A Natural Adaptive Process for Collective Decision-Making

Felix Brandt

(joint work with Florian Brandl)

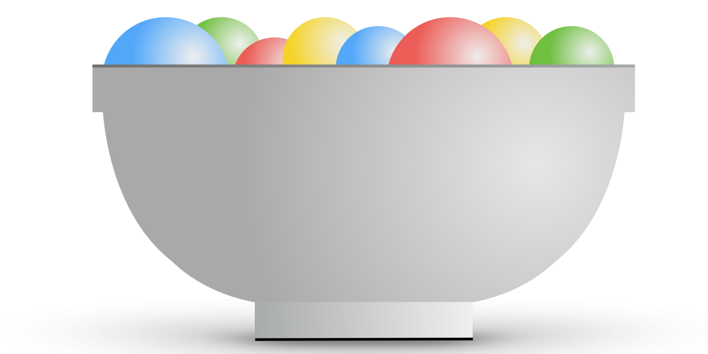

Symposium on Artificial Intelligence and the Social Sciences

Hi! Paris, June 2022

DSS

Decision Sciences & Systems

Voting

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

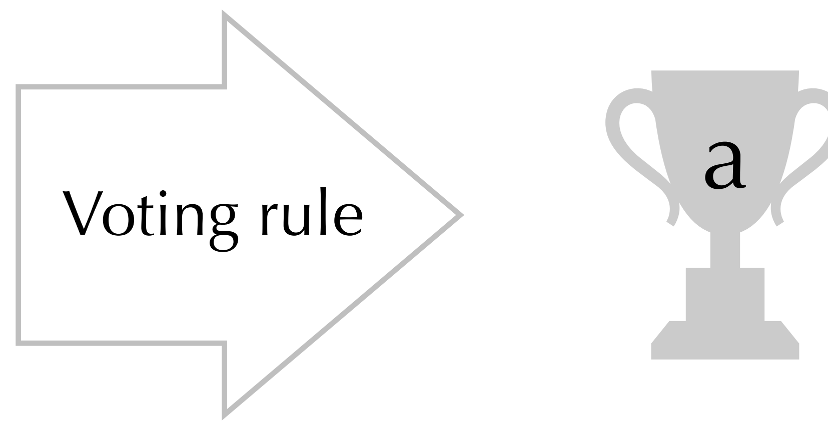

- Consider an ongoing dynamic voting process that aims for
  - Myopic strategyproofness
    - each round one voter chooses between two alternatives
  - Minimal preference elicitation
    - isolated pairwise comparisons, privacy protection
  - Verifiability
    - simple physical procedure, no trusted authority
  - Flexibility
    - voters may arrive, leave, and change their preferences

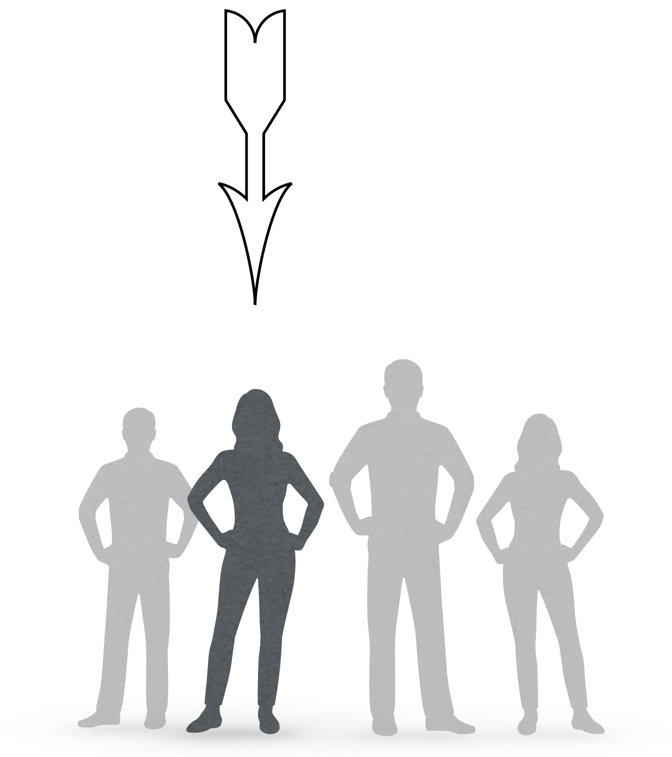

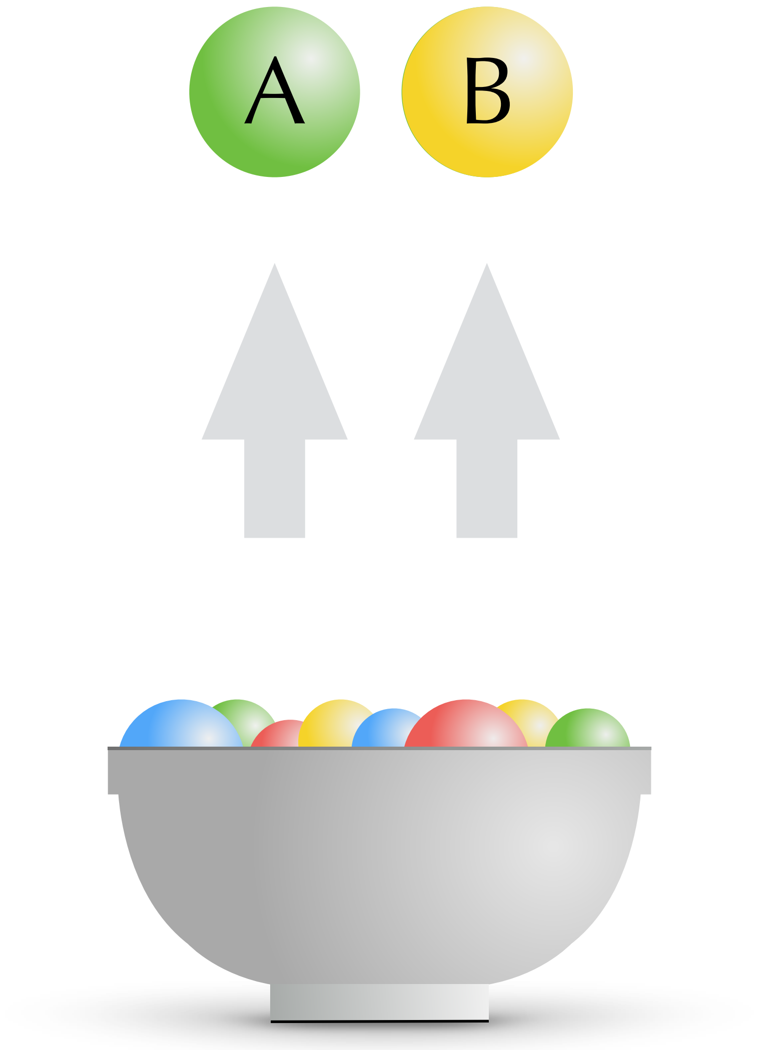
A
B

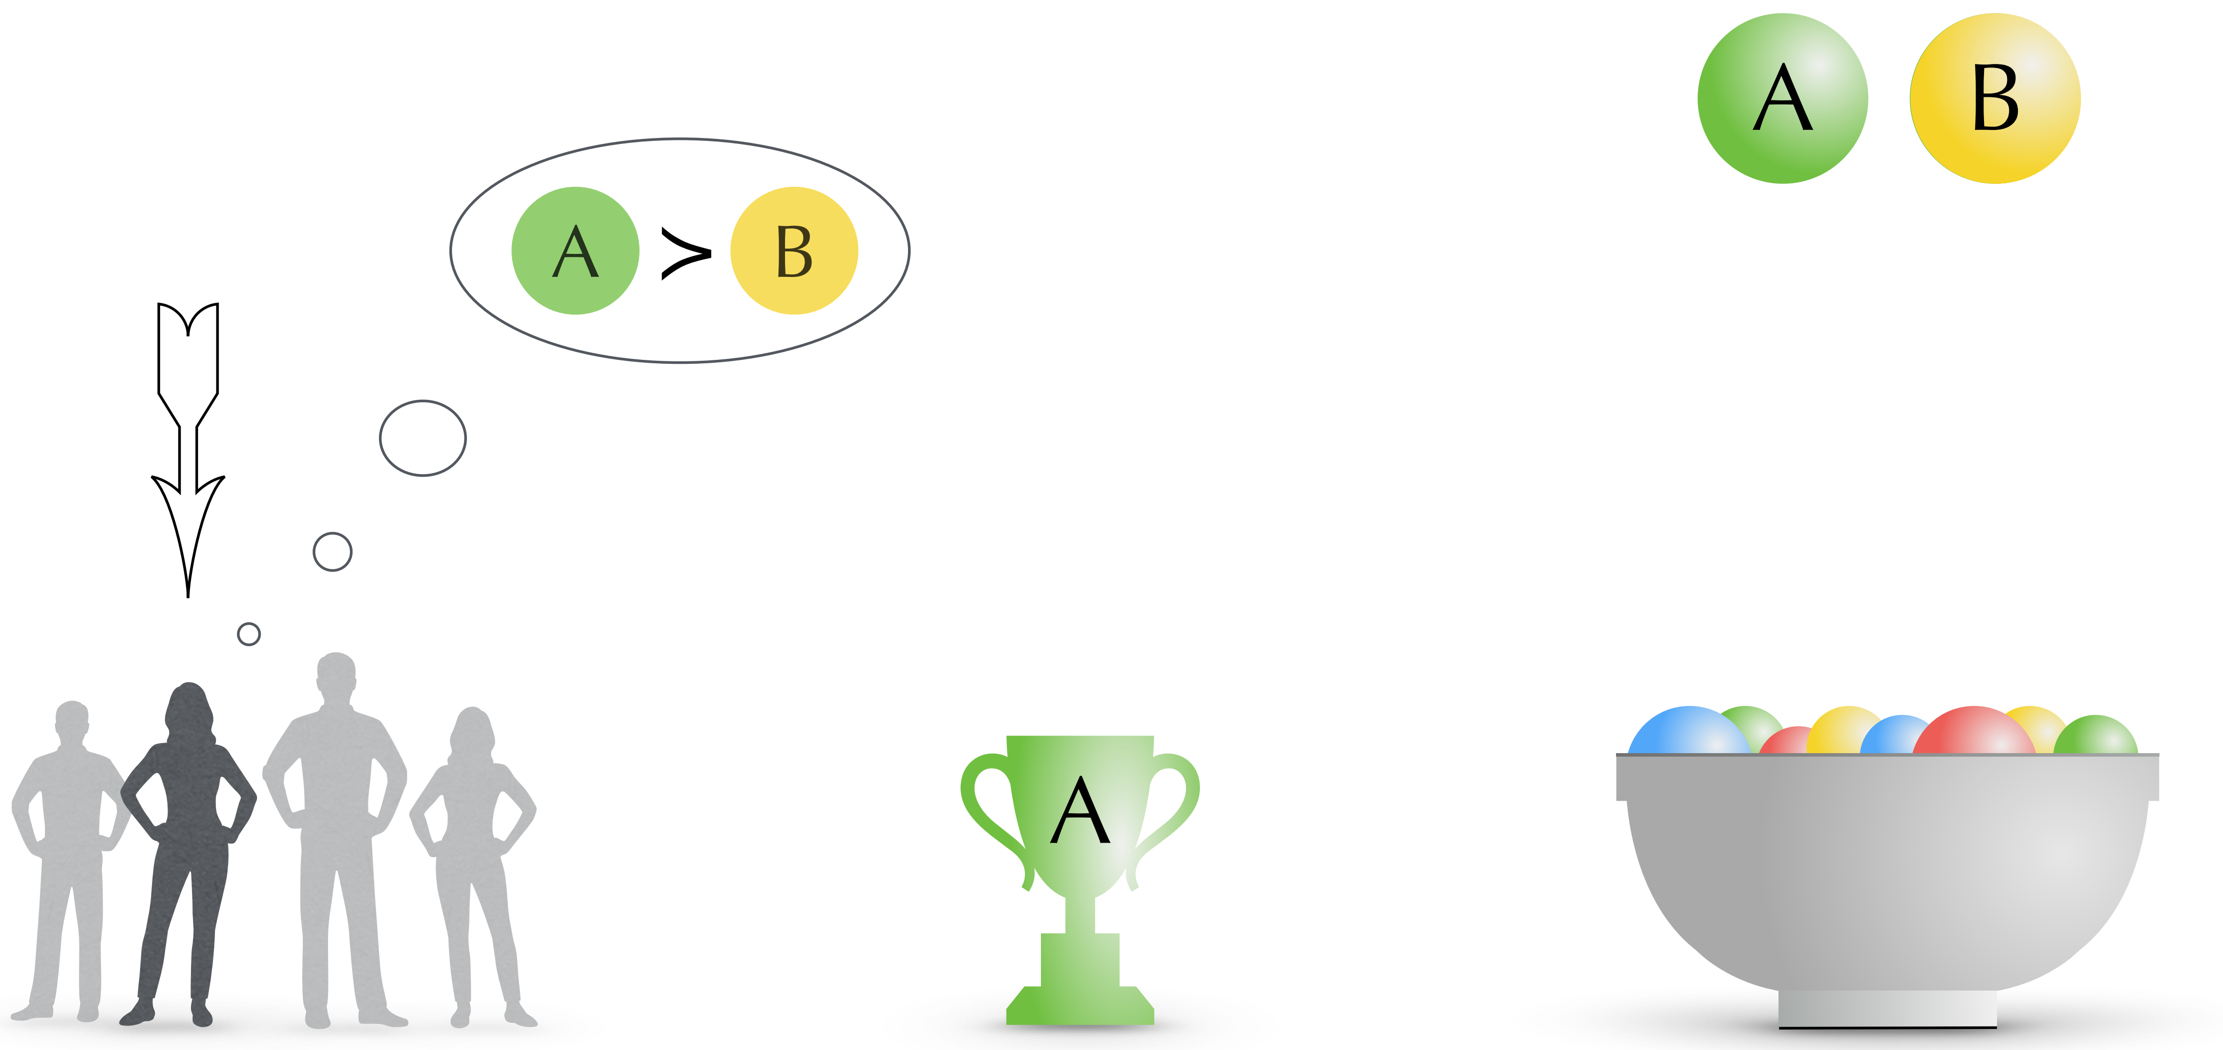
A ≻ B
A
A
B


A

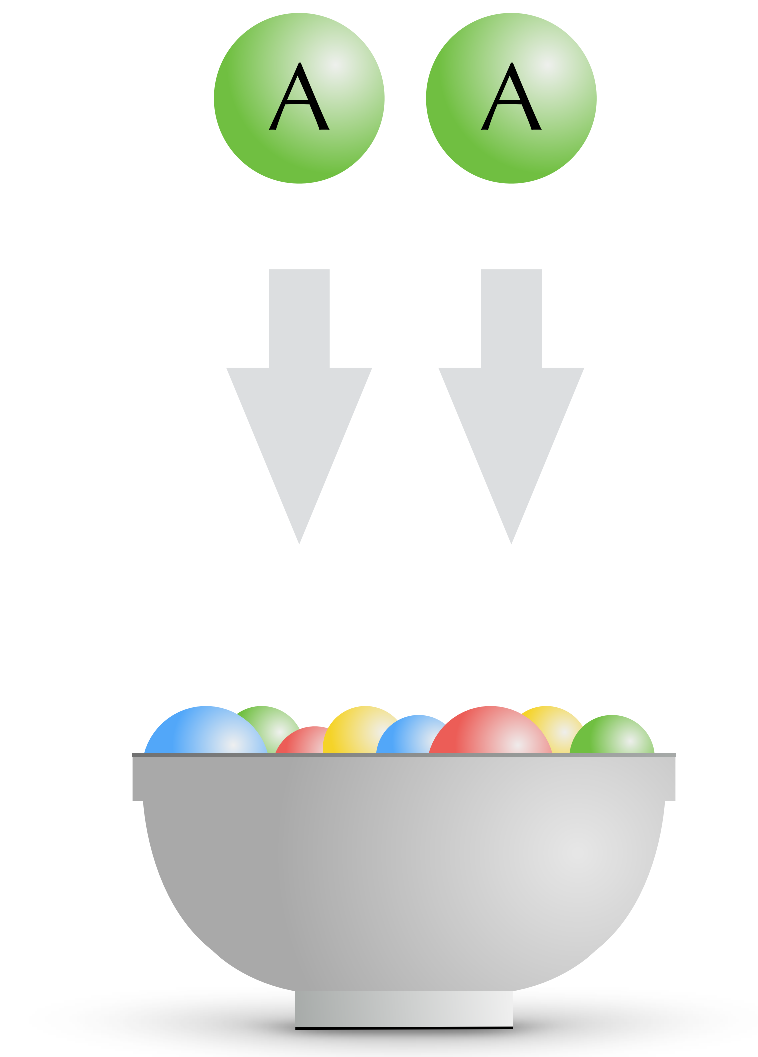
A
A

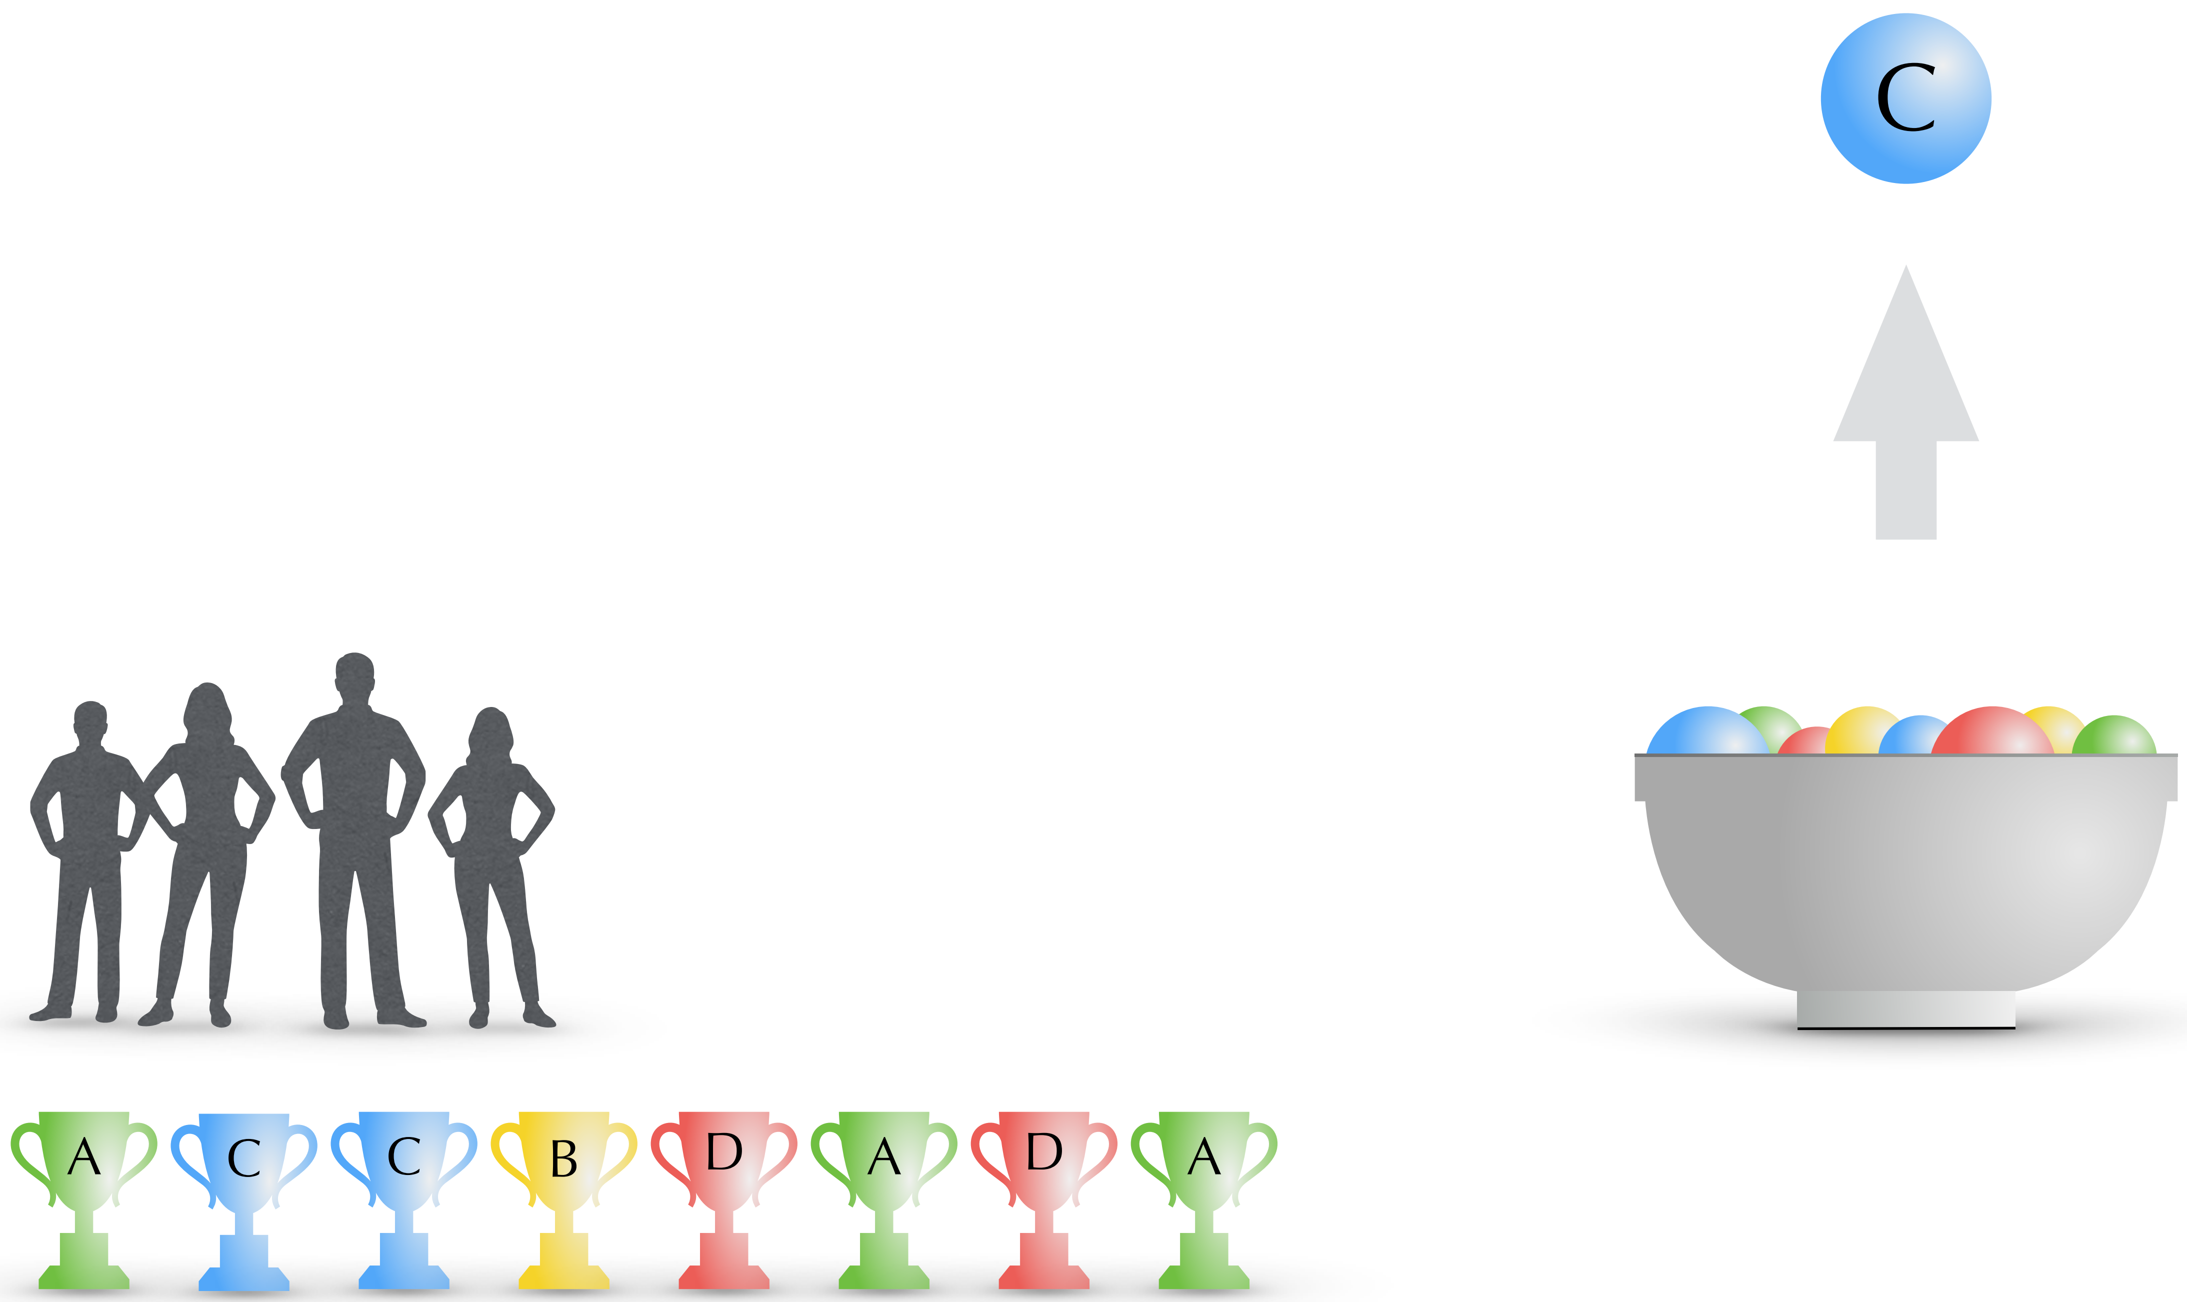
C
A
C
C
B
D
A
D
A

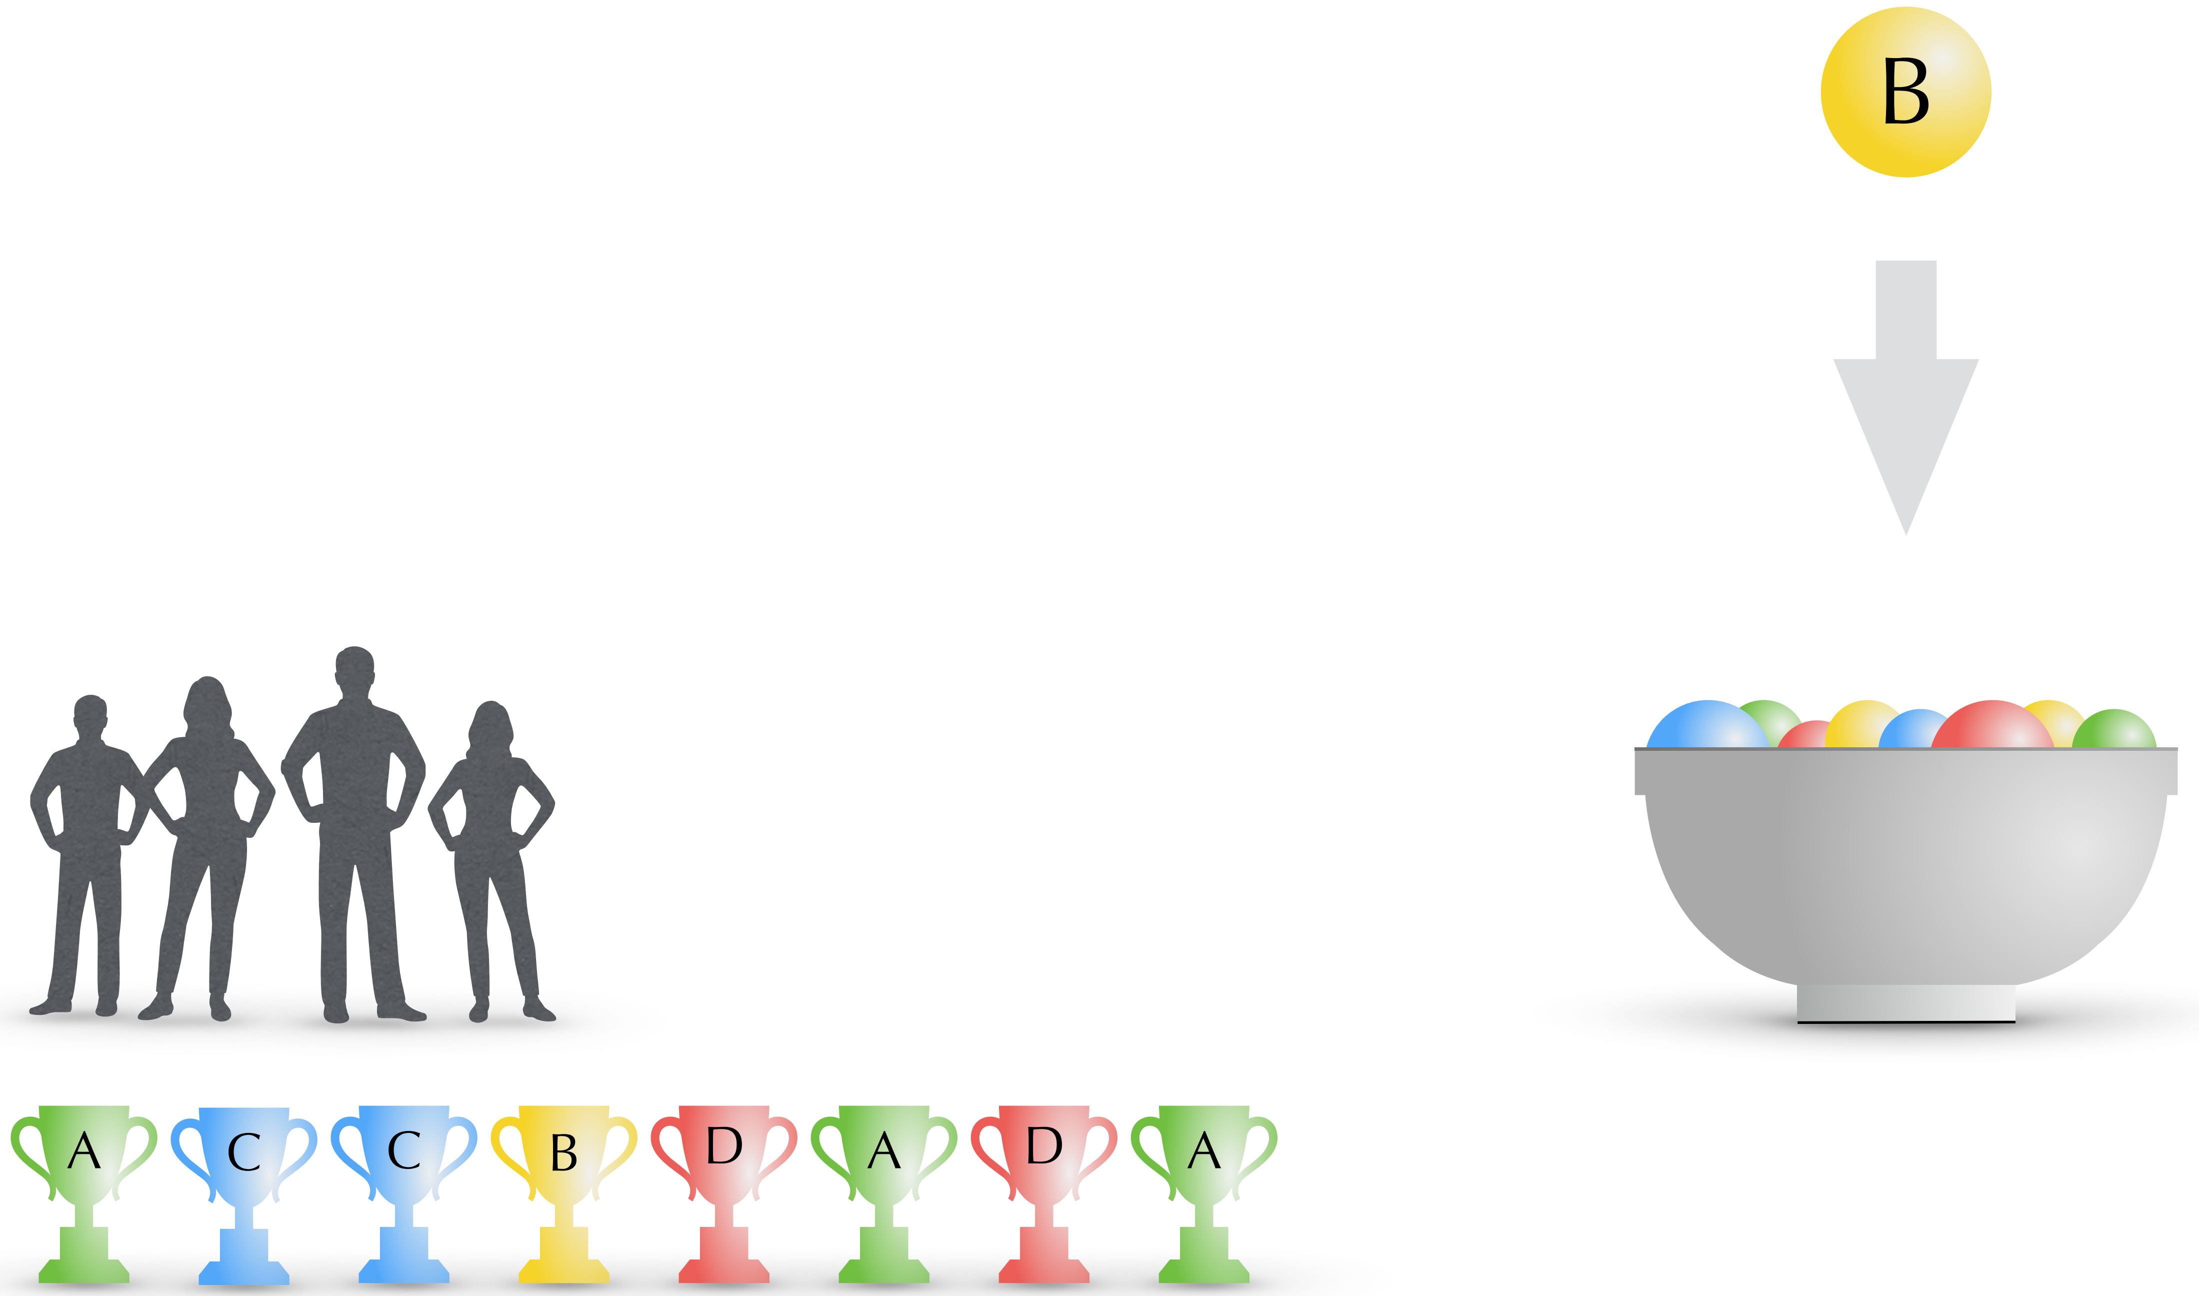
B
A
C
C
B
D
A
D
A


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

- What can be said about the sequence of winners?
- How about the empirical distribution of winners?
- 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 will draw 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

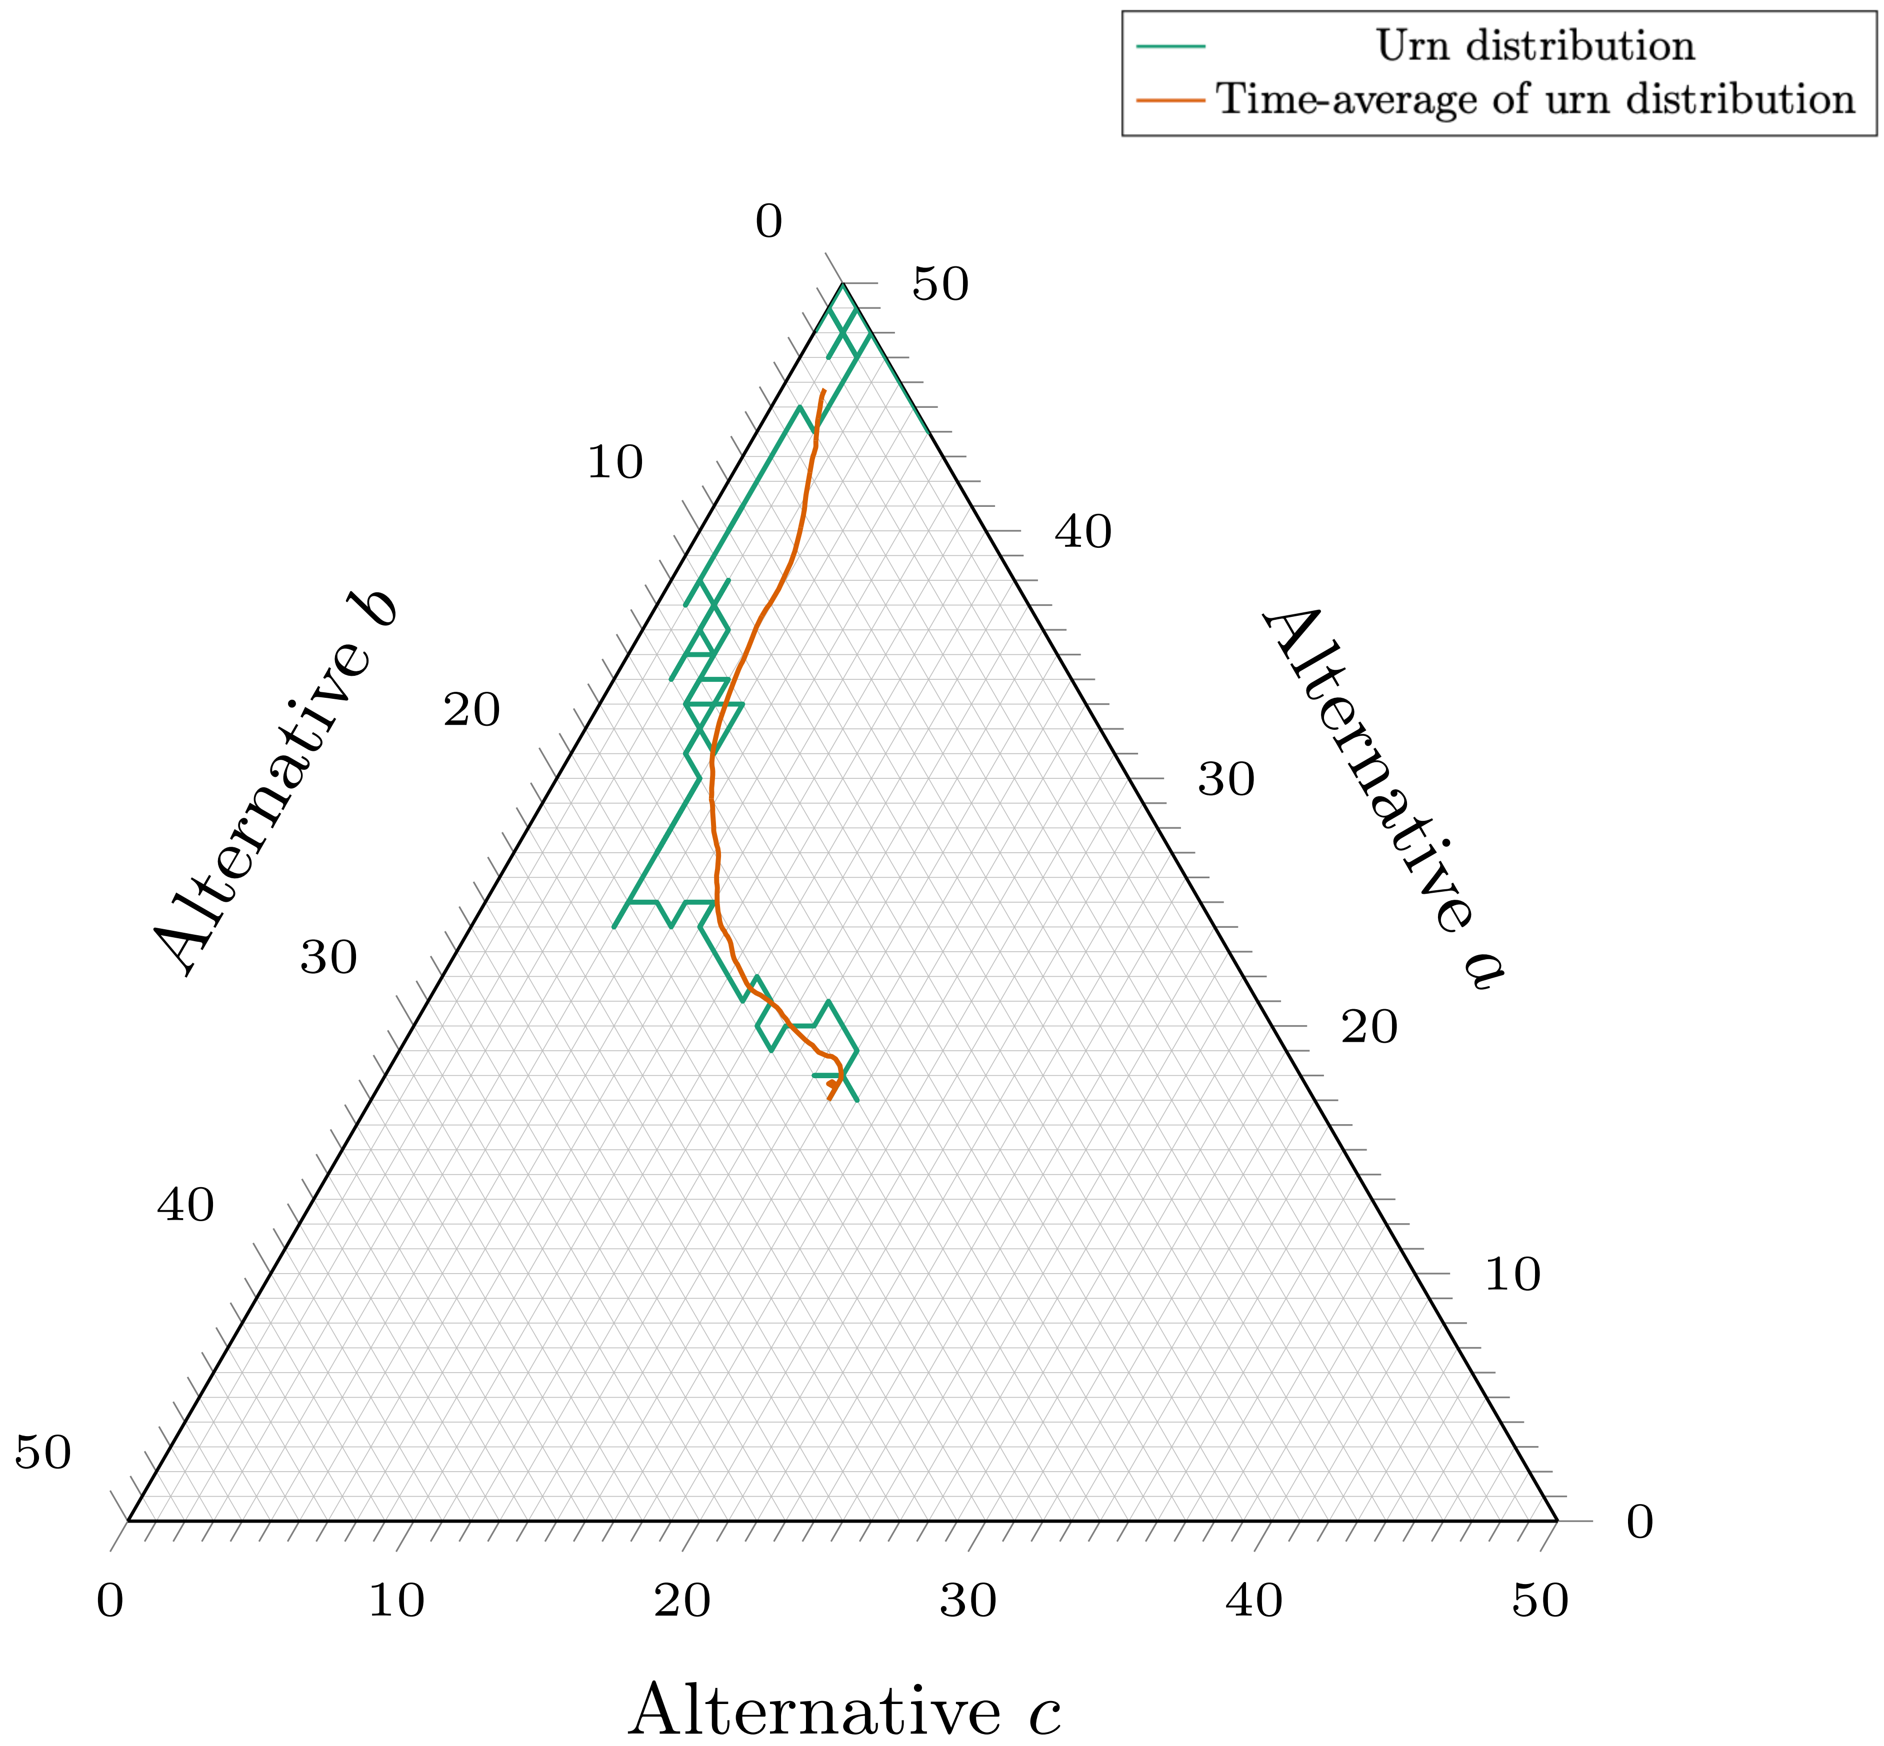

100	**100**	**100**
a	*b*	*c*
b	*c*	*a*
c	*a*	*b*

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

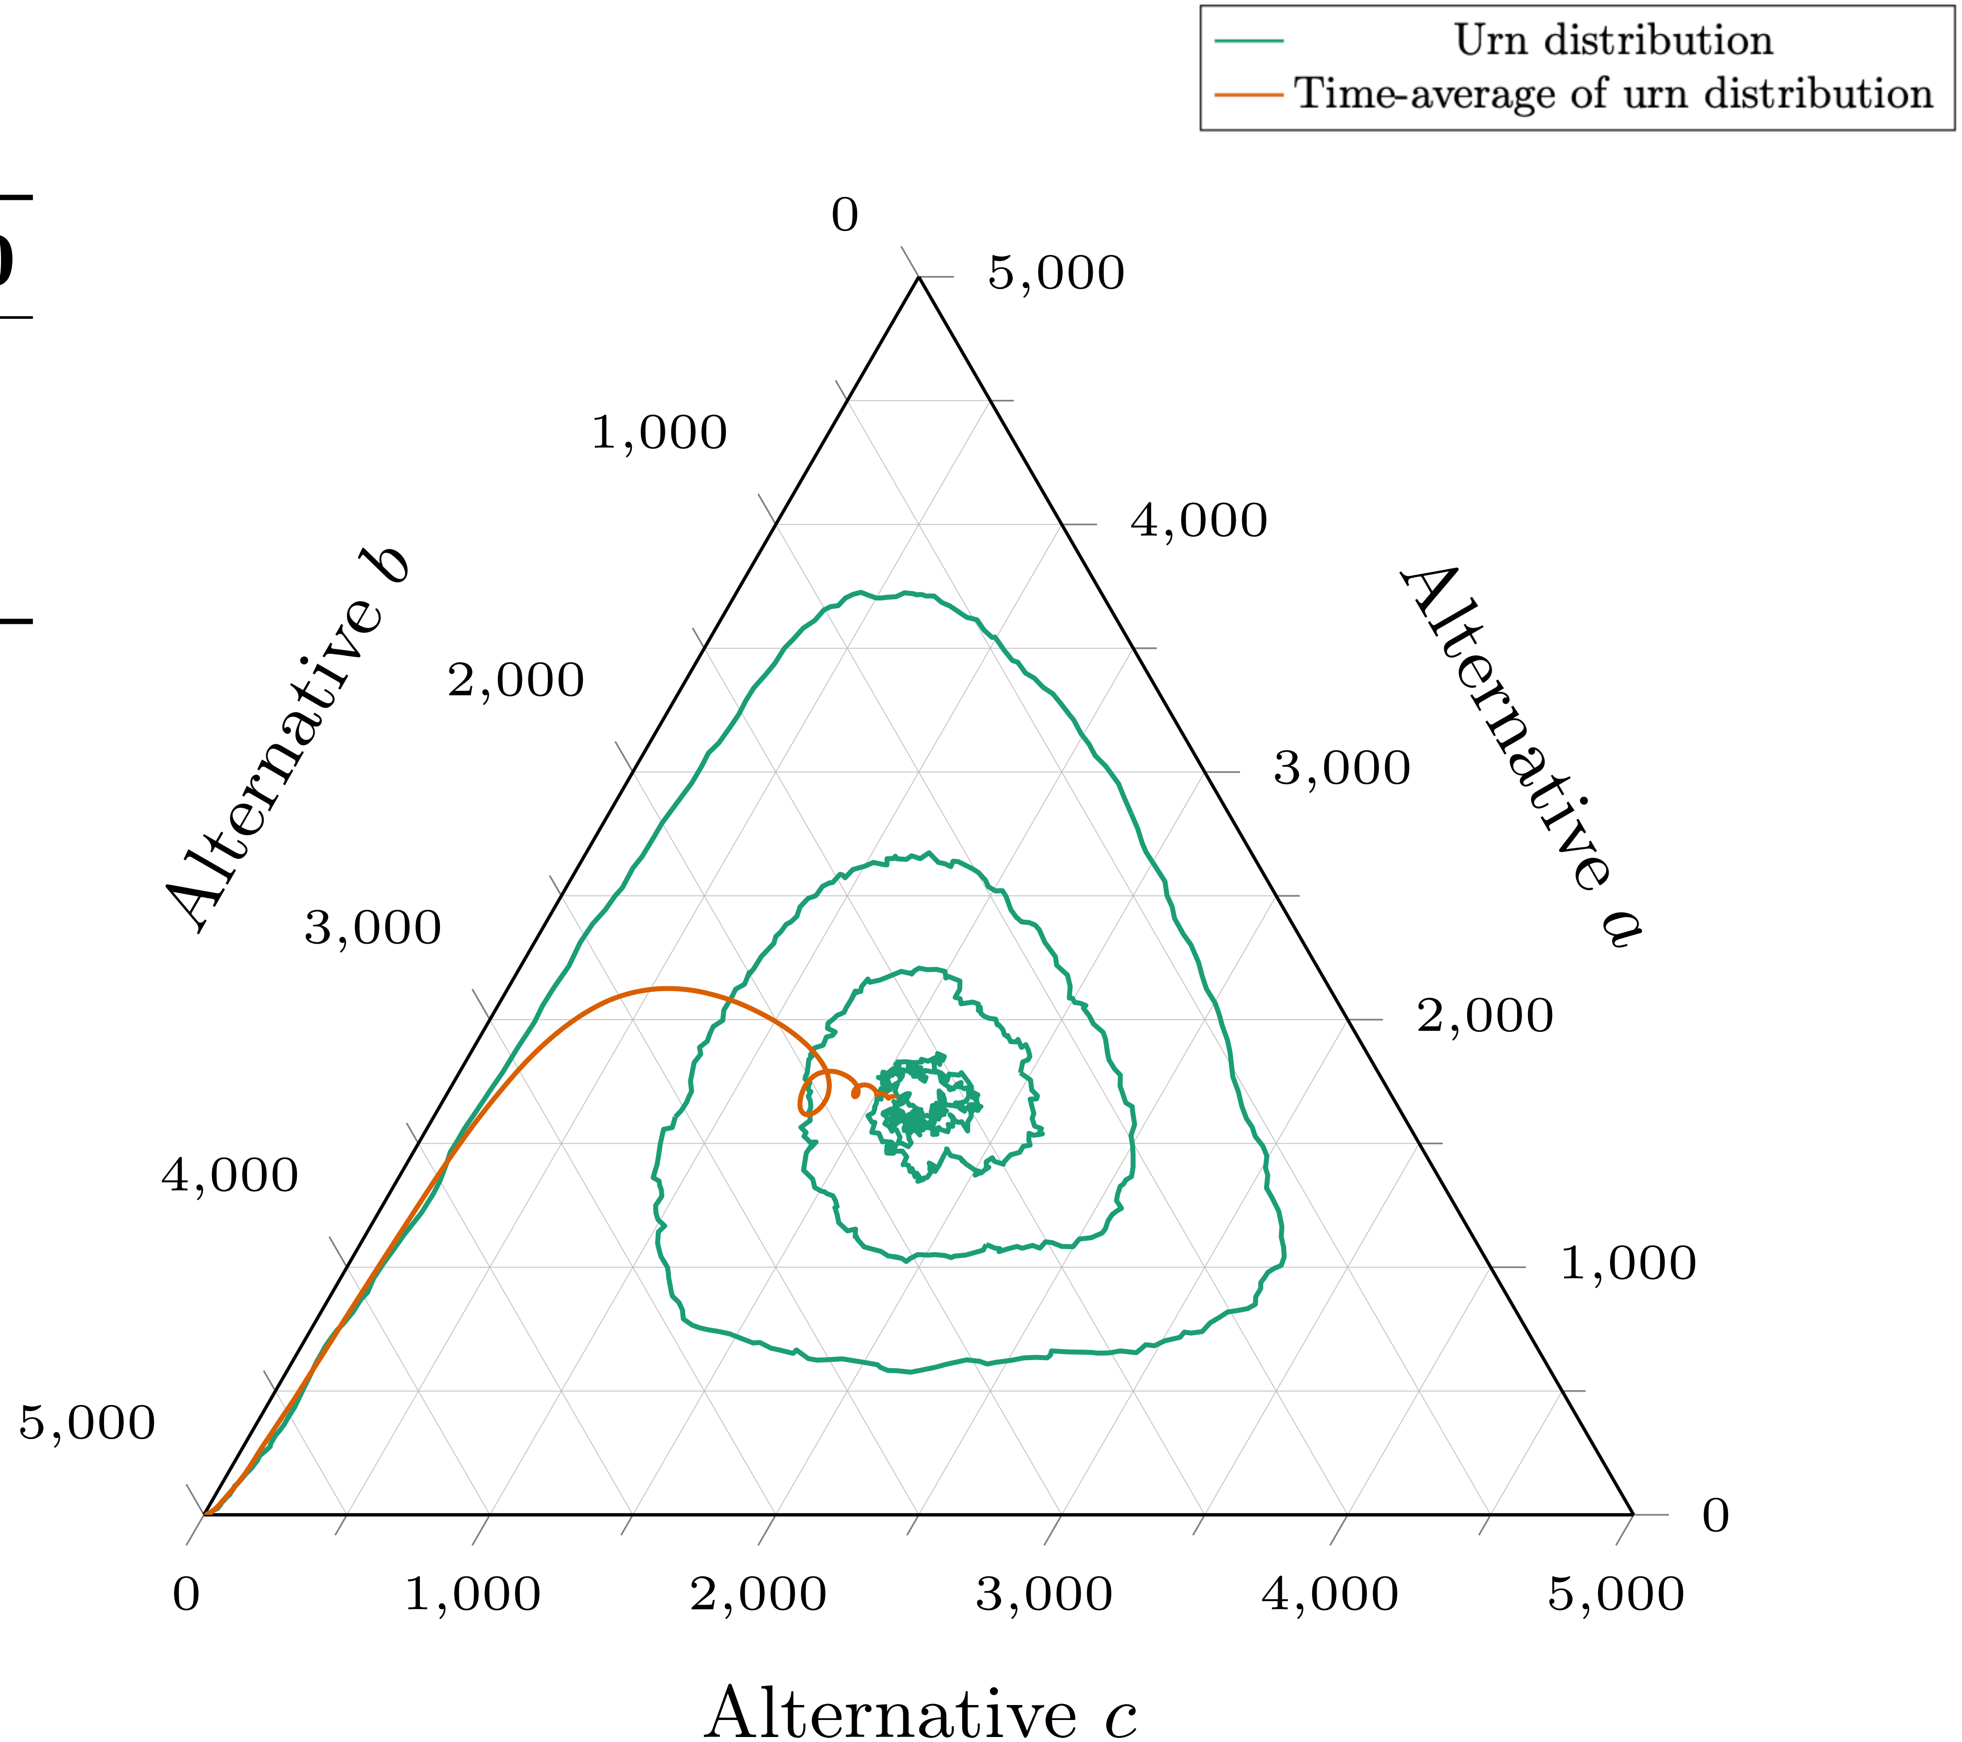

Convergence Result

- The empirical distribution of winners $W^{(N,r)}$ almost surely converges.
  - Let $\delta > 0$. Then there is $r_0 > 0$ such that for all $0 < r \leq r_0$, there is $N_0 \in \mathbb{N}$ such that for all $N \geq N_0$ and initial distributions s_0,
    $$\mathbb{P}\left(\left|\lim_{n\to\infty} W^{(N,r)}(n, s_0) - p^*\right| \leq \delta\right) = 1$$
    where p^* is a **maximal lottery** of the preference profile.
- More generally, we show that the relative urn distribution $X^{(N,r)}$ is almost surely close to a maximal lottery most of the time.
- The probability that the relative urn distribution is close to a maximal lottery gets arbitrarily close to 1 and converges exponentially fast.
  - Let $\delta, \varepsilon > 0$. Then there is $r_0 > 0$ such that for all $0 < r \leq r_0$, there is $N_0 \in \mathbb{N}$ and $C > 0$ such that for all $N \geq N_0$, s_0, and $n \in \mathbb{N}$,
    $$\mathbb{P}\left(\left|X^{(N,r)}(n, s_0) - p^*\right| \leq \delta\right) \geq 1 - \varepsilon - e^{-\lfloor Cn \rfloor}.$$

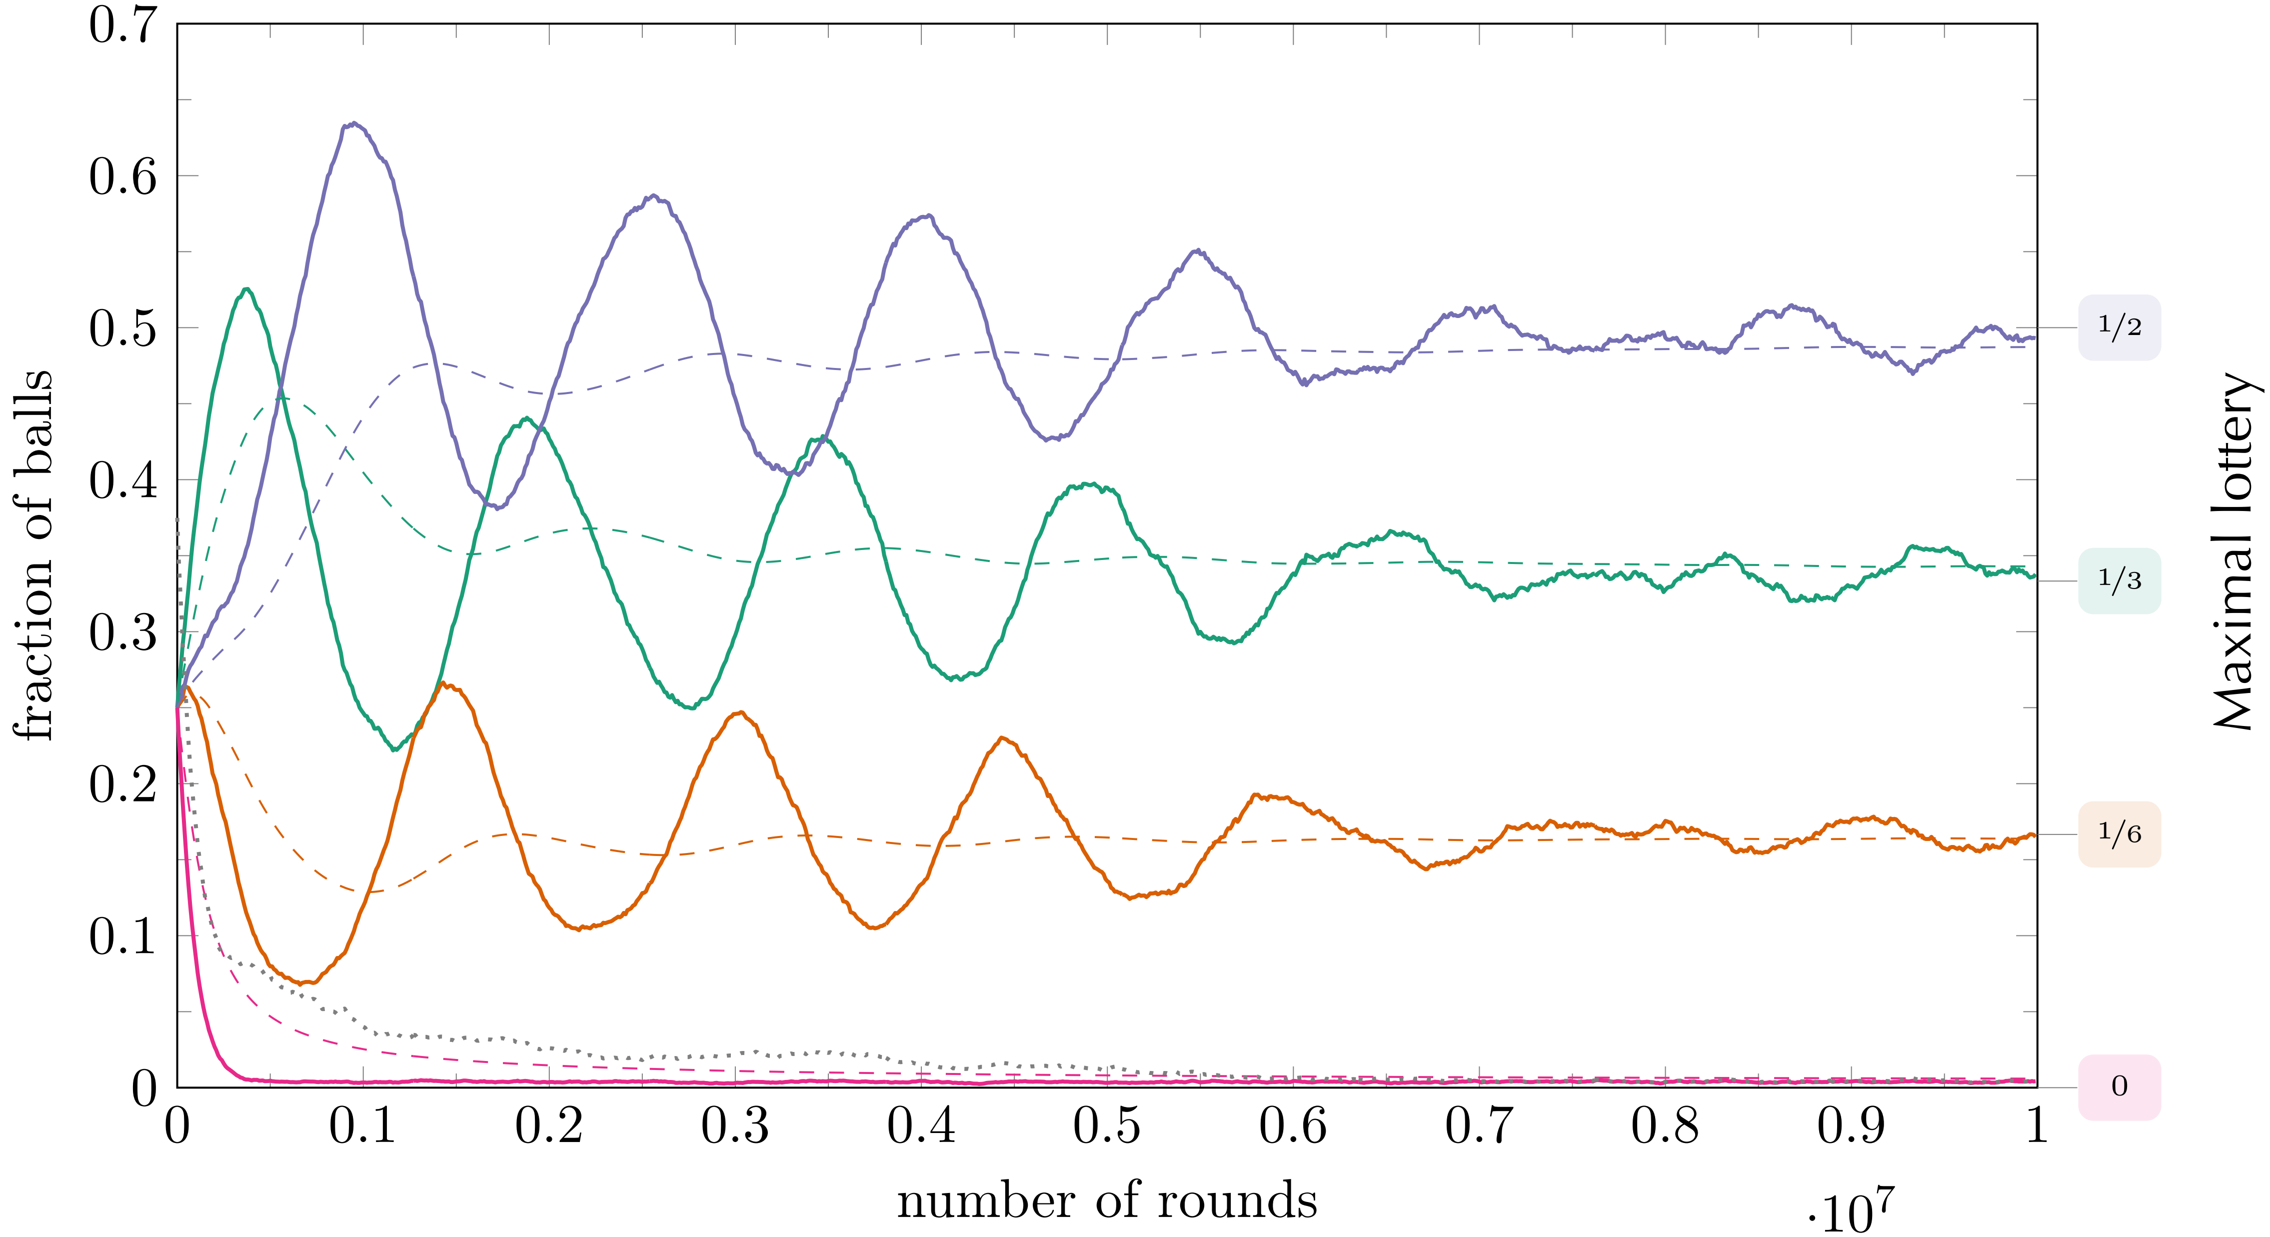

5	**4**	**3**
a	*c*	*d*
b	*a*	*b*
c	*b*	*c*
d	*d*	*a*

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

Germain Kreweras

Maximal Lotteries

Peter C. Fishburn

- Randomized voting rule proposed independently by Kreweras (1965) and Fishburn (1984).
- Let $M_{x,y}$ be the fraction of voters who prefer x to y.
- Matrix M induces a skew-symmetric matrix $\tilde{M} = M - M^{\mathsf{T}}$.
- A lottery p is maximal if $p^{\mathsf{T}}\tilde{M} \geq \mathbf{0}$.
    - mixed equilibrium strategy of the symmetric zero-sum game $\tilde{M}$
    - no other lottery q is preferred by an expected majority ($p^{\mathsf{T}}\tilde{M}\, q \geq 0$)
    - randomized Condorcet winner
    - almost always unique
        - e.g., for odd number of voters (Laffond et al., 1997)

Germain Kreweras

Maximal Lotteries

Peter C. Fishburn

4	3	5
a	b	c
b	c	a
c	a	b

$$M = \begin{array}{c} \\ a \\ b \\ c \end{array}\!\!\begin{array}{c} \begin{array}{ccc} a & b & c \end{array} \\ \begin{pmatrix} 0 & 3/4 & 1/3 \\ 1/4 & 0 & 7/12 \\ 2/3 & 5/12 & 0 \end{pmatrix} \end{array}$$

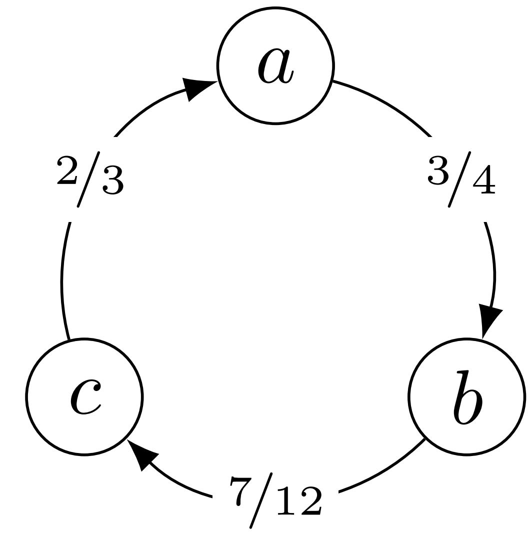

$$\underbrace{\begin{pmatrix} 1/6 \\ 1/3 \\ 1/2 \end{pmatrix}^{\mathsf{T}}}_{\text{maximal lottery}} \underbrace{\begin{pmatrix} 0 & 1/2 & -1/3 \\ -1/2 & 0 & 1/6 \\ 1/3 & -1/6 & 0 \end{pmatrix}}_{\tilde{M}=M-M^{\mathsf{T}}} = \begin{pmatrix} 0 & 0 & 0 \end{pmatrix} \geq \mathbf{0}$$

Stochastic Choice

$$\underbrace{M = \begin{array}{c c} & \begin{array}{ccc} a & b & c \end{array} \\ \begin{array}{c} a \\ b \\ c \end{array} & \begin{pmatrix} 0 & 1 & 1/6 \\ 0 & 0 & 2/3 \\ 5/6 & 1/3 & 0 \end{pmatrix} \end{array}}_{\text{comparison matrix}}$$

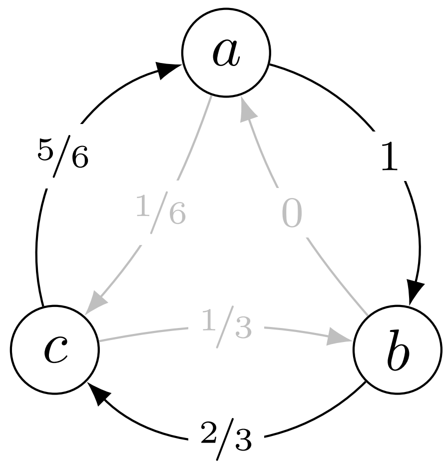

$$\underbrace{\begin{pmatrix} 1/6 \\ 1/3 \\ 1/2 \end{pmatrix}^{\intercal}}_{\text{maximal lottery}} \underbrace{\begin{pmatrix} 0 & 1 & -2/3 \\ -1 & 0 & 1/3 \\ 2/3 & -1/3 & 0 \end{pmatrix}}_{\tilde{M}=M-M^{\intercal}} = \begin{pmatrix} 0 & 0 & 0 \end{pmatrix} \geq \mathbf{0}$$

- Comparison matrices appear in various contexts and maximal lotteries have been repeatedly identified as attractive choice rules.
  - **Tournament Solutions**: *Bipartisan set* (Laffond et al., 1993), *Essential set* (Dutta & Laslier, 1999)
  - **Voting**: *Maximal lottery* (Fishburn, 1984), *Game theory procedure* (Felsenthal & Machover, 1992), *Game theory method* (Rivest and Shen, 2010)
  - **Matching Markets**: *Popular mixed matching* (Kavitha et al., 2011)
  - **Multi-Armed Bandits**: *von Neumann winner* (Dudík et al., 2015)
  - **Google DeepMind's AlphaStar**: *Nash averaging* (Balduzzi et al., 2018)

Desirable Properties

- A lottery remains maximal when removing unchosen alternatives or changing the dominance probabilities between such alternatives.
- A lottery that is maximal for two comparison matrices is also maximal for any convex combination of both matrices.
- The selection probability of an alternative is unaffected by cloning other alternatives.
- Classic social choice impossibilities have been turned into complete axiomatic characterizations of maximal lotteries, e.g.,
  - ➔ Brandl & B., *Arrovian Aggregation of Convex Preferences* (ECMA 2020)
  - ➔ Brandl et al., *Consistent Probabilistic Social Choice* (ECMA 2016)
  - ➔ Brandl et al., *Welfare Maximization Entices Participation* (GEB 2018)

$$\begin{array}{c} \\ a \\ b \\ c \end{array} \begin{array}{c} \begin{array}{ccc} a & b & c \end{array} \\ \begin{pmatrix} 0 & 2/3 & 1/3 \\ 1/3 & 0 & 2/3 \\ 2/3 & 1/3 & 0 \end{pmatrix} \end{array}$$

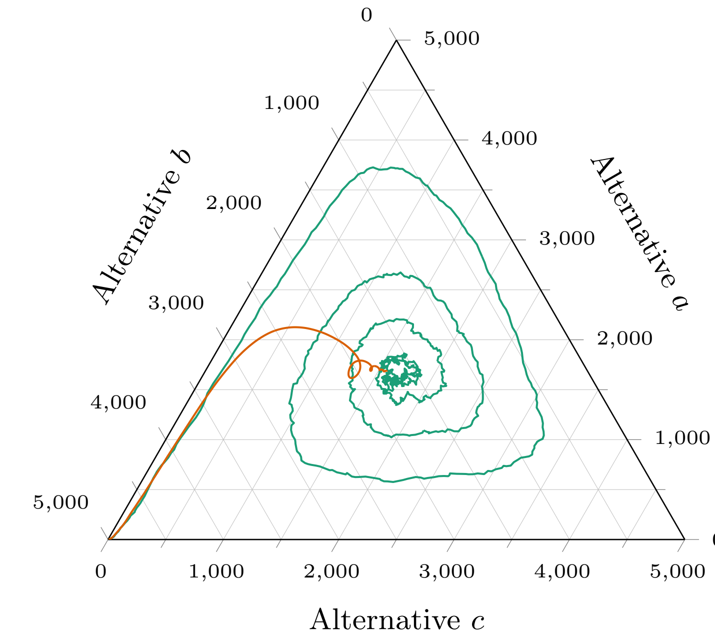

- Similar dynamic processes with equilibrium convergence
  - Population biology: coexistence of species
  - Quantum physics: condensation of bosons
  - Chemical kinetics: reactions of molecules
  - Plasma physics: scattering of plasmons
  - E.g., Allesina and Levine (*PNAS* 2011), Knebel et al. (*Nat Commun* 2015), Laslier & Laslier (*Ann Appl Probab* 2017), Grilli et al. (*Nature* 2017)
- Differences of our model and result
  - discrete (not continuous)
  - stochastic (not deterministic) interactions between pairs (not triples)
  - mutations
  - bound on sojourn time (rather than only convergence of time avg.)

Conclusion

- Advantages 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.