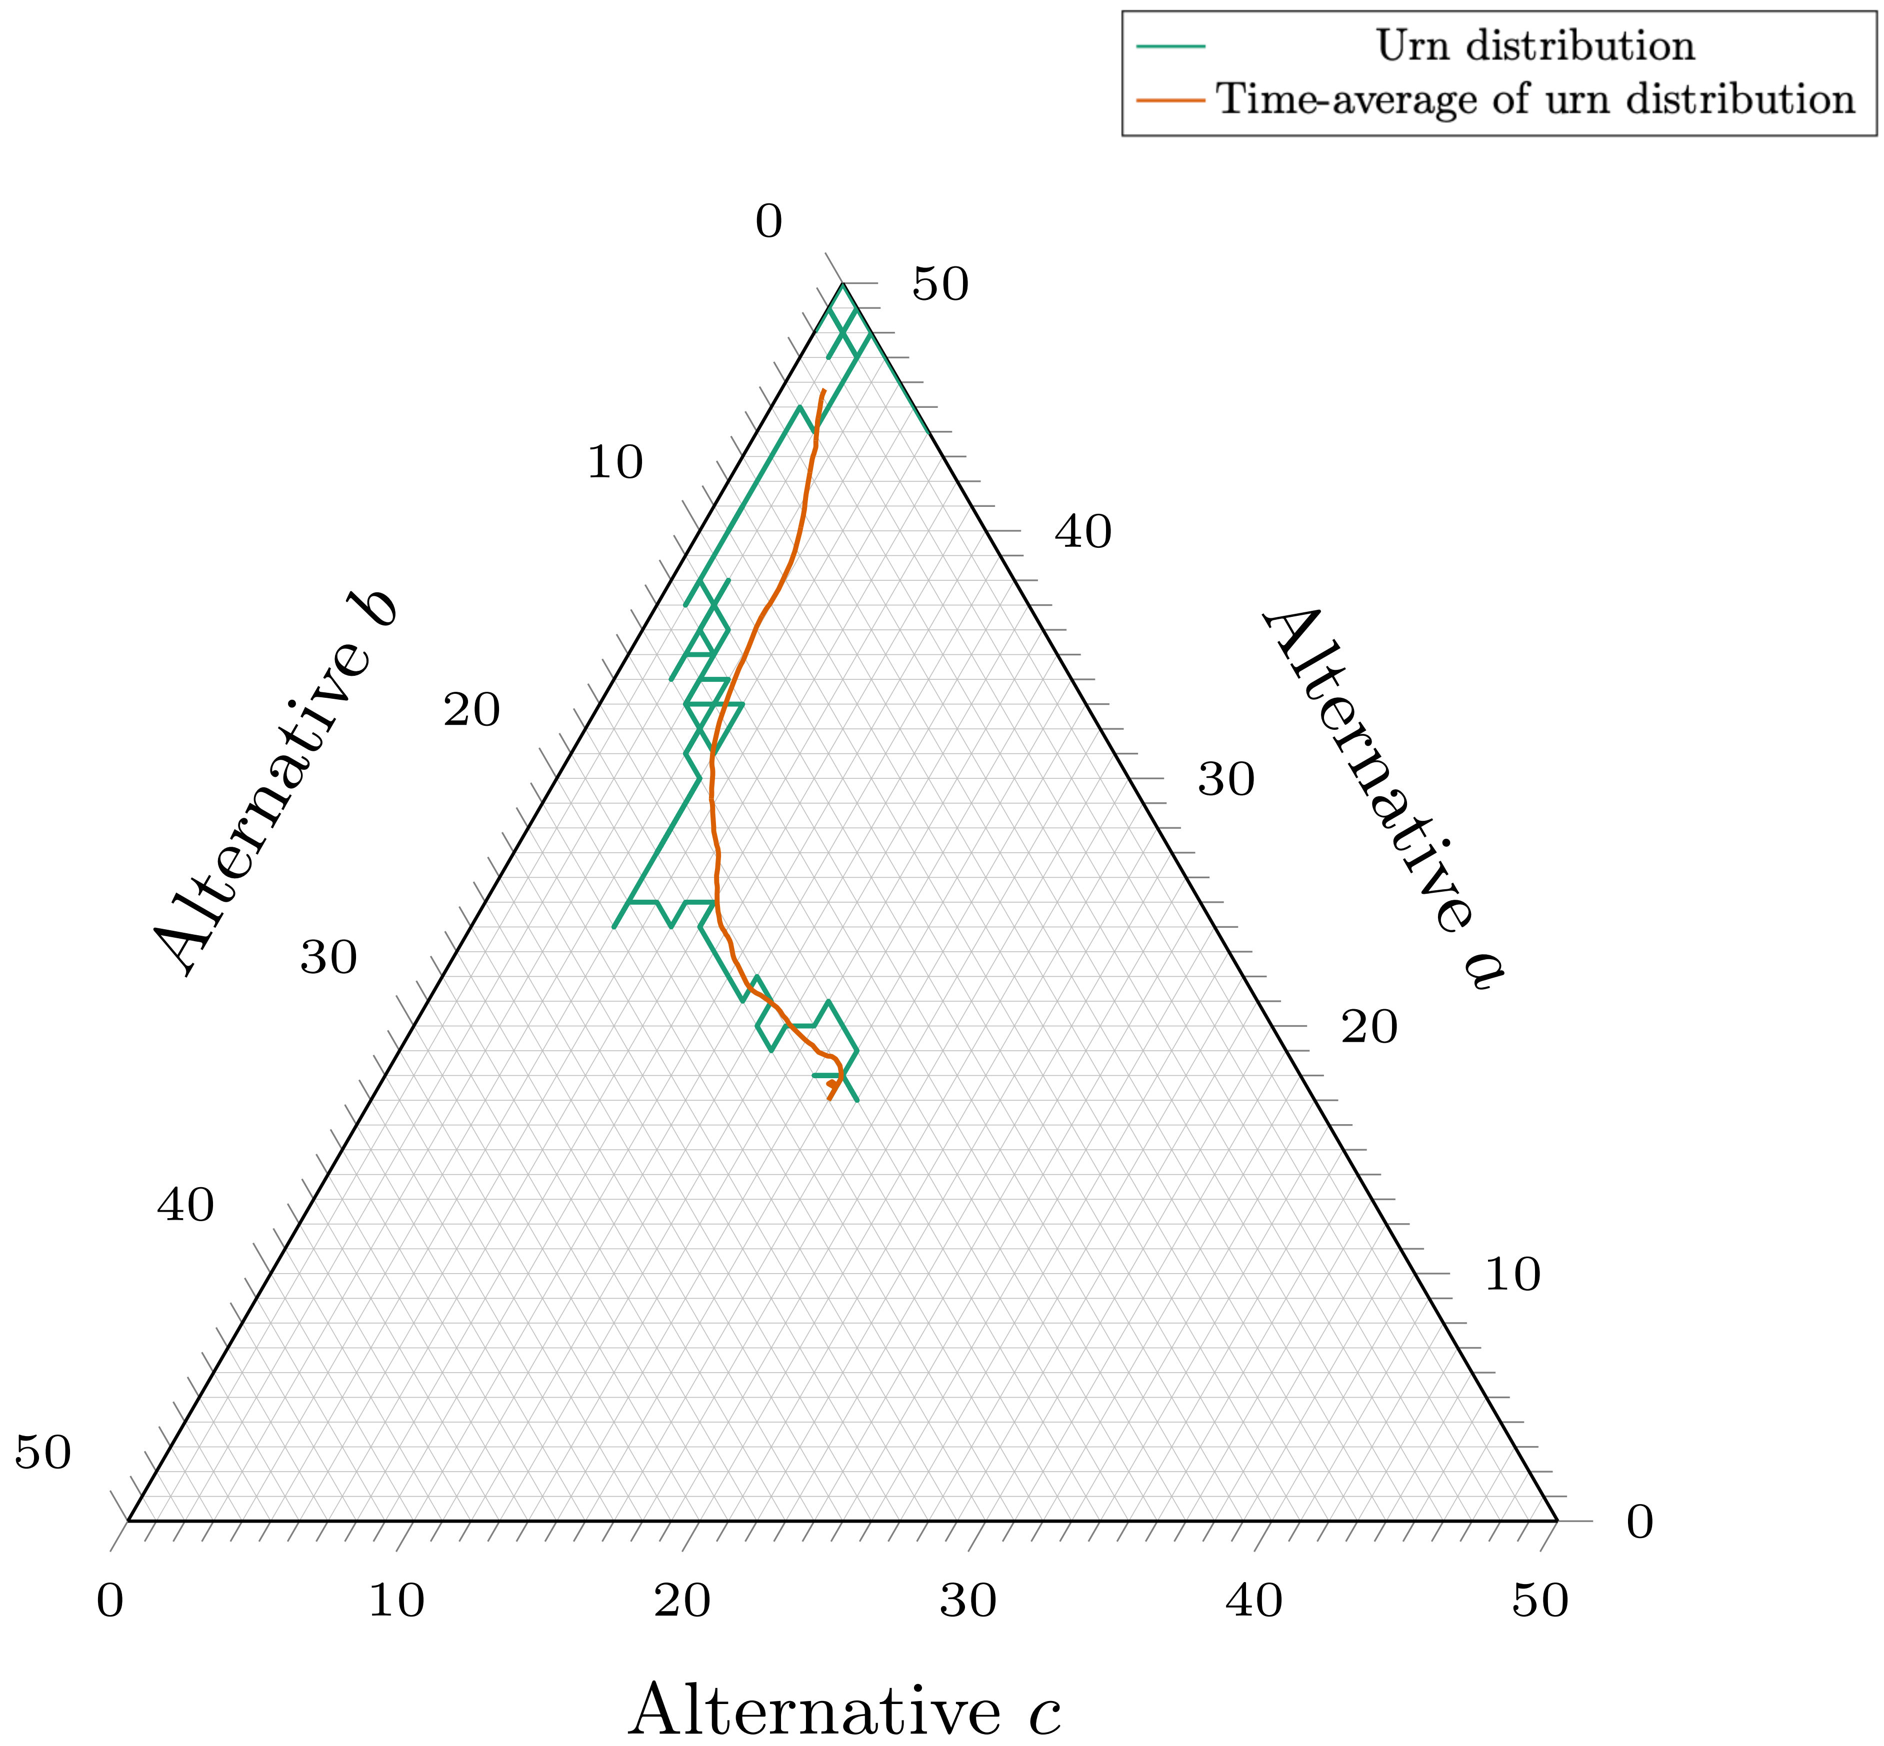

Main Result

- If N is sufficiently large, the relative urn distribution is almost surely close to a **maximal lottery** most of the time.
  - Hence, the distribution of winners is almost surely close to a maximal lottery.
- Moreover, the probability that the relative urn distribution is often close to a maximal lottery converges exponentially fast.

The Model

- Finite sets of alternatives A and voters
- Voters have complete and transitive preferences over alternatives → preference profile R
- $M_R \in \mathbb{R}^{A \times A}$ comparison matrix
  - $M_R(x, y)$ fraction of voters preferring x to y
  - $\tilde{M}_R = M_R - M_R^T$ matrix of majority margins
- Δ set of lotteries over A

R

2	**1**	**1**
a	b	c
b	a	a
c	c	b

$$\tilde{M}_R = \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ -\frac{1}{2} & 0 & \frac{1}{2} \\ -\frac{1}{2} & -\frac{1}{2} & 0 \end{pmatrix}$$

Germain Kreweras
† March 4, 1998

Maximal Lotteries

Peter C. Fishburn
† June 10, 2021

- Randomized voting rule proposed by Kreweras (1965) and Fishburn (1984).
  - rediscovered by Laffond et al. (1993), Felsenthal and Machover (1992), Fisher and Ryan (1995), Rivest and Shen (2010)
- A lottery p is maximal for a profile R if $p^T \tilde{M}_R q \geq 0$ for any other lottery q.
  - equivalently, $p^T \tilde{M}_R \geq \mathbf{0}$
  - mixed equilibrium strategy of the symmetric zero-sum game $\tilde{M}_R$
  - randomized Condorcet winner

Germain Kreweras
† March 4, 1998

Maximal Lotteries

Peter C. Fishburn
† June 10, 2021

5	4	3	2
a	*e*	*d*	*b*
c	*b*	*c*	*d*
b	*c*	*b*	*e*
d	*d*	*e*	*c*
e	*a*	*a*	*a*

$$\begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}^T \begin{array}{c} \\ a \\ b \\ c \\ d \\ e \end{array} \begin{array}{c} \begin{array}{ccccc} a & b & c & d & e \end{array} \\ \begin{pmatrix} 0 & -4 & -4 & -4 & -4 \\ 4 & 0 & -2 & 8 & 6 \\ 4 & 2 & 0 & 4 & 2 \\ 4 & -8 & -4 & 0 & 6 \\ 4 & -6 & -2 & -6 & 0 \end{pmatrix} \end{array} = \begin{pmatrix} 4 & 2 & 0 & 4 & 2 \end{pmatrix} \geq 0$$

- A lottery p is maximal if $p^T \tilde{M}_R \geq \mathbf{0}$.

Germain Kreweras
† March 4, 1998

Maximal Lotteries

Peter C. Fishburn
† June 10, 2021

5	**4**	**3**	**2**
a	*e*	*d*	*b*
c	*b*	*c*	*d*
b	*d*	*b*	*e*
d	*c*	*e*	*c*
e	*a*	*a*	*a*

$$\begin{pmatrix} 0 \\ 2/7 \\ 4/7 \\ 1/7 \\ 0 \end{pmatrix}^T \begin{matrix} \\ a \\ b \\ c \\ d \\ e \end{matrix} \overset{\begin{matrix} a & b & c & d & e \end{matrix}}{\begin{pmatrix} 0 & -4 & -4 & -4 & -4 \\ 4 & 0 & -2 & 8 & 6 \\ 4 & 2 & 0 & -4 & 2 \\ 4 & -8 & 4 & 0 & 6 \\ 4 & -6 & -2 & -6 & 0 \end{pmatrix}} = (4\ 0\ 0\ 0\ 26/7) \geq 0$$

- A lottery p is maximal if $p^T \tilde{M}_R \geq \mathbf{0}$.

Maximal Lotteries

Germain Kreweras
† March 4, 1998

Peter C. Fishburn
† June 10, 2021

- Desirable properties
  - Condorcet-consistency
  - Consistency with respect to variable electorates
  - Invariance under removal of losing alternatives
  - Independence of clones
- Axiomatic Characterizations
  - ➔ B. & Brandt, *Arrovian Aggregation of Convex Preferences* (ECMA 2020)
  - ➔ B. et al., *Consistent Probabilistic Social Choice* (ECMA 2016)
  - ➔ B. et al., *Welfare Maximization Entices Participation* (GEB 2018)

Convergence Result

If N is sufficiently large, the relative urn distribution is almost surely close to a maximal lottery most of the time.

- Let $\delta, \varepsilon > 0$. Then there is $r_0 = r_0(\delta, \epsilon) > 0$ such that for all $0 < r \leq r_0$, there is $N_0 = N_0(\delta, \epsilon, r) \in \mathbb{N}$ such that for all $N \geq N_0$, and initial distributions p_0,
$$\lim_{n\to\infty} \frac{1}{n} \left| \left\{ k \leq n\colon \ \left| X^{(N,r)}(k, p_0) - p^* \right| \leq \delta \right\} \right| \geq 1 - \epsilon$$
almost surely, where p^* is a maximal lottery of the preference profile.
- The probability that the relative urn distribution is close to a maximal lottery gets arbitrarily close to 1 exponentially fast.
- The empirical distribution of winners almost surely approximates a maximal lottery.

100	100	100
a	b	c
b	c	a
c	a	b

- N=5,000 balls
- Mutation rate r=0.04
- 500,000 rounds

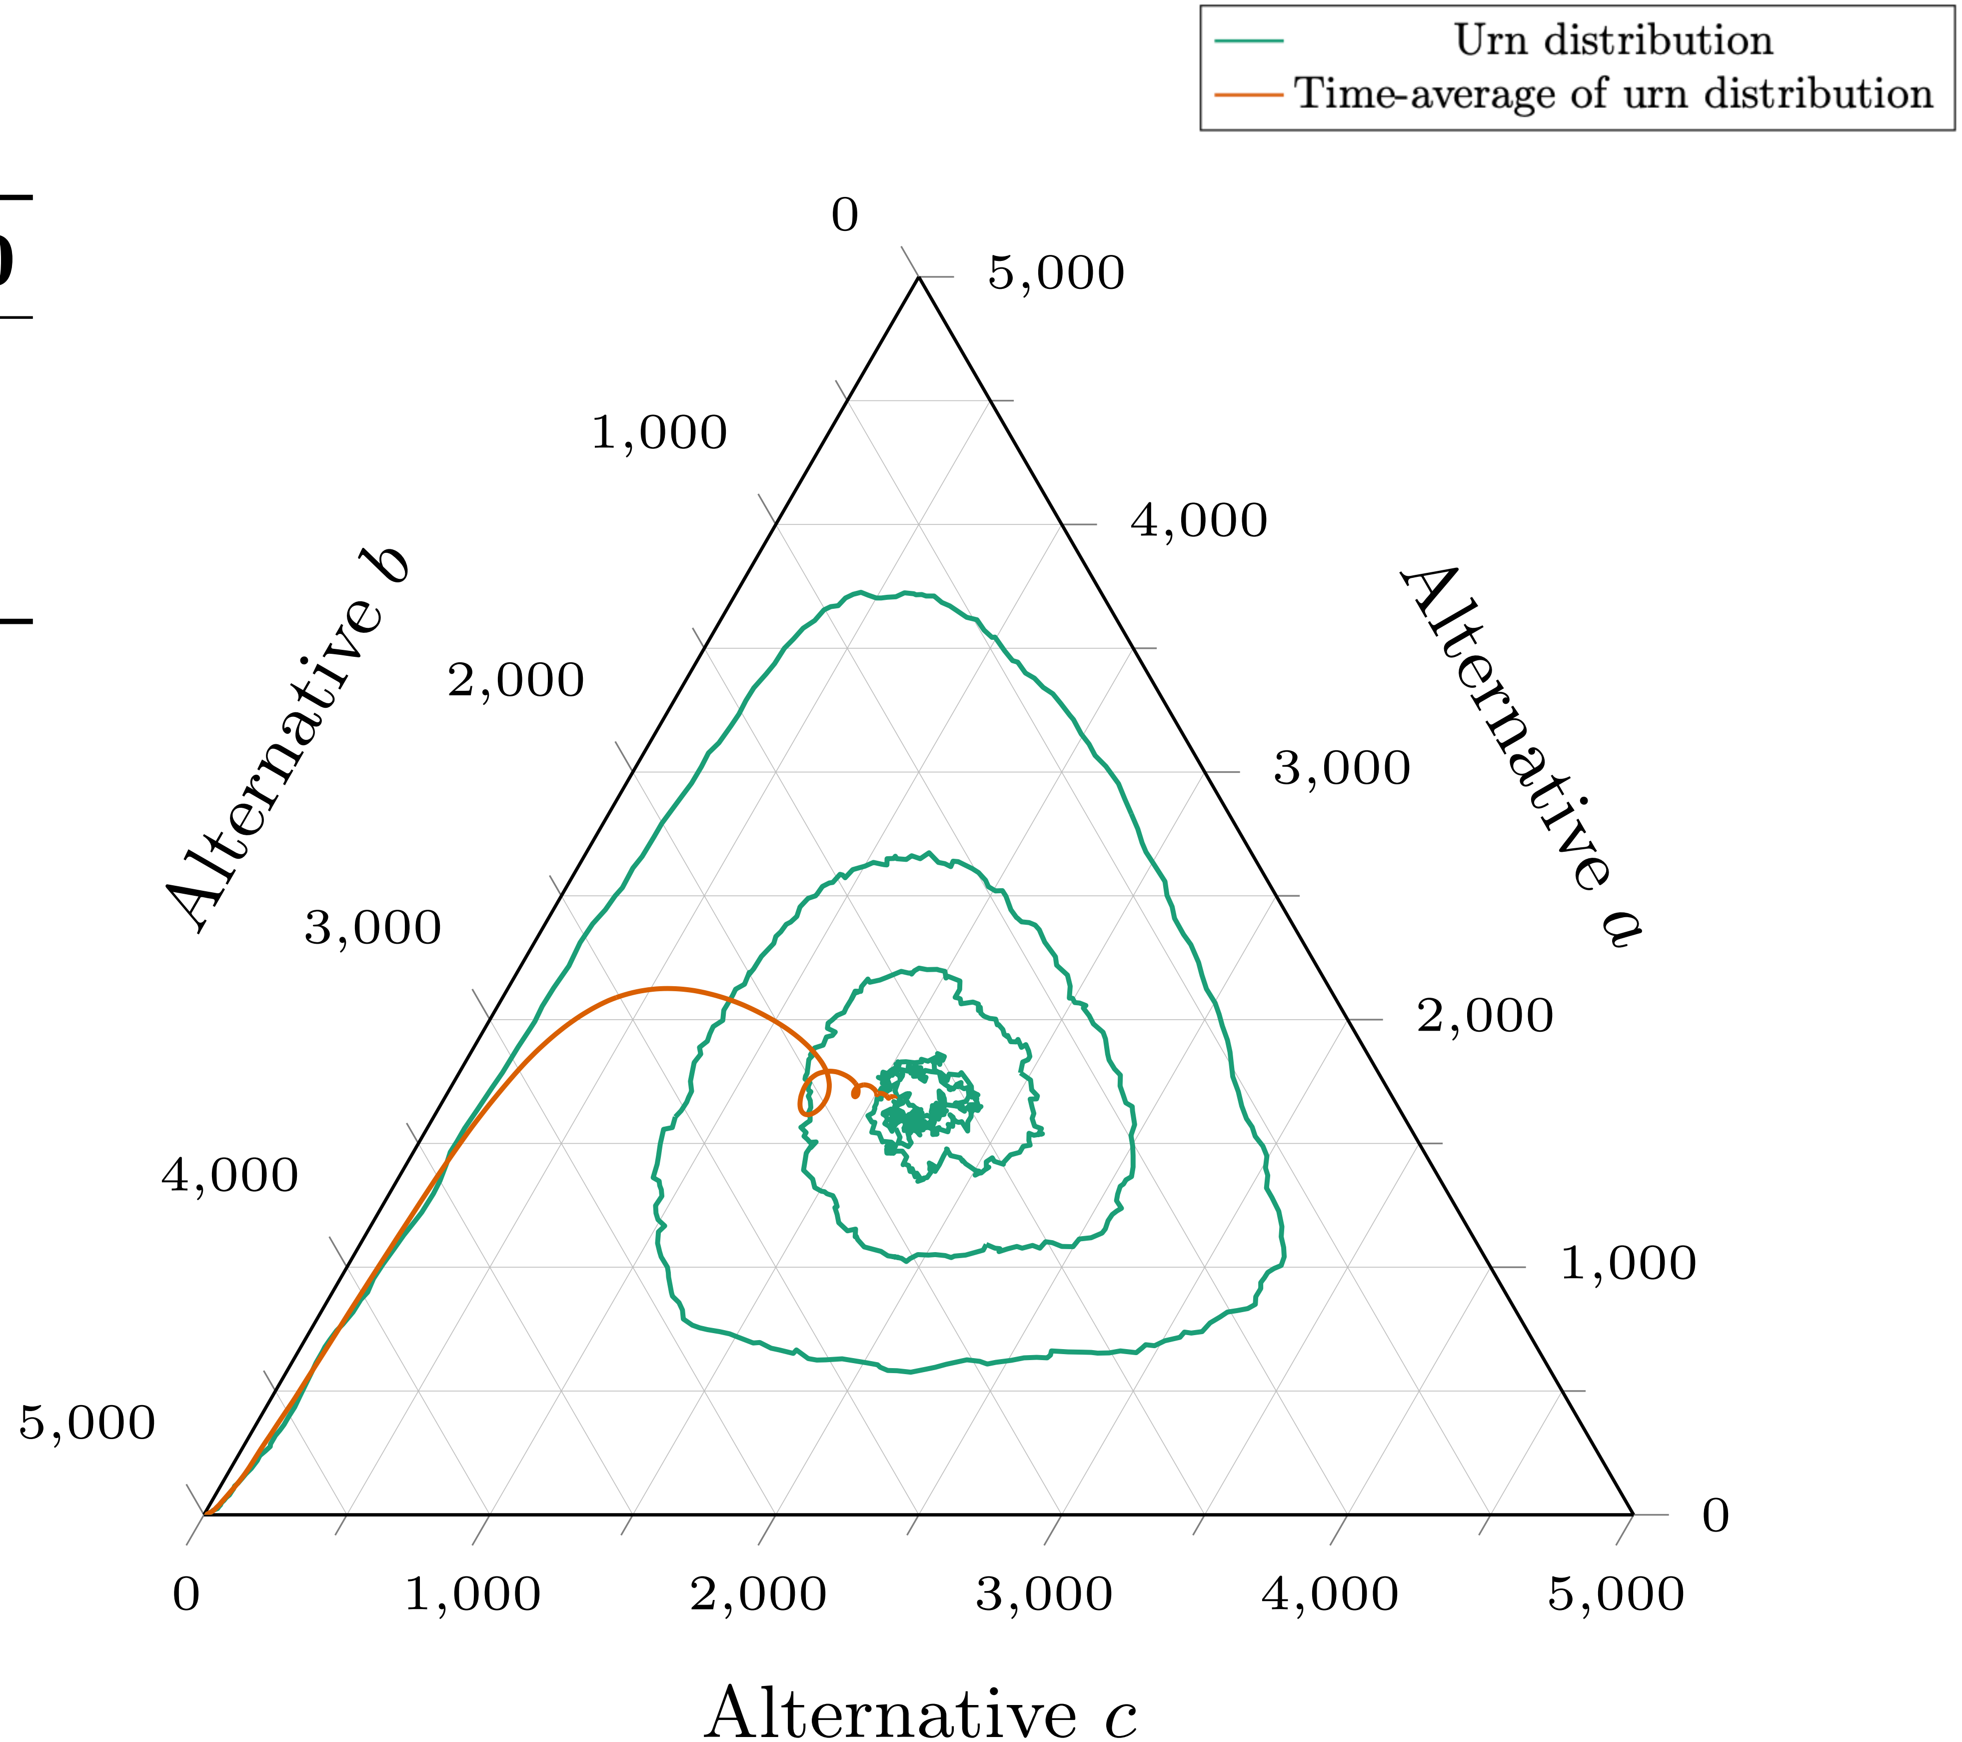

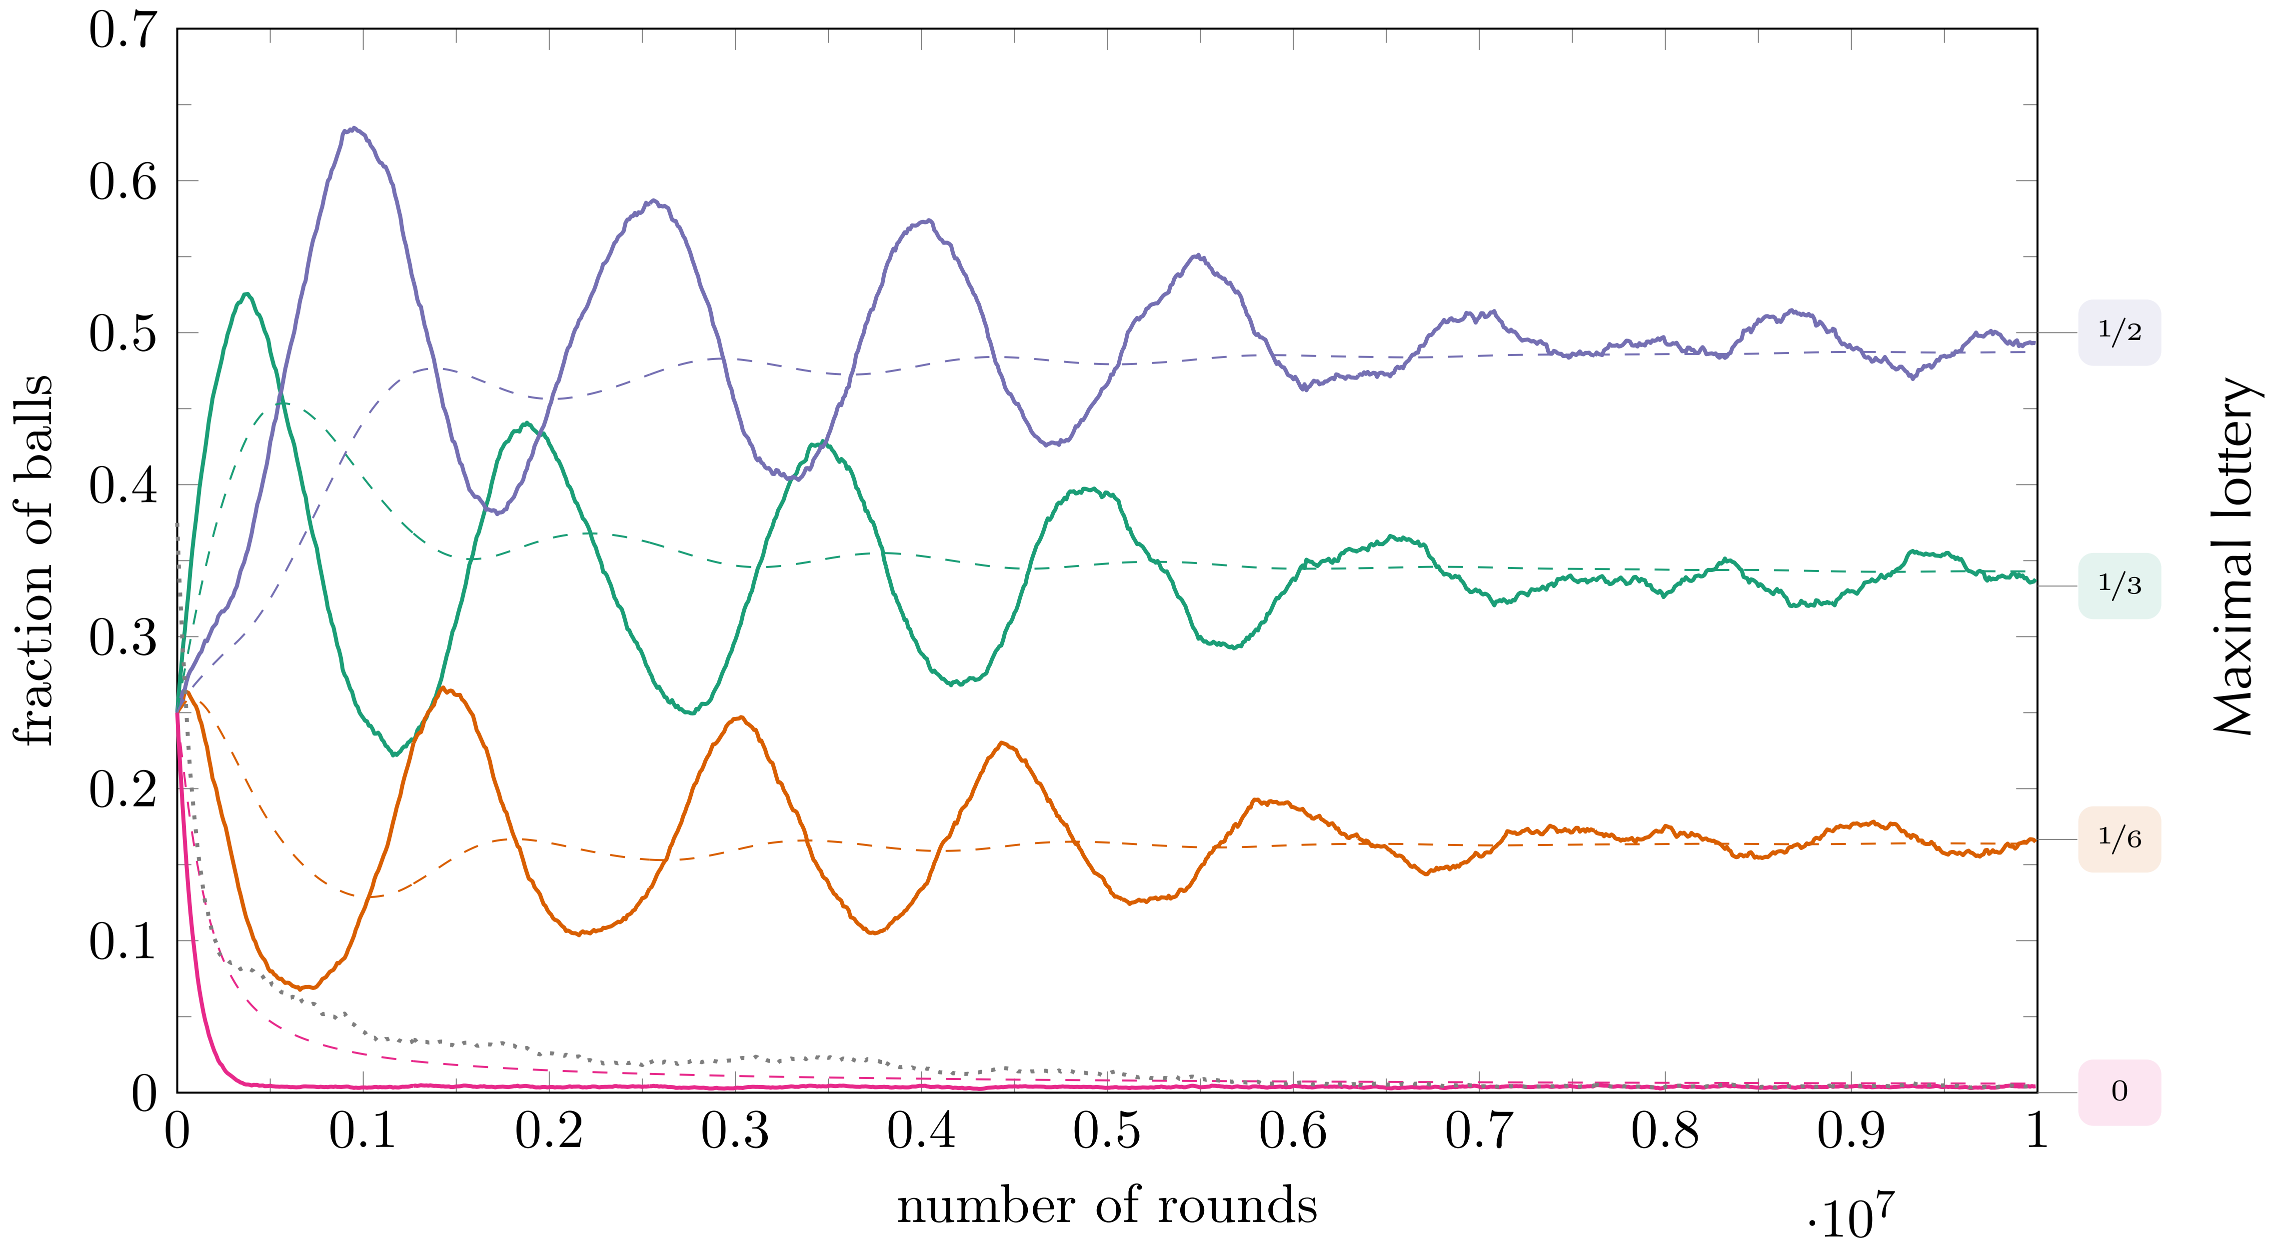

5	**4**	**3**
a	*c*	*d*
b	*a*	*b*
c	*b*	*c*
d	*d*	*a*

- N=50,000 balls
- Mutation rate r=0.01
- 10,000,000 rounds

Proof Sketch: Condorcet Winner Case

100	**100**	**100**
a	*a*	*b*
b	*c*	*c*
c	*b*	*a*

- N=50 balls
- Mutation rate r=0.02
- 1000 rounds

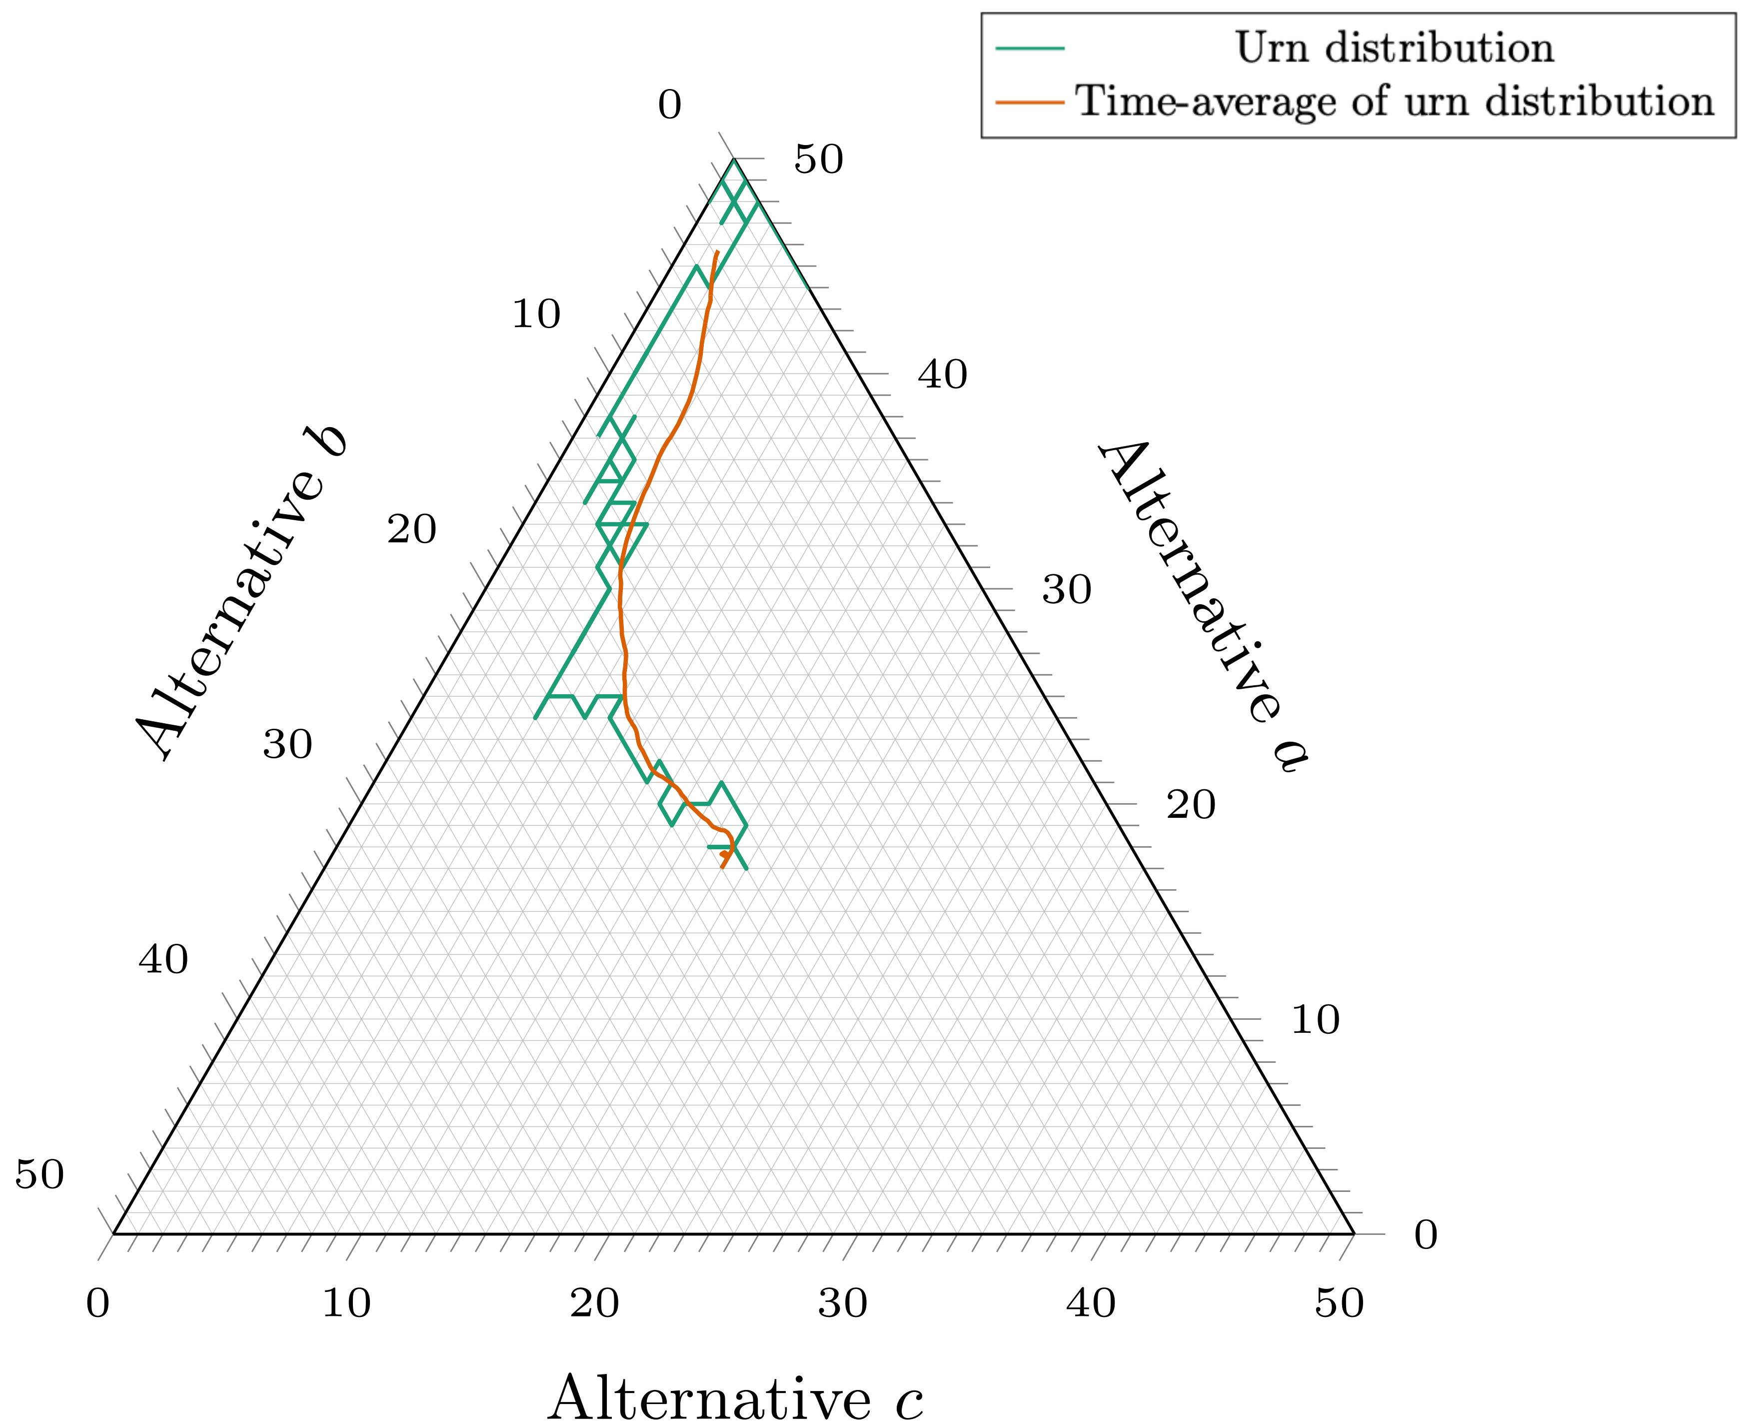

Proof Sketch (1/4)

- Consider the expected change of $X^{(N,r)}(n)$.
  $\mathbb{E}\left(X^{(N,r)}(n+1) - X^{(N,r)}(n) \,|\, X^{(N,r)}(n) = p\right)$
- Induces a continuous vector field $f^{(r)}\colon \Delta \to \mathbb{R}^A$ with
  $$f_x^{(r)}(p) = (1-r)2p_x(\tilde{M}_R p)_x + r\left(\frac{1}{|A|} - p_x\right).$$
- For $r > 0$, $f^{(r)}$ has a unique zero $p^{(r)}$, which is close to a maximal lottery p^* if r is small.

Proof Sketch (2/4)

- The differential equation
  $\dot{y}(t) = f^{(r)}(y(t)) \qquad y(0) = p$
  has a unique solution $y^{(r)}(\,\cdot\,, p) \colon [0,\infty) \to \Delta$.
- The relative entropy of $p, q \in \Delta$ is $D(p \mid q) = \sum_{x \in A} p_x \log\left(p_x / q_x\right)$.
- Considering the time-derivative of the relative entropy between $p^{(r)}$ and $y^{(r)}$ gives
  $\dot{D}(p^{(r)} \mid y^{(r)}(t,p)) \leq -\, C \left| p^{(r)} - y^{(r)}(t,p) \right|^2$
  for some $C > 0$.
  $\rightarrow$ $y^{(r)}(t,p)$ converges to $p^{(r)}$ uniformly in p.
- Choose $T > 0$ so that $y^{(r)}(t,p)$ is close to $p^{(r)}$ for all but a small fraction of $t \in [0,T]$ for all $p \in \Delta$.

Proof Sketch (3/4)

- For large N, $X^{(N,r)}(Nt, p_0)$ approximately solves the differential equation
  $\dot{y}(t) = f^{(r)}(y(t)) \qquad y(0) = p$
  on bounded time intervals with high probability.
  $\rightarrow X^{(N,r)}(Nt, p)$ is close to $y^{(r)}(t, p)$ for all $t \in [0,T]$ with probability close to 1 for large N.
- Thus, almost surely, $X^{(N,r)}(Nt, p_0)$ is close to $y^{(r)}(t, p)$ within all but a small fraction of the intervals $[0,T]$, $[T,2T]$, $[2T,3T]$, … provided they start at the same point at the beginning of the interval.

Proof Sketch (4/4)

- $p^{(r)}$ is close to a maximal lottery p^* if r is small.
- $y^{(r)}(t, p)$ is close to $p^{(r)}$ for all but a small fraction of $t \in [0,T]$ for all $p \in \Delta$.
- $X^{(N,r)}(Nt, p_0)$ is close to $y^{(r)}(t, p)$ within all but a small fraction of the intervals $[0,T]$, $[T,2T]$, $[2T,3T]$, … almost surely, provided they start at the same point at the beginning of the interval.

➡ $X^{(N,r)}(n, p_0)$ is close to p^* almost surely for all but a small fraction of $n \in \mathbb{N}$.

Related Work

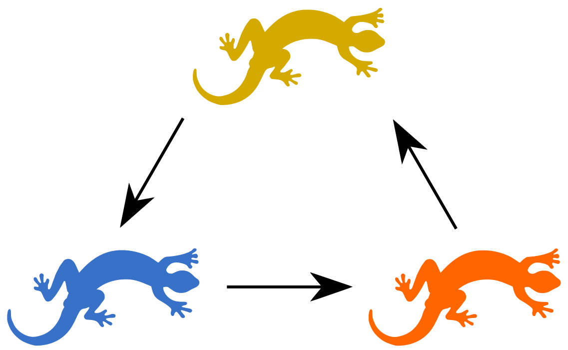

- Replicator equation in evolutionary game theory
- Similar processes studied in the natural sciences
  - Quantum physics: condensation of bosons
  - Population biology: coexistence of species
  - Chemical kinetics: reactions of molecules
  - Plasma physics: scattering of plasmons
  - E.g., Allesina and Levine (*PNAS* 2011), Knebel et al. (*Nat Commun* 2015), Laslier and Laslier (*Ann Appl Probab* 2017), Grilli et al. (*Nature* 2017)
- Differences of our model and result
  - discrete and stochastic (not continuous and deterministic)
  - interactions between pairs (not triples)
  - mutations
  - bound on occupation time (rather than only convergence of temp. avg.)

Conclusion

- Properties of urn process
  - Myopic strategyproofness
    - each round a randomly selected voter chooses between 2 alternatives
  - Minimal preference elicitation
    - isolated pairwise comparisons, privacy protection
  - Verifiability
    - simple physical procedure, no trusted authority
  - Flexibility
    - voters may arrive, leave, and change their preferences
- Alternative descriptive interpretation: opinion formation
  - Agents come together in random pairwise interactions, in which they try to convince each other of their opinion.
- The urn process approximately solves a linear program.