# A Natural Adaptive Process for Collective Decision-Making

Florian Brandl
Princeton University

Felix Brandt
Technische Universität München

Consider an urn filled with balls, each labeled with one of several possible collective decisions. Now, draw two balls from the urn, let a random voter pick her more preferred as the collective decision, relabel the losing ball with the collective decision, put both balls back into the urn, and repeat. In order to prevent the permanent disappearance of some types of balls, a randomly drawn ball is labeled with a random collective decision once in a while. We prove that the empirical distribution of collective decisions converges towards the outcome of a celebrated probabilistic voting rule proposed by Peter C. Fishburn (Rev. Econ. Stud., 51(4), 1984). The proposed procedure has analogues in nature studied in biology, physics, and chemistry. It is more flexible than traditional voting rules because it does not require a central authority, elicits very little information, and allows voters to arrive, leave, and change their preferences over time.

## 1. Introduction

The question of how to collectively select one of many alternatives based on the preferences of multiple agents has occupied great minds from various disciplines. Its formal study goes back to the Age of Enlightenment, in particular during the French Revolution, and the important contributions by Jean-Charles de Borda and Marie Jean Antoine Nicolas de Caritat, better known as the Marquis de Condorcet. Borda and Condorcet agreed that plurality rule—then and now the most common collective choice procedure—has serious shortcomings. This observation remains a point of consensus among social choice theorists and is largely due to the fact that plurality rule merely asks each voter for her most-preferred alternative (see, e.g., Brams and Fishburn, 2002; Laslier, 2011).[1] When eliciting more fine-grained preferences such as complete rankings over all alternatives from the voters, much more attractive choice procedures are available. As a matter of fact, since Arrow's (1951) seminal work, the standard assumption in social

[1]For example, plurality rule may select an alternative that an overwhelming majority of voters consider to be the worst of all alternatives.

choice theory is that preferences are given in the form of binary relations that satisfy completeness, transitivity, and often anti-symmetry. Despite a number of results which prove critical limitations of choice procedures for more than two alternatives (e.g., Arrow, 1951; Gibbard, 1973; Satterthwaite, 1975), there are many encouraging results (e.g. Young, 1974; Young and Levenglick, 1978; Brams and Fishburn, 1978; Laslier, 2000). In particular, when allowing for randomization between alternatives, some of the traditional limitations can be avoided and there are appealing choice procedures that stand out (Gibbard, 1977; Brandl et al., 2016; Brandl and Brandt, 2020).

The standard framework in social choice theory rests on a number of rigid assumptions that confine its applicability: there is a fixed set of voters, a fixed set of alternatives, and a single point in time when preferences are to be aggregated; all voters are able to rank-order all alternatives; there is a central authority that has access to all preferences, computes the outcome, and convinces voters of the outcome's correctness, etc. On top of that, computing the outcome of many attractive choice procedures is a demanding task that requires a computer, which can render the process less transparent to voters.[2]

In this paper, we propose a simple urn-based procedure that implements a celebrated choice procedure called *maximal lotteries* (Fishburn, 1984). Our goal is to devise a continuous process in which voters may arrive, leave, and change their preferences over time. Moreover, voters are never asked for their complete preference relations, but rather reveal minimal information about their preferences by choosing between two randomly drawn alternatives from time to time. No central voting authority is required. The process can be executed via a simple physical device: an urn filled with balls that allows for two primitive operations: *(i)* randomly sampling a ball and *(ii)* replacing a sampled ball of one kind with a ball of another kind.

More precisely, the process works as follows (see Figure 1). Consider an urn filled with balls that each carry the label of one alternative. The initial distribution of balls in the urn is irrelevant. In each round, a randomly selected voter will draw two balls from the urn at random. Say these two balls are labeled with alternatives 1 and 2, and the voter prefers 1 to 2. She will then change the label of the second ball to 1 and return both balls to the urn. Alternative 1 is declared the winner of this round. Once in a while, with some small probability $r$, which we call *mutation rate*, a randomly drawn ball is labeled with a random alternative.

We show that if the number of balls in the urn is sufficiently large, then the empirical distribution of the winners converges to a lottery close to a maximal lottery almost surely, that is, with probability 1. How far the limiting distribution will be from a maximal lottery depends on $r$. As $r$ goes to 0, the limiting distribution converges to a maximal lottery. We can, however, not set $r$ to 0 as then with probability 1, all alternatives except one will permanently disappear from the urn and the limiting distribution will be degenerate. Our proof not only shows convergence of the limiting distribution but also that the distribution of balls in the urn is close to maximal lottery most of the time.

[2]In some cases, computing the outcome was even shown to be NP-hard, i.e., the running time of all known algorithms for computing election winners increases exponentially in the number of alternatives (see, e.g., Bartholdi, III et al., 1989; Brandt et al., 2016b).

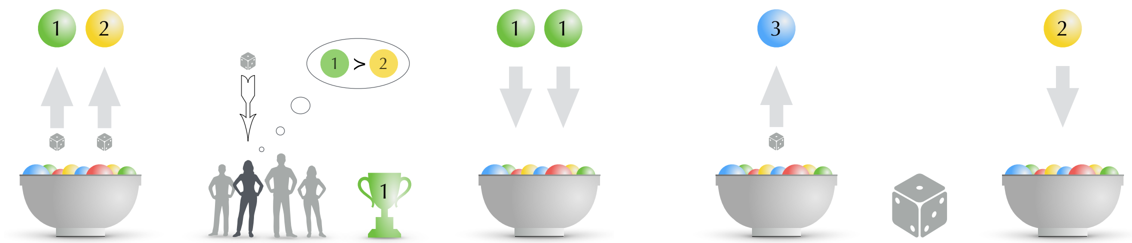

*(i)* Draw two balls, random voter picks winner, relabel losing ball.

*(ii)* Randomly relabel one drawn ball with small probability $r$ (mutation).

*(iii)* The distribution of winners converges to a maximal lottery w.r.t. the voters' preferences.

Figure 1: Illustration of one round of the urn process (*i* and *ii*) and the main result (*iii*).

Maximal lotteries are known to satisfy a number of desirable properties that are typically considered in social choice theory. For example, Condorcet winners (i.e., alternatives that defeat every other alternative in a pairwise majority comparison) will be selected with probability 1, and Condorcet losers (i.e., alternatives that are defeated in all pairwise majority comparisons) will never be selected. No group of voters benefits by abstaining from an election, removing losing alternatives does not affect maximal lotteries, and each alternative's probability is unaffected by cloning other alternatives. Maximal lotteries have been axiomatically characterized using Arrow's independence of irrelevant alternatives and Pareto efficiency (Brandl and Brandt, 2020) as well as population-consistency and composition-consistency (Brandl et al., 2016). The dynamic procedure described above implements maximal lotteries while providing

- myopic strategyproofness (in each round, a randomly selected voter chooses between two alternatives),
- minimal preference elicitation and thus increased privacy protection,
- verifiability realized via a simple physical procedure, and
- increased flexibility in the sense that agents may arrive, leave, and change their preferences over time; similarly, the set of alternatives can be modified while the process is running.

The axiomatic characterizations of maximal lotteries not only imply that maximal lotteries satisfy several desirable axioms, but also that any deviation from maximal lotteries leads to a violation of at least one of the axioms. Hence, a process that only guarantees an approximation of a maximal lottery will not enjoy the same axiomatic properties. However, rather than insisting on stringent axioms, one could relax them by only requiring them to hold in an approximate sense. For example, a natural notion of approximate Condorcet-consistency would require that a Condorcet winner receives probability close to 1 whenever one exists. Since the empirical distribution of winners according to our process is almost surely close to a maximal lottery and maximal lotteries are Condorcet-consistent, the process is approximately Condorcet-consistent in the above sense. More generally, many of the axioms that maximal lotteries satisfy such as population-consistency, composition-consistency, agenda-consistency, and efficiency also hold for approximate maximal lotteries. This follows from the fact that the correspondence returning the set of maximal lotteries depends continuously on the underlying preference profile (see Brandl et al., 2016).

Remarkably, dynamic processes similar to the process we describe here have recently been studied in population biology, quantum physics, chemical kinetics, and plasma physics to model phenomena such as the coexistence of species, the condensation of bosoms, the reactions of molecules, and the scattering of plasmons. In each of these cases, simple interactions between randomly sampled entities can be connected to equilibrium strategies in symmetric zero-sum games. In fact, maximal lotteries are precisely the mixed Nash equilibrium (or maximin) strategies of the symmetric two-player zero-sum game given by the pairwise majority margins of the voters' preferences. We discuss these relationships, including those to evolutionary game theory, in detail in Section 6.

An alternative interpretation of our result can be used to describe the formation of opinions. In this model, there is a population of agents, each of which entertains one of many possible opinions. Agents come together in random pairwise interactions, in which they try to convince each other of their opinion. The probabilities with which one opinion beats another are given as a square matrix and, with some small probability, an agent randomly changes her opinion. In other words, the agents correspond to the balls in the urn, the opinions correspond to the alternatives, and there are neither voters nor preference profiles as transition probabilities are given directly. Our main theorem then shows that, if the population is large enough, the distribution of opinions within the population is close to a maximal lottery of the probability matrix most of the time.

The remainder of the paper is structured as follows. After defining our model in Section 2, we state the main result (Theorem 1) and a rough proof sketch in Section 3. The full proof is given in the Appendix. In Section 4, we analyze the instructive special case of preference profiles that admit a Condorcet winner, which allows for a more elementary proof. As shown in Section 5, the urn process converges exponentially fast. In Section 6, we extensively discuss related work from various disciplines, including a continuous version of our main theorem (Theorem 2) that may be of independent interest.

## 2. The Model

Let $[d] = \{1, \dots, d\}$ be a set of alternatives and $\Delta$ the $d-1$-dimensional unit simplex in $\mathbb{R}^d$, that is, $\Delta = \{x \in \mathbb{R}^d_{\geq 0} \colon \sum_{i=1}^d x_i = 1\}$. We refer to elements of $\Delta$ as lotteries. Throughout the paper, for a vector $x \in \mathbb{R}^k$ for some $k$, $|x| = \sum_{l=1}^k |x_l|$ denotes its $L^1$-norm. For a finite set $S$, we write $\#S$ for the number of elements of $S$.

A *preference relation* $\succ$ is an asymmetric binary relation over $[d]$.[3] By $\mathcal{R}$ we denote

[3]Preferences need not be transitive. The definition of maximal lotteries and the urn process we describe

the set of all preference relations. Let $V$ be a finite set of voters. A *preference profile* $R \in \mathcal{R}^V$ specifies a preference relation for each voter. With each preference profile $R$, we can associate a matrix $M_R \in [0,1]^{d\times d}$ that states for each ordered pair of alternatives the fraction of voters who prefer the first to the second. That is, $M_R(i,j) = \#\{v \in V\colon i \succ_v j\}/\#V$. This matrix induces a skew-symmetric matrix $\tilde{M}_R = M_R - M_R^\intercal$, which we call the matrix of *majority margins*.[4]

## 2.1. Maximal Lotteries

A lottery $p \in \Delta$ is a *maximal lottery* for a profile $R$ if $\tilde{M}_R p \leq 0$. By $ML(R)$ we denote the set of all lotteries that are maximal for $R$.

**Example 1.** Consider, for example, three voters ($V = \{v_1, v_2, v_3\}$), three alternatives ($d = 3$), and a preference profile $R$ given by the following table (each column contains the preference ranking of the corresponding voter).

| $v_1$ | $v_2$ | $v_3$ |
|---|---|---|
| 1 | 1 | 2 |
| 2 | 3 | 3 |
| 3 | 2 | 1 |

Then,

$$M_R = \begin{pmatrix} 0 & 2/3 & 2/3 \\ 1/3 & 0 & 2/3 \\ 1/3 & 1/3 & 0 \end{pmatrix} \quad \text{and} \quad \tilde{M}_R = \begin{pmatrix} 0 & 1/3 & 1/3 \\ -1/3 & 0 & 1/3 \\ -1/3 & -1/3 & 0 \end{pmatrix}.$$

The set of maximal lotteries $ML(R) = \{(1,0,0)^\intercal\}$ only contains the degenerate lottery with probability 1 on the first alternative. This alternative is a *Condorcet winner*, i.e., an alternative that is preferred to every other alternative by some majority of voters.

## 2.2. Markov Chains

We need some basic concepts for Markov chains. Let $S$ be a finite set, called the state space. A *Markov chain* with state space $S$ is a sequence of random variables $\{X(n)\colon n \in \mathbb{N}\}$ with values in $S$ so that for all $n \in \mathbb{N}$ and states $s, s_0, \ldots, s_n \in S$,

$$\mathbb{P}\left(X(n+1) = s \mid X(n) = s_n\right) = \mathbb{P}\left(X(n+1) = s \mid X(n) = s_n, \ldots, X(0) = s_0\right)$$

The defining property of Markov chains is that the probability of transitioning to any state from time $n$ to time $n+1$ depends only on the state at time $n$. Conditional on the state at time $n$, it is independent of the states at times $0, \ldots, n-1$. The Markov chain is time-homogeneous if the probability on the left-hand side is independent of $n$. With

only depend on the pairwise majority margins.

[4]A matrix $M$ is skew-symmetric if $M = -M^\intercal$.

a time-homogeneous Markov chain, we can associate its *transition probability matrix* $P \in [0,1]^{S\times S}$ with

$$P(s,s') = \mathbb{P}\left(X(n+1) = s \mid X(n) = s'\right)$$

for all $s, s' \in S$. $P$ is a (row) stochastic matrix, which means that it has non-negative values each of its rows sums to 1. We will frequently write $X(n, s_0)$ for the random variable $X(n)$ conditioned on $X(0) = s_0 \in S$ and call $s_0$ the *initial state*. All Markov chains we consider will be time-homogeneous and have a finite state space.

Let $\{X(n) \colon n \in \mathbb{N}\}$ be a Markov chain with transition probability matrix $P$. The period of a state $s \in S$ is the greatest common divisor of the return times $\{n \in \mathbb{N} \colon (P^n)(s,s) > 0\}$. A Markov chain is *aperiodic* if every state has period 1. Aperiodicity requires that the return times of each state are not all multiples of the same prime. A Markov chain is *irreducible* if every state is reached from any other state with positive probability. That is, for any two states $s, s' \in S$, there is a positive integer $n$ so that $(P^n)(s,s') > 0$. If $\{X(n) \colon n \in \mathbb{N}\}$ is irreducible and aperiodic, it has a unique *stationary distribution* $\pi \in \Delta S$ so that

$$\pi^\intercal = \pi^\intercal P$$

Hence, $\pi$ is a left-eigenvector of the transition probability matrix $P$ for the eigenvalue 1.

## 2.3. The Urn Process

Consider an urn with $N \in \mathbb{N}$ balls, each labeled with some alternative. Viewing balls with the same label as indistinguishable, we can identify each state of the urn with an element of $S^{(N)} = \{s \in \mathbb{N}^d \colon \sum_{i=1}^d s_i = N\}$. Fix a *mutation rate* $r \in [0,1]$.

We are interested in a Markov chain with state space $S^{(N)}$, which can be informally described as follows. First, we flip a coin that has probability $1-r$ of landing heads. If the coin shows heads, we choose one voter $v \in V$ uniformly at random and ask the voter to draw two balls from the urn. Say these two balls are labeled with alternatives 1 and 2. If $1 \succ_v 2$, the label of the second ball is changed to 1. Likewise, if $2 \succ_v 1$, the first ball is relabeled with label 2. If both balls carry the same label, the labels remain unchanged. If the coin shows tails, we draw a single ball from the urn, relabel it with an alternative chosen uniformly at random, and put it back into the urn.

In order to formally capture this process, we define a transition probability matrix $P^{(N,r)}$ that specifies for every pair of states the probability that the distribution of the urn transitions from the first to the second. Denote by $e_i$ the $i$th unit vector in $\mathbb{N}^d$. For $s \in S^{(N)}$ and $i, j \in [d]$ with $s' = s + e_i - e_j \in S^{(N)}$, let

$$P^{(N,r)}(s,s') = \begin{cases} (1-r)\frac{s_i s_j}{\binom{N}{2}} M_R(i,j) + \frac{r}{d}\frac{s_j}{N} & \text{if } i \neq j \\ (1-r)\sum_{k=1}^d \frac{\binom{s_k}{2}}{\binom{N}{2}} + \frac{r}{d} & \text{if } i = j \end{cases}$$

be the probability of transitioning from $s$ to $s'$. For the remaining pairs of states $s, s' \in S^{(N)}$, let $P^{(N,r)}(s,s') = 0$. Then, $P^{(N,r)}$ has non-negative values and its rows sum to 1 so that it is a valid transition probability matrix. For an initial state $s_0 \in S^{(N)}$, we consider a Markov chain $\{X^{(N,r)}(n, s_0)\colon n \in \mathbb{N}\}$ with transition probability matrix $P^{(N,r)}$. The distribution of $X^{(N,r)}(n, s_0)$ over $S^{(N)}$ is given by the row of $\left(P^{(N,r)}\right)^n$ with index $s_0$. If $r > 0$, this Markov chain is irreducible and aperiodic. It corresponds to the urn process described above when the initial state of the urn is $s_0$.

Continuing Example 1, consider an urn containing $N = 5$ balls. Then, the transition probability matrix $P^{(N,r)}$ is an $\binom{3+5-1}{5} = 21$-dimensional square matrix. Let the mutation rate be $r = 0.1$ and the initial state $s_0 = (1, 2, 2)^\intercal$. Then, the probability that one of the balls of the second type is replaced with one of the first type is

$$P^{(5,0.1)}(s_0, (2, 1, 2)^\intercal) = 0.9 \cdot \frac{2}{10} \cdot \frac{2}{3} + 0.1 \cdot \frac{1}{3} \cdot \frac{2}{5} = \frac{2}{15} \sim 0.133.$$

## 3. The Result

We show that for any initial state $s_0 \in S^{(N)}$, the distribution of alternatives in the state $X^{(N,r)}(k, s_0)$ is close to a maximal lottery for all but a small fraction of rounds $k$ for some $r > 0$ provided the number of balls $N$ is large enough. More precisely, take any $\delta, \tau > 0$. Then we can find $N_0 \in \mathbb{N}$ and $r > 0$ so that for every $N \geq N_0$ and every initial state $s_0 \in S^{(N)}$, the following event has probability 1: the lower density of the $k \in \mathbb{N}$ for which $\frac{1}{N}X^{(N,r)}(k, s_0)$ is no more than $\delta$ away from $ML(R)$ is at least $1 - \tau$. For the formal statement, let $B_\delta(ML(R)) = \{p \in \Delta\colon \inf\{|p - q|\colon q \in ML(R)\} < \delta\}$ denote the set of lotteries with distance less than $\delta$ to some maximal lottery.[5]

**Theorem 1.** *Let $\delta, \tau > 0$. Then, there are $N_0 \in \mathbb{N}$ and $r > 0$ such that for all $N \geq N_0$ and $s_0 \in S^{(N)}$,*

$$\mathbb{P}\left(\liminf_{n\to\infty} \frac{1}{n} \# \left\{k \in [n]\colon \frac{1}{N}X^{(N,r)}(k, s_0) \in B_\delta\left(ML(R)\right)\right\} \geq 1 - \tau\right) = 1$$

We outline the main steps in the proof of Theorem 1. Consider the process $x^{(N,r)}(k, \frac{s_0}{N}) = \frac{1}{N}X^{(N,r)}(k, s_0)$ obtained by scaling to the discrete unit simplex $\Delta^{(N)} = \{p \in \Delta\colon Np \in \mathbb{N}_0^d\}$. For any $p \in \Delta^{(N)}$, we consider the expected value of $x^{(N,r)}(k + 1, p) - x^{(N,r)}(k, p)$, which, conditional on $x^{(N,r)}(k, p)$, is independent of $k$ since $x^{(N,r)}$ is time-homogeneous. These expected values induce a continuous function $f^{(N,r)}\colon \Delta \to \mathbb{R}^d$. For $r > 0$, $f^{(N,r)}$ has a unique zero $p^{(N,r)}$, which is close to the set of maximal lotteries $ML(R)$ if $r$ is small.

Consider now the following differential equation with $p \in \Delta$, $t \in \mathbb{R}_+$, and $y(\cdot, p)\colon \mathbb{R}_+ \to \Delta$.

$$\begin{aligned} \frac{d}{dt}y(t, p) &= f^{(N,r)}(y(t, p)) \\ y(0, p) &= p \end{aligned} \tag{1}$$

[5] Theorem 1 implies that the stationary distribution of $X^{(N,r)}$ assigns probability at least $1-\tau$ to states that are in a $\delta$-neighborhood of $ML(R)$. Conversely, this property of the stationary distribution implies Theorem 1 by the ergodic theorem for Markov chains. The proof does however not derive the above property of the stationary distribution as an intermediate step. It is only a by-product of the final result.

A solution to (1) is a *deterministic* dynamic process that can be interpreted as the *stochastic* process we consider with a continuum of balls. We show that the unique solution $y^{(N,r)}(\cdot,p)$ of (1) converges to $p^{(N,r)}$ for any initial state $p \in \Delta$ as $t$ goes to infinity and the convergence is uniform in $p$. This is done by showing that the entropy of $p^{(N,r)}$ relative to $y^{(N,r)}(t,p)$ decrease monotonically at a rate proportional to the square of the distance between $p^{(N,r)}$ and $y^{(N,r)}(t,p)$.

To relate the discrete-time process $x^{(N,r)}$ to the continuous-time process $y^{(N,r)}$, we extend the former to the real time axis by letting $\bar{x}^{(N,r)}(t,p) = x^{(N,r)}(k,p)$ for $t \in [\frac{k-1}{N}, \frac{k}{N})$. Given any $T > 0$, one can show that with probability close to 1, $\bar{x}^{(N,r)}$ *approximately* satisfies the integral equation corresponding to (1) for $t$ between 0 and $T$ and uniformly in $p \in \Delta^{(N)}$ if $N$ is large. Using Grönwall's inequality, we show that with probability close to 1, $\bar{x}^{(N,r)}(t,p)$ and $y^{(N,r)}(t,p)$ are close to each other for all $t$ from 0 to $T$. However, for $t$ larger than $T$, they may (and with probability 1 will) be arbitrarily far apart.

To deal with this, we partition the time axis into consecutive intervals of length $T$ and synchronize the continuous process with the discrete process at the beginning of each interval. More precisely, take any $\delta, \tau > 0$. Since $y^{(N,r)}(t,p)$ converges to $p^{(N,r)}$ as $t$ goes to infinity uniformly in $p$, we can find $T > 0$ such that $y^{(N,r)}(t,p)$ is no more than $\frac{\delta}{2}$ away from $p^{(N,r)}$ for all but possibly a $1 - \frac{\tau}{2}$ fraction of the interval $[0,T]$ for all $p$. Moreover, we can choose $N$ large enough so that with probability at least $1 - \frac{\tau}{2}$, the distance between $\bar{x}^{(N,r)}$ and $y^{(N,r)}$ is less than $\frac{\delta}{2}$ for all $t$ in an interval of length $T$ provided both processes start at the same point at the beginning of the interval. We chop up the time axis into intervals $[0,T]$, $[T,2T]$, .... On the interval $[(k-1)T, kT]$, we compare $\bar{x}^{(N,r)}(t,p) = \bar{x}^{(N,r)}(t-(k-1)T, \bar{x}_{k-1})$ to $y^{(N,r)}(t-(k-1)T, \bar{x}_{k-1})$, where $\bar{x}_{k-1} = x^{(N,r)}((k-1)T, p)$. That is, we reset $y^{(N,r)}$ to the position of $\bar{x}^{(N,r)}$ at the beginning of the interval. In those intervals where the distance between both processes is never more than $\frac{\delta}{2}$, $\bar{x}^{(N,r)}$ is no more than $\frac{\delta}{2} + \frac{\delta}{2} = \delta$ away from $p^{(N,r)}$ for all but a $\frac{\tau}{2}$ fraction of the interval. By the choice of $N$, those intervals make up at least a $\frac{\tau}{2}$ fraction of all intervals in expectation. Summing over all intervals, this is enough to conclude that $\bar{x}^{(N,r)}$ is no more than $\delta$ away from $p^{(N,r)}$ at least a $1-\tau$ fraction of the time. Since $p^{(N,r)}$ is close to $ML(R)$ when $r$ is small, we can get the same conclusion with $ML(R)$ in the place of $p^{(N,r)}$. Translating this statement back to $X^{(N,r)}$ gives Theorem 1.

Recall that the collective decision in each round is the winner of the pairwise comparison between the two drawn balls. Theorem 1 implies that the empirical distribution of the winners is close to a maximal lottery. For suppose the distribution of balls in the urn is $p \in \Delta^{(N)}$. Then the probability that $i \in [d]$ is the collective decision is

$$w_i = p_i \left( p_i + 2 \sum_{j \neq i} M_R(i,j) p_j \right) = p_i \left( p_i + \sum_{j \neq i} (\tilde{M}_R(i,j) + 1) p_j \right) = p_i \left( 1 + \sum_{j \neq i} \tilde{M}_R p_j \right)$$

where we used that $2M_R(i,j) = \tilde{M}_R(i,j) + 1$ and $\sum_{j \in [d]} p_j = 1$. Since $\tilde{M}_R$ is skew-symmetric, $w \in \Delta$. If $p \in B_\delta(ML(R))$, then $(\tilde{M}_R p)_i \leq \delta$ for all $i \in [d]$. Hence, $w_i \in [p_i - \delta, p_i + \delta]$ for all $i$, so that $w \in B_{\delta d}(ML(R))$. For every $\delta' > 0$, choosing

$\delta = \tau = \frac{\delta'}{2d}$ in Theorem 1 thus shows that the empirical distribution of collective decisions is almost surely in a ball of radius $\delta'$ around $ML(R)$.

Another straightforward corollary of Theorem 1 is that the temporal average of the $x^{(N,r)}(k, s_0)$ is almost surely close to $ML(R)$ for some small $r$ provided that $N$ is large enough. We define the temporal averages of the $x^{(N,r)}(k, s_0)$.

$$z^{(N,r)}(n, s_0) = \frac{1}{n} \cdot \sum_{k=0}^{n-1} x^{(N,r)}(k, s_0).$$

By the ergodic theorem for Markov chains, $z^{(N,r)}$ converges almost surely to the stationary distribution $\pi \in \Delta(\Delta^{(N)})$ of the Markov chain $\{x^{(N,r)}(n, s_0) \colon n \in \mathbb{N}_0\}$. Theorem 1 shows that $x^{(N,r)}$ is almost surely in a ball of radius $\delta$ around $ML(R)$ for all but a $\tau$-fraction of rounds. Hence, by choosing $\tau \leq \frac{\delta}{2}$ and using that $|\cdot|$ is bounded by 2 on $\Delta^{(N)}$, we get that almost surely, the limit $\lim_{n\to\infty} z^{(N,r)}(n, s_0)$ exists and lies in a ball of radius $2\delta$ around $ML(R)$. In particular, $z^{(N,r)}(n, s_0)$ is almost surely in $B_\delta(ML(R))$ for all but a finite number of $n$.

**Corollary 1.** *Let $\delta > 0$. Then, there are $N_0 \in \mathbb{N}$ and $r > 0$ such that for all $N \geq N_0$ and $s_0 \in S^{(N)}$,*

$$\mathbb{P}\left(\lim_{n\to\infty} z^{(N,r)}(n, s_0) \in B_\delta(ML(R))\right) = 1$$

Before illustrating these results via examples, we briefly discuss possible variations of the urn process.

**Remark 1.** It is not necessary to synchronize the drawing of winners with the process in which balls are replaced (via comparisons and mutations). For example, it may be more practical to only draw a winner after a certain number of rounds of has passed.

**Remark 2.** The results still hold if we require that the probability of a random mutation from one alternative to another is different for different pairs of alternatives. For example, we could ask that the mutation rate for the pair $(i, j)$ is proportional to the number $M_R(i, j)$ of agents who prefer $i$ to $j$. The proof can be adapted at the expense of more book-keeping.

**Remark 3.** Rather than letting only a single voter decide on the pairwise comparison between the two randomly drawn balls, it is possible to ask all voters which alternative they prefer and replace the alternative which is less preferred by a majority of voters. This variant is equivalent to the original process for a single voter with intransitive preferences (given by the majority relation of the entire population of voters) and converges to a so-called C1 maximal lottery of the preference profile (see Brandl et al., 2021, for a comparison of maximal lottery schemes). Randomized voting rules based on Markov chains defined via the pairwise majority relation have been studied in the literature on tournament solutions (Laslier, 1997; Brandt et al., 2016a).

**Remark 4.** When the initial distribution of balls in the urn is uniform and remains fixed (i.e., no balls are replaced over time), then the empirical distribution of winners converges to the lottery returned by the proportional Borda rule (see, e.g., Barberà, 1979; Heckelman, 2003; Brandt, 2017). When adding a new ball labeled with the winning alternative rather than replacing the losing one (i.e., the number of balls increases over time), neither the relative distribution in the urn nor the temporal average converges (see Section 6).

Figure 2 (left) shows a simulation of the urn process for the preference profile and corresponding majority margin matrix given in Example 1. The urn process corresponds to a random walk within the shown triangle starting from the center (an almost uniform distribution). The first alternative in this profile is a Condorcet winner. In round 205, 48 of the 50 balls are labeled with the Condorcet winner (and hence the distance to the maximal lottery falls below 0.1) for the first time. While the distribution in the urn moves away from the Condorcet winner for short periods of time, it remains within a distance of 0.1 for 683 of the remaining 794 rounds. The path is tilted to the left because a majority of voters prefer alternative 2 to alternative 3. Note that the process only depends on the majority margins and is thus independent of the number of voters. Hence, if there are three million—rather than three—voters and the preferences are equally distributed among the columns shown in Example 1, the process could turn out exactly as shown in Figure 2. In particular, the overwhelming majority of voters would never be queried for their preferences.

We now give two other examples, for which the unique maximal lottery is not degenerate.

**Example 2.** Consider three voters ($V = \{v_1, v_2, v_3\}$), three alternatives ($d = 3$), and the following preference profile $R$, known as the Condorcet cycle or Condorcet paradox.

| $v_1$ | $v_2$ | $v_3$ |
|---|---|---|
| 1 | 2 | 3 |
| 2 | 3 | 1 |
| 3 | 1 | 2 |

Then,

$$M_R = \begin{pmatrix} 0 & 2/3 & 1/3 \\ 1/3 & 0 & 2/3 \\ 2/3 & 1/3 & 0 \end{pmatrix} \quad \text{and} \quad \tilde{M}_R = \begin{pmatrix} 0 & 1/3 & -1/3 \\ -1/3 & 0 & 1/3 \\ 1/3 & -1/3 & 0 \end{pmatrix}.$$

The set of maximal lotteries $ML(R) = \{(1/3, 1/3, 1/3)\}$ consists of the uniform lottery over the three alternatives. A simulation of an urn process for this profile is given in Figure 2 (right). This time, the initial distribution is degenerate with all balls being of type 2. It can be seen how the distribution of balls in the urn closes in on the maximal lottery and remains in its neighborhood most of the time while the temporal average converges to the maximal lottery.

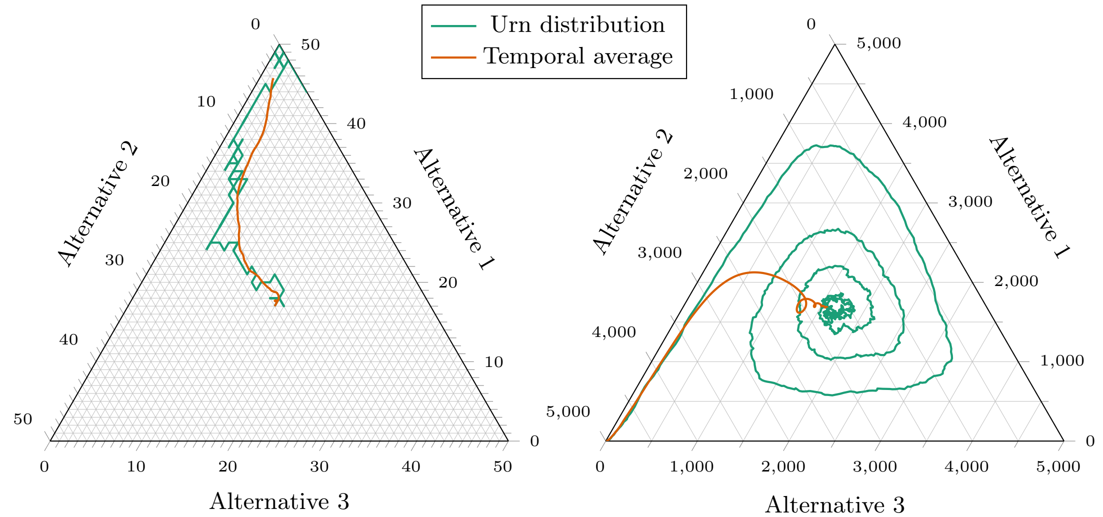


Figure 2: Simulations of the urn process.
The left diagram shows the urn process for the profile given in Example 1 using an urn with $N = 50$ balls for 1,000 rounds and mutation rate $r = 0.02$, starting from an almost uniform distribution. Each intersection of the grid lines corresponds to a configuration of the urn. The right diagram shows the urn process for the profile given in Example 2 using an urn with $N = 5,000$ balls for 500,000 rounds and mutation rate $r = 0.04$, starting from the degenerate distribution in which all balls are labeled with Alternative 2.

**Example 3.** Consider the following preference profile $R$ with 12 voters and 4 alternatives.

| $v_1, \ldots, v_5$ | $v_6, \ldots, v_9$ | $v_{10}, \ldots, v_{12}$ |
|---|---|---|
| 1 | 3 | 4 |
| 2 | 1 | 2 |
| 3 | 2 | 3 |
| 4 | 4 | 1 |

Then,

$$\tilde{M}_R = \begin{pmatrix} 0 & 1/3 & -1/9 & 1/3 \\ -1/3 & 0 & 2/9 & 1/3 \\ 1/9 & -2/9 & 0 & 1/3 \\ -1/3 & -1/3 & -1/3 & 0 \end{pmatrix}.$$

The set of maximal lotteries $ML(R) = \{(1/3, 1/6, 1/2, 0)\}$ consists of a single lottery, which is supported on the first three alternatives. A simulation of an urn process for this profile starting from the uniform distribution is given in Figure 3. The figure shows the distribution in the urn, the temporal average of urn distributions, and the difference of the urn distribution and the maximal lottery in terms of the relative entropy.[6]

[6] We use the relative entropy (rather than the distance $|p - q|$) to measure how much the distribution in the urn diverges from the maximal lottery since the proof of Theorem 1 shows that the entropy of the maximal lottery relative to the continuous approximation of the discrete process converges monotonically to 0.

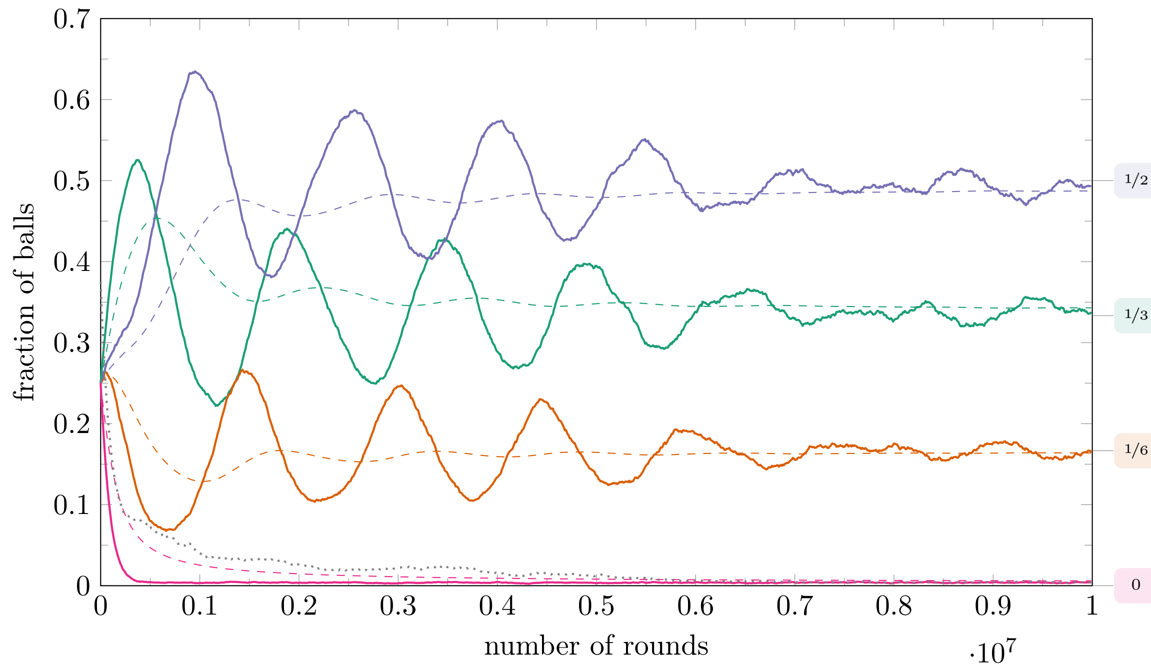


Figure 3: Simulation of the urn process for the profile in Example 3 on an urn with $N = 50,000$ balls for $10^7$ rounds and mutation rate $r = 0.01$. The solid lines show the fraction of balls in the urn. The dashed lines show the temporal average of the fraction of balls in the urn. The unique maximal lottery is $p = (1/3, 1/6, 1/2, 0)$. The dotted line shows relative entropy $D(p \mid q) = \sum_{i\in[d]} p_i \log(\frac{p_i}{q_i})$ of $p$ with respect to the distribution in the urn $q$.

## 4. The Case of a Condorcet Winner

We give an elementary proof of Theorem 1 for profiles that admit a Condorcet winner. In that case, we can directly analyze the stationary distribution $\pi \in \Delta(\Delta^{(N)})$ of the Markov chain induced by the urn process (which exists and is unique if $r > 0$). This allows us to give a concrete lower bound on the number of balls $N$ required so that $x^{(N,r)}$ is no more than $\delta > 0$ away from the maximal lottery (the lottery with probability 1 on the Condorcet winner) for at least a $1 - \tau$ fraction of rounds as a function of $\delta$ and $\tau$.

Let $M = M_R$ be the majority matrix of a profile $R$ with Condorcet winner $i \in [d]$. Hence, $M_{ij} > \frac{1}{2}$ for all $j \in [d]\setminus\{i\}$. Let $\alpha = \min\{M_{ij} : j \in [d]\setminus\{i\}\} - \frac{1}{2}$. We slice up $\Delta^{(N)}$ into the level sets of $i$. For $k \in \{0, \dots, N\}$, let $S_k = \{p \in \Delta^{(N)} : p_i = \frac{k}{N}\}$ be the states corresponding to distributions with $k$ of the $N$ balls of type $i$. Then $\sigma_k = \sum_{p\in S_k} \pi(p)$ is the limit probability that the urn is in a state in $S_k$ as the number of rounds goes to infinity. We want to show that if $r$ is sufficiently small and $N$ sufficiently large, $\pi$ has most of the probability on states in $S_k$ with $k$ close to $N$.

For 4 alternatives, one can illustrate the ensuing argument as follows. The set of states $\Delta^{(N)}$ corresponds to rooms in a tetrahedral-shaped pyramid. The rooms on the $k$th floor correspond to $S_k$, so that the tip of the pyramid is the state where all balls are of type $i$. The urn process is a random walk through the pyramid, moving from one room to an adjacent one (which could be on the same floor, the floor below, or the floor above). With the exception of few floors close to the tip, the probability of going up is always larger than the probability of going down. It is then intuitively clear that if the pyramid is large enough, one should expect to find the random walk close to the tip of the pyramid most of the time.

Recall that $P^{(N,r)}(p,q)$ is the probability of transitioning from state $p$ to state $q$. Since $\pi$ is a stationary distribution, we have $\pi^\intercal P^{(N,r)} = \pi^\intercal$. Consider any partition of $\Delta^{(N)}$ into two sets. For the stationary distribution, the probability of transitioning from the first set of the second is equal to the probability of transitioning from the second set to the first since the probabilities of both sets of conserved. Apply this to the sets $\bigcup_{l=0}^{k-1} S_l$ and $\bigcup_{l=k}^{N} S_l$ for $k \in [N]$ and notice that the only transitions between the two sets with positive probability are from $S^{k-1}$ to $S^k$ and vice versa. We get

$$\sum_{p\in S_{k-1}} \pi(p) \sum_{q\in S_k} P^{(N,r)}(p,q) = \sum_{p\in S_k} \pi(p) \sum_{q\in S_{k-1}} P^{(N,r)}(p,q). \tag{2}$$

That is, the probability of being in a state in $S_{k-1}$ and transitioning to a state in $S_k$ equals the probability of being in a state in $S_k$ and transitioning to a state in $S_{k-1}$.

Now observe that for $p \in S_k$, $k \in \{0,\dots,N-1\}$, we have

$$\sum_{q\in S_{k+1}} P^{(N,r)}(p,q) \geq (1-r)\frac{k(N-k)}{\binom{N}{2}}\left(\frac{1}{2}+\alpha\right) + \frac{r}{d}\frac{N-k}{N} =: u_k$$

where the left hand side is the probability of replacing a ball of type other than $i$ by one of type $i$ in state $p \in S_k$ (moving up one floor in the pyramid). Similarly, we find that for $p \in S_k$, $k \in \{1,\dots,N\}$, we have

$$\sum_{q\in S_{k-1}} P^{(N,r)}(p,q) \leq (1-r)\frac{k(N-k)}{\binom{N}{2}}\left(\frac{1}{2}-\alpha\right) + r\frac{d-1}{d}\frac{k}{N} =: d_k$$

for the probability of replacing a ball of type $i$ by one of type other than $i$ in state $p \in S_k$ (moving down one floor in the pyramid). Plugging this into (2), we get

$$\sigma_{k-1}u_{k-1} \leq \sigma_k d_k. \tag{3}$$

All terms in (3) are strictly positive if $r > 0$.

Let $N$ be so that $\frac{r}{Nd} \geq \frac{1-r}{\binom{N}{2}}$ (we choose $r > 0$ later). Then,

$$u_k \geq (1-r)\frac{k(N-k)}{\binom{N}{2}}\left(\frac{1}{2}+\alpha\right) + (1-r)\frac{N-k}{\binom{N}{2}} \geq (1-r)\frac{(k+1)(N-k-1)}{\binom{N}{2}}\left(\frac{1}{2}+\alpha\right)$$

where the last inequality uses $1 \geq \frac{1}{2}+\alpha$. Similarly, we find that for $r \leq \frac{1}{d}$ and $k \leq N\left(1-\frac{r}{\alpha}\right)$,

$$d_k \leq (1-r)\frac{k(N-k)}{\binom{N}{2}}\frac{1-\alpha}{2}$$

Hence, with this bound on $k$, we have

$$\frac{d_k}{u_{k-1}} \leq \frac{1-\alpha}{2\left(\frac{1}{2}+\alpha\right)} = \frac{1-\alpha}{1+2\alpha} =: \beta.$$

Thus, by (3), $\frac{\sigma_{k-1}}{\sigma_k} \leq \beta < 1$. We have shown that the cumulative probability $\sigma_k$ of the states $S_k$ decreases at least as fast the geometric series with parameter $\beta$ from some $k$ (close to $N$) downwards.

The maximal lottery for $R$ is the degenerate lottery with probability 1 on $i$. For given $\delta, \tau > 0$, we are aiming for a lower bound on $N$ so that the probability on states with at least $1-\delta$ fraction of balls of type $i$ in the stationary distribution $\pi$ is at least $1-\tau$. That is,

$$\sum_{k=\lceil N(1-\delta)\rceil}^{N} \sigma_k \geq 1-\tau.$$

First observe that

$$\sum_{k \geq k_0} \beta^k = \beta^{k_0} \frac{1}{1-\beta} \leq \tau \tag{4}$$

for $k_0 \geq \frac{\log(\tau(1-\beta))}{\log \beta}$. For our bound, $N$ needs to be large enough so that there are at least $k_0$ integers in the interval $\{\lceil (1-\delta)\, N \rceil, \ldots, \lfloor \left(1-\frac{r}{\alpha}\right) N \rfloor\}$. The probability on states in $S_k$ with $k < (1-\delta)\, N$ will then be below $\tau$ by (4) and the choice of $k_0$ (since the bound on $d_k$ assumes that $k \leq N(1-\frac{r}{\alpha})$). Choosing $r \leq \frac{\alpha\delta}{2}$ and

$$N \geq \frac{k_0}{\delta - \frac{r}{\alpha}} \geq \frac{1}{\delta} \left\lceil \frac{\log\left(\tau(1-\beta)\right)}{\log \beta} \right\rceil$$

achieves this.

In Example 1, there are three alternatives and three voters. Alternative 1 is a Condorcet winner as it is preferred to every other alternative by two of the three voters ($\alpha = \frac{2}{3} - \frac{1}{2} = \frac{1}{6}$, $\beta = \frac{5}{8}$). Suppose we want that at least 90% of the balls in the urn are of type 1 in at least 90% of rounds ($\delta = 0.1$, $\tau = 0.1$). Choosing $r = \frac{\alpha\delta}{2} = \frac{1}{120}$, we need $N \geq 140$ balls in the urn.

**Remark 5.** The calculations above suggest that when a Condorcet winner exists, a reasonable choice for simulations is $N \geq -\frac{1}{\delta}\log(\tau)$ and $\frac{1}{N} \leq r \leq \delta$.

**Remark 6.** If $r$ is too small compared to $N$, it will in general not be the case that the distribution in the urn is close to a maximal lottery for most rounds. For any long enough time interval, the distribution in the urn will for all $r$ with high probability degenerate within the interval, that is, it will only contain balls of one type. If $r$ is very small, it will stay in a degenerate state for a long time (compared to the chosen interval) with high probability. When the process leaves the degenerate state, the same will repeat itself (possibly with a different degenerate state), so that the process spends most rounds in a degenerate state.

## 5. Rate of Convergence

Recall that $\{x^{(N,r)}(n) \colon n \in \mathbb{N}_0\}$ is an irreducible and aperiodic Markov chain if $r > 0$. It is well-known (see, e.g., Levin et al., 2009, Theorem 4.9) that any such Markov chain converges to its unique stationary distribution $\pi \in \Delta(S^{(N)})$ in the sense that there are constants $C, \alpha > 0$ such that for all $n \geq 0$,

$$\max_{p \in \Delta^{(N)}} \left| \left(P^{(N,r)}\right)^n (p, \cdot) - \pi \right| \leq C\alpha^n \tag{5}$$

This means that the distance between the stationary distribution $\pi$ and the expected distribution at time $n$ decreases at least exponentially in $n$ uniformly over all initial states $p \in \Delta^{(N)}$.

Theorem 1 shows that $x^{(N,r)}$ is almost surely close to a maximal lottery in $ML(R)$ for all but a small fraction of iterations. It is a limit statement since it talks about the fraction of iterations for which some event occurs as the number of iterations goes to infinity. We now give a lower bound on the probability that $x^{(N,r)}$ is close to a maximal lottery in all but a small fraction of iterations as a function of the number of iterations. That is, for given $\delta, \tau > 0$, what is the probability that $x^{(N,r)}$ is no more than $\delta$ away from $ML(R)$ in all but a $1 - \tau$ fraction of the first $n$ iterations irrespective of the initial state?

By Theorem 1, we can choose $N$ and $r$ so that $\pi$ assigns probability at least $1 - \tau^2$ to states that are no more than $\delta$ away from a maximal lottery.

$$\sum_{p \in B_\delta(ML(R))} \pi(p) \geq 1 - \tau^2 \tag{6}$$

Let $C, \alpha$ be so that (5) holds and fix $p \in \Delta^{(N)}$. (The constants $C$ and $\alpha$ depend on $N$ and $r$ and, thus, on $\delta$ and $\tau$.) Let $q_n$ be the probability that a realization of $x^{(N,r)}$ is more than $\delta$ away from every maximal lottery in more than $n\tau$ of the first $n$ iterations when the initial state is $p$, i.e.,

$$q_n = \mathbb{P}\left( \# \left\{ k \in \{0, \dots, n-1\} \colon x^{(N,r)}(k,p) \notin B_\delta(ML(R)) \right\} \geq n\tau \right).$$

Hence, the probability that $x^{(N,r)}(k)$ is not in $B_\delta(ML(R))$ averaged over the first $n$ iterations is at least $\tau q_n$:

$$\tau q_n \leq \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{P}\left( x^{(N,r)}(k,p) \notin B_\delta(ML(R)) \right).$$

For any initial state $p$, the expected distribution of $x^{(N,r)}$ after $k$ iterations cannot have "too much " probability on states that are more than $\delta$ away from every maximal lottery since this would contradict (5). More precisely,

$$\mathbb{P}\left( x^{(N,r)}(k,p) \notin B_\delta(ML(R)) \right) \leq \left| \left(P^{(N,r)}\right)^k (p, \cdot) - \pi \right| + \sum_{p \notin B_\delta(ML(R))} \pi(p)$$

$$\leq C\alpha^k + \tau^2$$

The above bound is independent of $p$. Together, we get

$$q_n \leq \frac{1}{n\tau} \sum_{k=0}^{n-1} \left( C\alpha^k + \tau^2 \right) \leq \frac{C}{(1-\alpha)n\tau} + \tau$$

Thus, the upper bound for the probability of a "bad" realization after $n$ iterations approaches $\tau$ like $\frac{1}{n}$. The summand $\tau$ is arbitrary in the sense that requiring a different bound in (6) would give another summand. A smaller summand possibly comes at the cost of a larger $N$. The dependence of $q_n$ on $\delta$ is hidden in the fact that $C$ and $\alpha$ depend on $\delta$ and $\tau$. This estimate of $q_n$ is crude since it does not fully leverage the fact that our process is Markov. With more tailored methods, one can obtain the same asymptotic bound as above without the summand $\tau$.

# 6. Discussion of Related Work

Since the urn process we describe only depends on the transition probabilities induced by $\tilde{M}_R$ and $r$, it is applicable to various problems unrelated to collective decision-making.

## 6.1. Population dynamics

Our urn process is related to the *replicator equation* in population dynamics and evolutionary game theory (see, e.g., Taylor and Jonker, 1978; Schuster and Sigmund, 1983; Hofbauer and Sigmund, 1998). In its basic form, it states that the change in the relative frequency of a species equals the relative fitness of the species (that is, its fitness relative to the entire population) minus the change in the size of the entire population. When the fitness depends linearly on the relative frequencies of the species and the population size is constant, this is the differential equation (7) below with $r = 0$. The models in this stream of research are typically not discrete, which means that the distribution of species changes *continuously* with time and space

**Allesina and Levine (2011)** study the competition and coexistence of species in nature via a mathematical that is similar to our urn process. There is a fixed finite number of individuals, each of which is assigned to some species at random. In each round, two randomly selected individuals interact. The superior species will replace the individual of the inferior. Which species is superior to which species is given in the form of a tournament graph, which can be represented by a deterministic dominance matrix. Interestingly, these tournaments are sampled from distributions that are obtained via multiple rankings of the species called "limiting factors", similar to the preferences of voters. Simulations with large populations (e.g., 25,000 individuals) then show that the relative frequencies of the species oscillate around the equilibrium strategy of the dominance matrix. However, this phenomenon is an artifact of the population size and the limited time horizon. In the long run, as mentioned in Section 1, all species but one will become extinct with probability 1.

**Knebel et al. (2015)** study a similar process (including random mutations) in the context of quantum physics. Here, balls in the urn model correspond to bosons and alternatives to quantum states. The distribution of quantum states determines which states are condensates and are thus observed macroscopically. Since the number of particles in such systems is typically large, they focus on a process with a continuum of particles as described in Section 3 (which is deterministic). Leveraging a classic result from evolutionary game theory (Hofbauer and Sigmund, 1998, Theorem 5.2.3), they show that the *temporal average* of this process converges to an equilibrium strategy (i.e., a maximal lottery) of the zero-sum game induced by the transition probabilities between quantum states. All states with probability zero in the equilibrium strategy are depleted; the fractions of the remaining states are bounded away from 0 for all times. Knebel et al. neglect mutations for the continuous process, which may cause the process to cycle around the equilibrium strategy without converging to it. Part of our proof of Theorem 1 shows that the continuous process with mutations does converge (and not only its temporal average). Knebel et al. (2015, Supplementary Note 1) argue that the discrete process with mutations is well-approximated by the continuous process if the number of particles is large and mutations become vanishingly unlikely. Hence, they conclude that the temporal average of the discrete process converges to an equilibrium strategy, which is in the spirit of Corollary 1. Our understanding is that their arguments are heuristic and not intended to provide a rigorous derivation of this result. In particular, the arguments do not seem to use that mutations happen with non-zero probability. Without mutations, however, the discrete process almost surely enters a state with a degenerate distribution.

In earlier work, Knebel et al. (2013) have connected the survival and extinction of states to the Pfaffian of the transition matrix. This is reminiscent of a statement by Kaplansky (1995) about the support of equilibrium strategies in symmetric zero-sum games.Reichenbach et al. (2006) study the extinction probabilities for three states with cyclical dominance ("rock-paper-scissors") for finite populations.

**Laslier and Laslier (2017)** consider a discrete urn process that is similar to ours, but in which the number of balls in the urn increases over time. Two balls are drawn at random and a deterministic dominance matrix specifies which alternative wins against which alternative (this could be seen as a single voter with possibly intransitive preferences in our model). Rather than replacing the losing ball, a new ball of the same type as the winning ball is added to the urn. They show that the distribution in the urn does not converge unless one alternative beats all alternatives (which corresponds to the Condorcet winner case). However, the fraction of alternatives not contained in the support of the maximal lottery of the dominance matrix goes to zero. They then consider a modified process, in which three balls are drawn from the urn. Whenever one of three balls beats both other balls, a new ball of the same type is added to the urn. Otherwise, one of the three types is chosen at random and a ball of that type is added. Their main result is that, for this modified process, the distribution in the urn converges towards the (unique) maximal lottery of the dominance matrix. Laslier and

| | Model | Interaction | Mutations | Pop. Size | Convergence |
|---|---|---|---|---|---|
| Allesina et al. (2011) | discrete | pairs, det. | no | fixed | —[a] |
| Knebel et al. (2015) | cont. | pairs, det. | no[b] | fixed | temp. avg. |
| Laslier et al. (2017) | discrete | pairs, det. | no | increasing | supp. of distr. |
| Laslier et al. (2017) | discrete | triples, det. | no | increasing | distr. |
| Grilli et al. (2017) | cont. | triples, stoch. | no | fixed | distr. |
| Theorem 1 | discrete | pairs, stoch. | yes | fixed | frequ. of approx. distr. |
| Corollary 1 | discrete | pairs, stoch. | yes | fixed | temp. avg. |
| Theorem 2 | cont. | pairs, det. | yes | fixed | distr. |

Table 1: Comparison of related models and results.
*a*: In simulations, Allesina and Levine (2011) observe that the temporal average of their process comes close to a maximal lottery after a finite number of rounds. However, when the process is run long enough, the distribution will with probability 1 degenerate since there are no mutations.
*b*: While Knebel et al. (2015) consider a discrete process with mutations, the continuous process they study has no mutations.

Laslier neither consider the empirical distribution of winners nor the temporal average of the distribution in the urn. For the process with two drawn balls, it can be shown that not even the temporal average converges. Since the number of balls in the urn increases, convergence is generally very slow.

**Grilli et al. (2017)** consider a dynamical process in population biology to explain the stable coexistence of multiple species. Based on Laslier and Laslier's findings, Grilli et al. adapt the replicator equation to interactions of triples of individuals: in each round, a randomly chosen individual dies; it is replaced by the winner of a comparison between three randomly selected individuals, where the winner is determined as in Laslier and Laslier's process with three drawn balls based on a stochastic dominance matrix between species. Hence, the number of individuals remains constant. They show that with a continuum of individuals, this process converges to an equilibrium strategy of the zero-sum game corresponding to the dominance matrix. For a finite number of individuals, permanent coexistence of multiple species is a probability zero event. However, they argue that interactions of three or more individuals can prolong coexistence compared to pairwise interactions.

**Comparison.** Table 1 shows the key differences between the above mentioned results. In comparison, the main contribution of our work is that we are able to show for a discrete (rather than continuous) process based on stochastic (rather than discrete) interactions between pairs (rather than triples) that the actual distribution in the urn is close to a maximal lottery most of the time (rather than convergence of the temporal average). Methodologically, the approach we take to cope with the discrete process is related to that of Benaim and Weibull (2003), who study more general population processes in

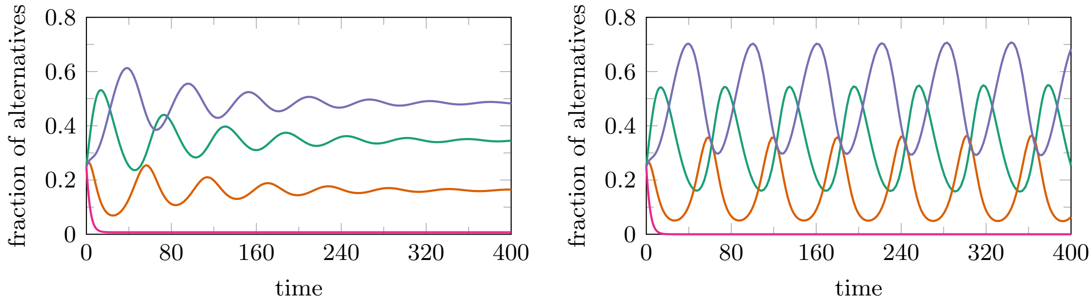


Figure 4: The continuous deterministic process $y(t)$ solving Equation (7) for the profile in Example 3 with $r = 0.01$ on the left and $r = 0$ on the right. For strictly positive $r$, $y(t)$ converges to a zero of $f^{(r)}$ (see Theorem 2). For $r = 0$, it approaches an orbit of constant entropy relative to a zero of $f^{(r)}$.

$n$-player games.[7]

It has been observed repeatedly that without mutations (i.e., $r = 0$), the deterministic process described by the differential equation (7) does not in general converge, but only approach an orbit of constant entropy relative to a zero of the fitness function $f$. When introducing mutations, the limiting behavior of the process changes qualitatively (see Figure 4). As Theorem 2 shows, it then converges to a zero of the fitness function (see Appendix B).

**Theorem 2.** *Let $f\colon \Delta \to \mathbb{R}^d$ be defined by*

$$f_i^{(r)}(p) = (1-r)p_i(\tilde{M}p)_i + r(\frac{1}{d} - p_i)$$

*If $r > 0$, $f^{(r)}$ has a unique zero $p^{(r)}$ and the unique solution $y(t)$ of*

$$\begin{aligned} \frac{d}{dt}y(t) &= f^{(r)}(y(t)) \\ y(0) &= p \end{aligned} \tag{7}$$

*converges to $p^{(r)}$ as $t \to \infty$. Moreover, if $r$ goes to infinity, then $p^{(r)}$ converges to $ML(R)$.*

We believe that this result as well as Theorem 1 and Corollary 1 are also of relevance for the natural sciences. In particular, a discrete model may describe the mentioned natural phenomena more accurately than continuous ones.

[7] In Benaim and Weibull's model, each player has a population of $N$ individuals who play pure strategies. In each round, one individual of one player can update their strategy based on their payoff against the strategies of randomly drawn individuals from the other players. An update rule induces a deterministic process described by a differential equation similar to (1). They show that if $N$ is large, the distributions of strategies among the individuals of each role in this probabilistic process approximate the deterministic process described by the differential equation. Our setting corresponds to a symmetric two-player zero-sum game and an update rule based on the matrix of majority margins $\tilde{M}$. The special properties of this instance allow us to make more precise statements about the behavior of the deterministic process, and, thus, of the probabilistic process for large $N$.

### 6.2. Equilibrium dynamics

When interpreting the majority margin matrix as a symmetric two-player zero-sum game and maximal lotteries as equilibrium strategies, our result can be phrased as a result about a dynamic process that converges towards equilibrium play. Equilibrium dynamics have been extensively studied in game theory and, in particular for zero-sum games, a number of simple and attractive processes have been proposed. The earliest of these is *fictitious play* (Brown, 1951). More recently, the *multiplicative weights update algorithm* (e.g., Freund and Schapire, 1999; Arora et al., 2012) and *regret matching* (Hart and Mas-Colell, 2000, 2013) have been celebrated in game theory, optimization, and machine learning. When translating the multiplicative weights update algorithm to our setting, one obtains a dynamic urn process, in which voters need to compare a drawn ball to all possible alternatives and adjust the distribution in the urn accordingly. It does not suffice to replace a single ball and the total number of balls does not remain constant.

### 6.3. Population Protocols

The process we describe approximately computes a mixed Nash equilibrium of a symmetric zero-sum game. This problem is known to be equivalent to linear programming. In fact, deciding whether an action is played with positive probability in an equilibrium of a symmetric zero-sum game is P-complete, even when all payoffs are $-1$, $0$, or $1$ (Brandt and Fischer, 2008, Theorem 5), which roughly means that the problem is at least as hard as any problem that can be solved in polynomial time. The urn process can thus be seen as a probabilistic algorithm that approximates polynomial-time computable functions. In contrast to traditional computing devices such as Turing machines, the urn process is based on unordered elementary entities that randomly interact according to very simple replacement rules. Related decentralized models of computation with applications to sensor networks and molecular computing are studied under the name “population protocols” in computer science (e.g., Angluin et al., 2006; Aspnes and Ruppert, 2009). While the urn process has the same *modus operandi* as population protocols, the input-output behavior is different. The input of population protocols is given by the initial distribution of balls in the urn and the output has been reached if all balls belong to a certain subset of types. By contrast, the input for our urn process is encoded in the matrix describing the replacement rules and the (approximate) output is given by the distribution of balls in the urn after sufficiently many rounds.

## Acknowledgments

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-1 and BR 5969/1-1. The authors thank Stefano Allesina, Vincent Conitzer, Erwin Frey, Philipp Geiger, and Johannes Knebel for stimulating discussions.

## References

S. Allesina and J. M. Levine. A competitive network theory of species diversity. *Proceedings of the National Academy of Sciences (PNAS)*, 108(14):5638–5642, 2011.

D. Angluin, J. Aspnes, Y. Diamadi, M. J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, 2006.

S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8:121–164, 2012.

R. Arratia and L. Gordon. Tutorial on large deviations for the binomial distribution. *Bulletin of Mathematical Biology*, 51(1):125–131, 1989.

K. J. Arrow. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition, 1951. 2nd edition 1963.

J. Aspnes and E. Ruppert. An introduction to population protocols. In *Middleware for Network Eccentric and Mobile Applications*, chapter 5, pages 97–120. Springer-Verlag, 2009.

S. Barberà. Majority and positional voting in a probabilistic framework. *Review of Economic Studies*, 46(2):379–389, 1979.

J. Bartholdi, III, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(3):157–165, 1989.

M. Benaim and J. Weibull. Deterministic approximation of stochastic evolution in games. *Econometrica*, 71(3):873–903, 2003.

S. J. Brams and P. C. Fishburn. Approval voting. *The American Political Science Review*, 72(3):831–847, 1978.

S. J. Brams and P. C. Fishburn. Voting procedures. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 4. Elsevier, 2002.

F. Brandl and F. Brandt. Arrovian aggregation of convex preferences. *Econometrica*, 88 (2):799–844, 2020.

F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.

F. Brandl, F. Brandt, and C. Stricker. An analytical and experimental comparison of maximal lottery schemes. *Social Choice and Welfare*, 2021. Forthcoming.

F. Brandt. Rolling the dice: Recent results in probabilistic social choice. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 1, pages 3–26. AI Access, 2017.

F. Brandt and F. Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.

F. Brandt, M. Brill, and P. Harrenstein. Tournament solutions. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3. Cambridge University Press, 2016a.

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016b.

G. W. Brown. Iterative solutions of games by fictitious play. In *Activity Analysis of Production and Allocation*, pages 374–376. John Wiley and Sons, Inc., 1951.

T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons Inc., 2006.

P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984.

Y. Freund and R. E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1–2):79–103, 1999.

A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4): 587–601, 1973.

A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45 (3):665–681, 1977.

J. Grilli, G. Barabás, M. J. Michalska-Smith, and S. Allesina. Higher-order interactions stabilize dynamics in competitive network models. *Nature*, 548:210–213, 2017.

S. Hart and A. Mas-Colell. A simple and adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

S. Hart and A. Mas-Colell. *Simple Adaptive Strategies*, volume 4 of *Economic Theory*. World Scientific Publishing Company, 2013.

J. C. Heckelman. Probabilistic Borda rule voting. *Social Choice and Welfare*, 21:455–468, 2003.

J. Hofbauer and K. Sigmund. *Evolutionary games and population dynamics*. Cambridge University Press, 1998.

I. Kaplansky. A contribution to von Neumann's theory of games. II. *Linear Algebra and its Applications*, 226-228:371–373, 1995.

J. Knebel, T. Krüger, M. F. Weber, and E. Frey. Coexistence and survival in conservative Lotka-Volterra networks. *Physical Review Letters*, 110, 2013.

J. Knebel, M. F. Weber, T. Krüger, and E. Frey. Evolutionary games of condensates in coupled birth–death processes. *Nature Communications*, 6(6977), 2015.

T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7(1):49–58, 1970.

B. Laslier and J.-F. Laslier. Reinforcement learning from comparisons: Three alternatives is enough, two is not. *The Annals of Applied Probability*, 27(5):2907–2925, 2017.

J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.

J.-F. Laslier. Aggregation of preferences with a variable set of alternatives. *Social Choice and Welfare*, 17(2):269–282, 2000.

J.-F. Laslier. And the loser is... plurality voting. In D. Felsenthal and M. Machover, editors, *Electoral Systems*, Studies in Choice and Welfare, pages 327–351. Springer, 2011.

D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2009.

T. Reichenbach, M. Mobilia, and E. Frey. Coexistence versus extinction in the stochastic cyclic Lotka-Volterra model. *Physical Review E*, 74(051907), 2006.

M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

P. Schuster and K. Sigmund. Replicator dynamics. *Journal of Theoretical Biology*, 100 (3):533–538, 1983.

P. D. Taylor and L. B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1–2):145–156, 1978.

H. P. Young. An axiomatization of Borda's rule. *Journal of Economic Theory*, 9(1): 43–52, 1974.

H. P. Young and A. Levenglick. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.

# APPENDIX: Proofs

## A. Preliminaries

Recall that $\{X^{(N,r)}(n, s_0)\colon n \in \mathbb{N}_0\}$ is a discrete-time, time-homogeneous Markov chain with state space $S^{(N)}$ and transition probability matrix

$$P^{(N,r)}(s, s + e_i - e_j) = \begin{cases} (1-r)\frac{s_i s_j}{\binom{N}{2}} M(i,j) + \frac{r}{d}\frac{s_j}{N} & \text{if } i \neq j \\ (1-r)\sum_{k=1}^{d} \frac{\binom{s_k}{2}}{\binom{N}{2}} + \frac{r}{d} & \text{if } i = j \end{cases}$$

for $s \in S^{(N)}$ and $i, j \in \{1, \dots, d\}$ with $s + e_i - e_j \in S^{(N)}$. All other transition probabilities are 0. If $r > 0$, it is irreducible and aperiodic and thus admits a unique stationary distribution in $\Delta(S^{(N)})$. We omit writing the initial state $s_0$ when it is convenient.

For $i \in [d]$, we calculate the expected change in the $i$th component of $X^{(N,r)}$ given that $X^{(N,r)}$ is in state $s \in S^{(N)}$.

$$\begin{aligned} &\mathbb{E}\left(X_i^{(N,r)}(n+1) - X_i^{(N,r)}(n) \mid X^{(N,r)}(n) = s\right) \\ &= \sum_{s' \in S^{(N)}} (s'_i - s_i) P^{(N,r)}(s, s') \\ &= (1-r)\sum_{j \neq i} \frac{s_i s_j}{\binom{N}{2}} \left(M(i,j) - M(j,i)\right) + \frac{r}{d}\sum_{j \neq i}\left(\frac{s_j}{N} - \frac{s_i}{N}\right) \\ &= (1-r)\frac{s_i}{\binom{N}{2}} \sum_{j \neq i} \tilde{M}(i,j) s_j + \frac{r}{d}\left(\frac{N - s_i}{N} - (d-1)\frac{s_i}{N}\right) \\ &= (1-r)\frac{s_i}{\binom{N}{2}} (\tilde{M}s)_i + r\left(\frac{1}{d} - \frac{s_i}{N}\right) \end{aligned}$$

For the last equality, recall that $\tilde{M}$ is skew-symmetric, so that $\tilde{M}(i,i) = 0$.

We start by studying the function $f^{(N,r)}\colon \Delta \to \mathbb{R}^d$ with

$$f_i^{(N,r)}(p) = (1-r)\frac{N^2}{\binom{N}{2}} p_i (\tilde{M}p)_i + r\left(\frac{1}{d} - p_i\right) = (1-r)\frac{2N}{N-1} p_i (\tilde{M}p)_i + r\left(\frac{1}{d} - p_i\right)$$

Note that for $s \in S^{(N)}$, $f_i^{(N,r)}(\frac{s}{N})$ is the expected change in the $i$th component of $X^{(N,r)}$ by the calculation above.

We show that for positive $r$, $f^{(N,r)}$ has a unique zero $p^{(N,r)}$ over $\Delta$. As $r$ goes to 0, $p^{(N,r)}$ converges to a maximal lottery for the profile that induces $\tilde{M}$.

**Lemma 1.** *For every $r > 0$, there is a unique $p^{(N,r)} \in \Delta$ such that $f^{(N,r)}(p^{(N,r)}) = 0$. Moreover, for every $\delta > 0$, there is $r_0$ so that $p^{(N,r)} \in B_\delta(ML(R))$ for all $r \leq r_0$.*

*Proof.* Let $r > 0$. We show that the function $g^{(N,r)}\colon \Delta \to \Delta$ with $g^{(N,r)}(p) = p + \frac{1}{4} f^{(N,r)}(p)$ has a fixed-point. We verify that $g^{(N,r)}$ maps to $\Delta$. For all $p \in \Delta$,

$\sum_{i\in[d]} f_i^{(N,r)}(p) = (1-r)\frac{2N}{N-1}p^\intercal \tilde{M} p + r\left(1 - \sum_{i\in[d]} p_i\right) = 0$ since $\tilde{M}$ is skew-symmetric and $p \in \Delta$. Moreover, since $N \geq 2$,

$$f_i^{(N,r)}(p) = (1-r)\frac{2N}{N-1}p_i \underbrace{(\tilde{M}p)_i}_{\geq -1} + r\left(\frac{1}{d} - p_i\right) \geq -4p_i$$

Thus,

$$g_i^{(N,r)}(p) \geq p_i + \frac{1}{4}(-4p_i) \geq 0$$

It follows that $g^{(N,r)}$ maps to $\Delta$. Moreover, $g^{(N,r)}$ is continuous since $f^{(N,r)}$ is continuous. Hence, $g^{(N,r)}$ has a fixed-point $p^{(N,r)}$ for which we have $f^{(N,r)}(p^{(N,r)}) = 4(g^{(N,r)}(p^{(N,r)}) - p^{(N,r)}) = 0$. This proves the existence of a solution to $f^{(N,r)}(p) = 0$.

Note that for all $p \in \Delta$ with $f^{(N,r)}(p) = 0$, we have for all $i \in [d]$, $p_i > 0$ since $p_i = 0$ implies $f_i^{(N,r)}(p) = r\frac{1}{d} > 0$. Hence, we can rewrite $f^{(N,r)}(p) = 0$ as follows: for all $i \in [d]$,

$$(1-r)\frac{2N}{N-1}(\tilde{M}p)_i = r\left(1 - \frac{1}{p_i d}\right) \tag{8}$$

To show that $f^{(N,r)}$ has a unique zero for $r > 0$, assume that $f^{(N,r)}(p) = f^{(N,r)}(q) = 0$ for $p, q \in \Delta$. We have

$$\begin{aligned}
0 &= (1-r)\frac{2N}{N-1}\left(p^\intercal \tilde{M} q + q^\intercal \tilde{M} p\right) \\
&= (1-r)\frac{2N}{N-1}\sum_{i\in[d]} p_i\left(\tilde{M}q\right)_i + q_i\left(\tilde{M}p\right)_i \\
&= r\sum_{i\in[d]} p_i\left(1 - \frac{1}{q_i d}\right) + q_i\left(1 - \frac{1}{p_i d}\right) \\
&= r\left(2 - \frac{1}{d}\sum_{i\in[d]}\frac{p_i}{q_i} + \frac{q_i}{p_i}\right) \\
&= -\frac{r}{d}\sum_{i\in[d]}\frac{(p_i - q_i)^2}{p_i q_i} \leq -\frac{r}{d}|p-q|_2^2
\end{aligned}$$

where the first equality uses that $\tilde{M}$ is skew-symmetric (hence, $p^\intercal \tilde{M} q = -q^\intercal \tilde{M} p$), the third equality follows from (8) and the fact that $p$ and $q$ are zeros of $f^{(N,r)}$, and the last two are algebra. ($|\cdot|_2$ denotes the $L^2$-norm.) This sequence of equalities implies that $p = q$. Hence, $p^{(N,r)}$ is the unique zero of $f^{(N,r)}$ for $r > 0$.

Lastly, let $\delta > 0$. By (8), for all $r > 0$ and $i \in [d]$,

$$\left(\tilde{M}p^{(N,r)}\right)_i = \frac{r}{(1-r)}\frac{N-1}{2N}\left(1 - \frac{1}{p_i^{(N,r)} d}\right) \leq \frac{r}{(1-r)}\frac{N-1}{2N} \tag{9}$$

Suppose for every $r_0 > 0$, there is $r < r_0$ so that $p^{(N,r)} \notin B_\delta(ML(R))$. Then we can find a sequence $(r_n)$ going to 0 so that $p^{(N,r_n)} \notin B_\delta(ML(R))$ for all $n$. By passing to a subsequence, we may assume that $p^{(N,r_n)} \to p \notin B_\delta(ML(R))$. But from (9) it follows that $\tilde{M}p \leq 0$ so that $p \in ML(R)$, which is a contradiction.

□

# B. Properties of the Deterministic Process

The function $f^{(N,r)}$ gives rise to a (first-order ordinary) differential equation for continuously differentiable functions from $[0,\infty)$ to $\Delta$, that is, functions in $\mathcal{C}^1([0,\infty),\Delta)$. For $y \in \mathcal{C}^1([0,\infty),\Delta)$ and $p_0 \in \Delta$, consider

$$\begin{aligned} \frac{d}{dt}y(t) &= f^{(N,r)}(y(t)) \\ y(0) &= p_0 \end{aligned} \tag{10}$$

We show that (10) has a unique global solution $y^{(N,r)}$ for all $r > 0$ and $p_0 \in \Delta$. Moreover, this solution converges to the zero $p^{(N,r)}$ of $f^{(N,r)}$. Since $N$ and $r$ remain fixed for now, we frequently omit the superscript $(N,r)$.

The proof that (10) has a unique local solution with values in $\mathbb{R}^d$ is standard. Only the fact that the solution does not leave the domain $\Delta$ of $f$ and can thus be extended to a global solution requires attention.

**Lemma 2.** *For every $p_0 \in \Delta$, (10) has a unique solution $y \in \mathcal{C}^1([0,\infty),\Delta)$ with $y(0) = p_0$.*

*Proof.* Note that $f$ is Lipschitz-continuous in a neighborhood of $\Delta$. It follows from the Picard-Lindelöf Theorem that for any $t_0 \in [0,\infty)$ and $p \in \Delta$, the system

$$\begin{aligned} \frac{d}{dt}y(t) &= f(y(t)) \\ y(t_0) &= p \end{aligned} \tag{11}$$

has a unique local solution, that is, a solution $y \in \mathcal{C}^1((t_0-\varepsilon, t_0+\varepsilon),\mathbb{R}^d)$.

We observe that $y$ maps to $\Delta$. First, by the same arguments as in the proof of Lemma 1, we have

$$\frac{d}{dt}\sum_{i\in[d]} y_i(t) = \sum_{i\in[d]} f_i(y(t)) = 0$$

whenever $y(t) \in \Delta$. Second, if $y_i(t) = 0$, then $\frac{d}{dt}y_i(t) = f_i(y(t)) > 0$. Hence, $y(t) \in \Delta$ for all $t \in (t_0-\varepsilon, t_0+\varepsilon)$. Since $t_0 \in [0,\infty)$ was arbitrary, it follows that $y$ can be uniquely extended to a global solution in $\mathcal{C}^1([0,\infty),\Delta)$. □

Denote by $y^{(N,r)}(t, p_0) \in \mathcal{C}^1([0,\infty), \Delta)$ the unique solution to (10) with $y^{(N,r)}(0, p_0) = p_0$. We will sometimes suppress the argument $p_0$ when it is clear from the context.

We want to show that if $r > 0$, $y^{(N,r)}(t, p_0)$ converges to the zero $p^{(N,r)}$ of $f^{(N,r)}$ as $t$ goes to infinity. In fact, the convergence is uniform in $p_0$. The proof of this fact in Lemma 4 uses the relative entropy (aka the Kulback-Leiber-Divergence) of $p, q \in \Delta$, which is defined as

$$D(p \mid q) = \sum_{i\in[d]} p_i \log\left(\frac{p_i}{q_i}\right)$$

Moreover, the following lower bound on the relative entropy will be helpful (see, e.g., Cover and Thomas, 2006, Lemma 11.6.1).

**Lemma 3.** *For all* $p, q \in \Delta$,

$$D(p \mid q) \geq \frac{1}{2\log 2} |p - q|^2$$

**Lemma 4.** *Let* $r > 0$. *Then,*

$$\lim_{t\to\infty} \sup\left\{\left|y^{(N,r)}(t, p_0) - p^{(N,r)}\right| : p_0 \in \Delta\right\} = 0$$

*Proof.* Since $N$ and $r$ stay fixed throughout the proof, they is suppressed except in $p^{(N,r)}$.

For now, fix $p_0$ in the interior of $\Delta$ and write $y = y(\cdot, p_0)$. We show that the entropy of $p^{(N,r)}$ relative to $y(t)$ decreases at a rate of at least $\frac{r}{d\sqrt{d}}\left|p^{(N,r)} - y(t)\right|_2^2$.

$$\begin{aligned}
\frac{d}{dt}D(p^{(N,r)} \mid y(t)) &= \frac{d}{dt}\sum_{i\in[d]} p_i^{(N,r)} \log\left(\frac{p_i^{(N,r)}}{y_i(t)}\right) = -\sum_{i\in[d]} p_i^{(N,r)} \frac{\frac{d}{dt}y_i(t)}{y_i(t)} \\
&\overset{(i)}{=} -\sum_{i\in[d]} p_i^{(N,r)} \frac{f_i(y(t))}{y_i(t)} \\
&= -\sum_{i\in[d]} p_i^{(N,r)} \frac{(1-r)\frac{2N}{N-1}y_i(t)(\tilde{M}y(t))_i + r\left(\frac{1}{d} - y_i(t)\right)}{y_i(t)} \\
&= -(1-r)\frac{2N}{N-1}\sum_{i\in[d]} p_i^{(N,r)}(\tilde{M}y(t))_i - r\sum_{i\in[d]} p_i^{(N,r)}\left(\frac{1}{y_i(t)d} - 1\right) \\
&\overset{(ii)}{=} (1-r)\frac{2N}{N-1}\sum_{i\in[d]} y_i(t)(\tilde{M}p^{(N,r)})_i - r\left(\sum_{i\in[d]}\frac{p_i^{(N,r)}}{y_i(t)d} - 1\right) \\
&\overset{(iii)}{=} \sum_{i\in[d]} y_i(t) r\left(1 - \frac{1}{p_i^{(N,r)}d}\right) - r\left(\sum_{i\in[d]}\frac{p_i^{(N,r)}}{y_i(t)d} - 1\right) \\
&= r\left(2 - \frac{1}{d}\sum_{i\in[d]}\frac{y_i(t)}{p_i^{(N,r)}} + \frac{p_i^{(N,r)}}{y_i(t)}\right)
\end{aligned}$$

$$\stackrel{(iv)}{=} -\frac{r}{d}\sum_{i\in[d]}\frac{(p_i^{(N,r)}-y_i(t))^2}{p_i^{(N,r)}y_i(t)} \le -\frac{r}{d\sqrt{d}}\left|p^{(N,r)}-y(t)\right|^2$$

Here, $(i)$ follows from the fact that $y$ satisfies (10), $(ii)$ uses the skew-symmetry of $\tilde{M}$ and $\sum_{i\in[d]} p_i^{(N,r)} = 1$, $(iii)$ uses (8), and $(iv)$ uses $a^2+b^2 = (a+b)^2-2ab$ for any $a, b \in \mathbb{R}$. It follows that for $t \ge t_0 \ge 0$,

$$0 \le D(p^{(N,r)} \mid y(t)) = D(p^{(N,r)} \mid y(t_0)) + \int_{t_0}^{t}\frac{d}{ds}D(p^{(N,r)} \mid y(s))ds \le D(p^{(N,r)} \mid y(t_0))$$

Combining this with the conclusion from the sequence of equalities above, we see that

$$0 \le \frac{r}{d}\sum_{i=1}^{d}\int_{t_0}^{t}\frac{(p_i^{(N,r)}-y_i(s))^2}{p_i^{(N,r)}y_i(s)}ds = -\int_{t_0}^{t}\frac{d}{ds}D(p^{(N,r)} \mid y(s))ds \le D(p^{(N,r)} \mid y(t_0)) \quad (12)$$

We want to prove that $y(t, p_0)$ convergences to $p^{(N,r)}$ uniformly in $p_0$ as $t$ goes to $\infty$. That is, for all $\varepsilon > 0$, there exists $T > 0$ such that for all $t \ge T$ and all $p_0 \in \Delta$, $|y(t, p_0) - p^{(N,r)}| < \varepsilon$.

To this end, first note that if $y_i(t, p_0) < \frac{r}{8d}$, then since $N \ge 2$,

$$\frac{d}{dt}y_i(t,p_0) = (1-r)\frac{2N}{N-1}y_i(t,p_0)\underbrace{\left(\tilde{M}y(t,p_0)\right)_i}_{\ge -1} + r\left(\frac{1}{d}-y_i(t,p_0)\right) \ge -\frac{r}{2d}+\frac{r}{d} \ge \frac{r}{2d}$$

Hence, for all $p_0 \in \Delta$, $i \in [d]$, and $t \ge 1$, $y_i(t, p_0) \ge \frac{r}{8d}$. We can thus upper bound $D(p^{(N,r)} \mid y(t, p_0))$ for all $p_0 \in \Delta$ and $t \ge 1$ by $C = \max_{p\in\Delta^r} D(p^{(N,r)} \mid p) < \infty$, where $\Delta^r = \{p \in \Delta\colon p_i \ge \frac{r}{8d} \text{ for all } i \in [d]\}$.

Now we prove that the convergence is uniform in $p_0$. Let $\varepsilon > 0$. It follows from (12) with $t_0 = 1$ that given $\delta > 0$, for all $p_0 \in \Delta^r$ and $\lambda$ the Lebesgue measure, $\lambda(\{t \ge 1\colon |y(t, p_0) - p^{(N,r)}| \ge \delta\}) \le \frac{Cd\sqrt{d}}{r\delta^2}$. Hence, for every $p_0 \in \Delta^r$, we can find $t_0(p_0, \delta) \in [1, 1 + \frac{Cd\sqrt{d}}{r\delta^2}]$ such that

$$\left|y(t_0(p_0,\delta),p_0) - p^{(N,r)}\right| < \delta \quad (13)$$

Using the estimate $\log(x - \delta) \ge \log(x) - \frac{\delta}{x-\delta}$ for the last inequality, we find that

$$\begin{aligned} D(p^{(N,r)} \mid y(t_0(p_0,\delta))) &\le \sum_{i\in[d]}\log\left(\frac{p_i^{(N,r)}}{p_i^{(N,r)}-\delta}\right) = \sum_{i\in[d]}\log\left(p_i^{(N,r)}\right) - \log\left(p_i^{(N,r)}-\delta\right) \\ &\le \sum_{i\in[d]}\frac{\delta}{p_i^{(N,r)}-\delta} \le \delta C' \end{aligned}$$

where $C' = 2d\max\{\frac{1}{p_i^{(N,r)}}: i \in [d]\}$ if $\delta \in (0, \frac{1}{2}\min\{p_i^{(N,r)}\colon i \in [d]\})$.

We use this bound and the fact that the relative entropy is non-increasing in $t$ to show that $|y(t,p_0) - p^{(N,r)}| < \varepsilon$ for $t \geq t_0(p_0,\delta)$ for sufficiently small $\delta$. By Lemma 3, we have for all $p \in \Delta$, $D(p^{(N,r)} \mid p) \geq \frac{1}{2\log 2}|p^{(N,r)} - p|^2$. Hence, $|p^{(N,r)} - y(t,p_0)| \leq \sqrt{2\log(2)\delta C'}$ for $t \geq t_0(p_0,\delta)$. Recalling that $t_0(p_0,\delta) \leq 1 + \frac{Cd\sqrt{d}}{r\delta^2} =: T$, we have for $\delta \in (0, \frac{\varepsilon^2}{2\log(2)C'})$ that $|p^{(N,r)} - y(t,p_0)| < \varepsilon$ for all $t \geq T$ and $p_0 \in \Delta$. Since $\varepsilon$ was arbitrary, this proves uniform convergence. □

The next lemma states that for any $\delta > 0$, if the process $y^{(N,r)}$ starts very close to $p^{(N,r)}$, it will never get go further than $\delta$ away from $p^{(N,r)}$.

**Lemma 5.** *Let $r > 0$ and $\delta > 0$. Then, there is $\eta > 0$ such that*

$$\sup\left\{\left|y^{(N,r)}(t,p) - p^{(N,r)}\right| : t \geq 0,\, p \in B_\eta(p^{(N,r)})\right\} < \delta$$

*Proof.* Recall first that $p_i^{(N,r)} > 0$ for all $i \in [d]$. By Lemma 3, if $p \notin B_\delta(p^{(N,r)})$, then $D(p^{(N,r)} \mid p) \geq \frac{1}{2\sqrt{2}}\delta^2 =: C$. Since $D(p^{(N,r)} \mid \cdot)$ is continuous on the interior of $\Delta$ and $D(p^{(N,r)} \mid p^{(N,r)}) = 0$, there is $\eta > 0$ such that $D(p^{(N,r)} \mid p) < C$ for all $p \in B_\eta(p^{(N,r)})$. In the proof of Lemma 4, we have seen that $D(p^{(N,r)} \mid y^{(N,r)}(t,p))$ is non-increasing in $t$. Hence, for $p \in B_\eta(p^{(N,r)})$, it follows that $|y^{(N,r)}(t,p) - p^{(N,r)}| < \delta$ for all $t \geq 0$. □

The proofs of Lemma 1, Lemma 2, and Lemma 4 carry over verbatim to give the following theorem. Note that $f^{(r)}$ is the uniform limit of $f^{(N,r)}$ as $N$ goes to infinity.

**Theorem 2.** *Let $f\colon \Delta \to \mathbb{R}^d$ be defined by*

$$f_i^{(r)}(p) = (1-r)p_i(\tilde{M}p)_i + r(\frac{1}{d} - p_i)$$

*If $r > 0$, $f^{(r)}$ has a unique zero $p^{(r)}$ and the unique solution $y(t)$ of*

$$\begin{aligned} \frac{d}{dt}y(t) &= f^{(r)}(y(t)) \\ y(0) &= p \end{aligned} \tag{7}$$

*converges to $p^{(r)}$ as $t \to \infty$. Moreover, if $r$ goes to infinity, then $p^{(r)}$ converges to $ML(R)$.*

## C. Properties of the Probabilistic Process

We study the behavior of the Markov chain $X^{(N,r)}$ by exploring its connections to the deterministic process $y^{(N,r)}$. To this end, it is more convenient to consider $x^{(N,r)}$ with $x^{(N,r)}(k, \frac{s_0}{N}) = \frac{1}{N}X^{(N,r)}(k,s_0)$, which lives in the unit simplex $\Delta$ rather than $S^{(N)}$.

We estimate the distance between $x^{(N,r)}$ and the set of maximal lotteries in several steps. First, we choose $T_0$ large enough so that $y^{(N,r)}(\cdot,p_0)$ is close to $p^{(N,r)}$ for all but a small fraction of every time interval of length $T_0$ for all initial states $p_0$. In Lemma 6, we show that if is $N$ is large enough, $x^{(N,r)}$ approximately solves (the integral equation equivalent to) the differential equation (10) with high probability on the interval $[0, T_0]$ for any initial state. From this we conclude in Lemma 7 that for large enough $N$, $x^{(N,r)}$ is close to $y^{(N,r)}$ with high probability on any interval of length $T_0$, provided both processes start with the same state at the beginning of that interval. Thus, $x^{(N,r)}$ is with high probability approximately equal to $p^{(N,r)}$ for all but a small fraction of iterations on any interval of length $T_0$. Now we chop up the time line into successive intervals of length $T_0$. In expectation, $x^{(N,r)}$ stays close to $y^{(N,r)}$ in a large fraction of intervals. Using an adaption of the strong law of large numbers, we show in Lemma 10 that $x^{(N,r)}$ is *almost surely* close to $p^{(N,r)}$ for all but a small fractions of iterations. Lastly, since by Lemma 1, $p^{(N,r)}$ is close to a maximal lottery if $r$ is small enough, Theorem 1 follows.

The integral equation equivalent to (10) is

$$\begin{aligned} y(t) - y(0) &= \int_0^t f^{(N,r)}(y(s))ds \\ y(0) &= p_0 \end{aligned} \tag{14}$$

We show that $x^{(N,r)}$ approximately satisfies (14) (with the integral replaced by a sum) for large $N$ on bounded time intervals. Lemma 6 below states that for any time $T$ and any $\delta > 0$, we can choose $N$ large enough so that with high probability, $x^{(N,r)}(n, p_0)$ does not violate (14) by more than $\delta$ within the first $NT$ iterations independently of the initial state $p_0 \in \Delta^{(N)}$. For the proof, we use the following proposition due to Kurtz (1970, Proposition 4.1). (The statement is adapted to the current setting.)

**Proposition 1** (Kurtz, 1970). *Let $(z^{(N)})_{N\in\mathbb{N}}$ be a sequence of discrete-time Markov chains with states spaces $A^{(N)}$ and probability transition matrices $Q^{(N)}$. Suppose there exist sequences of positive number $(\alpha_N)$ and $(\varepsilon_N)$,*

$$\lim_{N\to\infty} \alpha_N = \infty \qquad \text{and} \qquad \lim_{N\to\infty} \varepsilon_N = 0$$

*such that*

$$\sup_{N\in\mathbb{N}} \sup_{p\in A^{(N)}} \alpha_N \sum_{q\in A^{(N)}} |p-q| Q^{(N)}(p,q) < \infty \tag{15}$$

*and*

$$\lim_{N\to\infty} \sup_{p\in A^{(N)}} \alpha_N \sum_{q\in A^{(N)}, |p-q|>\varepsilon_N} |p-q| Q^{(N)}(p,q) = 0 \tag{16}$$

*Let*

$$G^{(N)}(p) = \alpha_N \sum_{q\in A^{(N)}} (q-p) Q^{(N)}(p,q)$$

*Then, for every $\delta > 0$ and $T > 0$,*

$$\lim_{N\to\infty} \sup_{p\in A^{(N)}} \mathbb{P}\left( \sup_{n\le\alpha_n T} \left| z^{(N)}(n) - z^{(N)}(0) - \sum_{k=0}^{n-1} \frac{1}{\alpha_N} G^{(N)}(z^{(N)}(k)) \right| > \delta \mid z^{(N)}(0) = p \right) = 0$$

The following lemma applies this result to $(x^{(N,r)})_{N\in\mathbb{N}}$ for a fixed $r$.

**Lemma 6.** *For every $T > 0$ and $\delta > 0$,*

$$\lim_{N\to\infty} \sup_{p\in\Delta^{(N)}} \mathbb{P}\left( \sup_{n\le NT} \left| x^{(N,r)}(n,p) - x^{(N,r)}(0,p) - \sum_{k=0}^{n-1} f^{(N,r)}\left(x^{(N,r)}(k,p)\right) \right| \ge \delta \right) = 0 \tag{17}$$

*Proof.* Recall that $P^{(N,r)}$ is the transition probability matrix of $X^{(N,r)}$ and, hence, of $x^{(N,r)}$. (Here, we view $P^{(N,r)}$ as a function from $\Delta^{(N)} \times \Delta^{(N)} \to \mathbb{R}$ instead of $S^{(N)} \times S^{(N)} \to \mathbb{R}$.) We apply Proposition 1 with $z^{(N)} = x^{(N,r)}$, $A^{(N)} = \Delta^{(N)}$, $Q^{(N)} = P^{(N,r)}$, $\alpha_N = N$, and $\varepsilon_N = \frac{2}{N}$. We check (15) and (16):

$$\begin{aligned} &\sup_{N\in\mathbb{N}} \sup_{p\in\Delta^{(N)}} N \sum_{q\in\Delta^{(N)}} |p-q| P^{(N,r)}(Np, Nq) \\ = &\sup_{N\in\mathbb{N}} \sup_{p\in\Delta^{(N)}} N \sum_{i,j=1}^{d} \frac{1}{N} |e_i - e_j| P^{(N,r)}(Np, Np - e_i + e_j) \le 2 \end{aligned}$$

and

$$\lim_{N\to\infty} \sup_{p\in\Delta^{(N)}} N \sum_{q\in\Delta^{(N)}:\ |p-q|>\frac{2}{N}} |p-q| P^{(N,r)}(Np, Nq) = 0$$

Recalling the definition of $f^{(N,r)}$ shows that $G^{(N)} = f^{(N,r)}$ for all $N$. Hence, (17) follows. □

Since we want to compare the discrete-time process $x^{(N,r)}$ to the continuous-time process $y^{(N,r)}$ solving (10), it is convenient to turn $x^{(N,r)}$ into a continuous-time process. To this end, let $\bar{x}^{(N,r)}(t,p) = x^{(N,r)}(\lfloor Nt \rfloor, p)$ for all $t \ge 0$ and $p \in \Delta$. $\bar{x}^{(N,r)}$ is a right-continuous step function, which takes steps of length $\frac{1}{N}|e_i - e_j| = \frac{2}{N}$ and is constant on time intervals $[\frac{k}{N}, \frac{k+1}{N})$. Thus, as $N$ grows, the steps become smaller and appear in shorter intervals. Lemma 6 shows that on any bounded time interval, $\bar{x}^{(N,r)}$ satisfies (14) up to an arbitrary some error with high probability when $N$ is large enough. That is, for every $T > 0$ and $\delta > 0$,

$$\lim_{N\to\infty} \sup_{p\in\Delta^{(N)}} \mathbb{P}\left( \sup_{t\le T} \left| \bar{x}^{(N,r)}(t,p) - \bar{x}^{(N,r)}(0,p) - \int_0^t f^{(N,r)}\left(\bar{x}^{(N,r)}(s,p) ds\right) \right| \ge \delta \right) = 0 \tag{18}$$

In Lemma 7, we show that this implies that the trajectories of $y^{(N,r)}(\cdot,p)$ and $\bar{x}^{(N,r)}(\cdot,p)$ stay close to each other with high probability on a given bounded time interval for any initial state $p$ for large $N$. Importantly for later use, the bound on the probability is uniform in $p$.

**Lemma 7.** *For every $T > 0$ and $\delta > 0$,*

$$\lim_{N\to\infty} \sup_{p\in\Delta^{(N)}} \mathbb{P}\left(\sup_{t\le T}\left|y^{(N,r)}(t,p) - \bar{x}^{(N,r)}(t,p)\right| \ge \delta\right) = 0 \tag{19}$$

*Proof.* First observe that since $f^{(N,r)}$ is continuously differentiable on the compact space $\Delta$, there is $C \in \mathbb{R}_{\ge 0}$ such that $f^{(N,r)}$ is Lipschitz-continuous with constant $C$. Let $T > 0$, $\delta > 0$, and $p \in \Delta$. If $\sup_{t\le T} |\bar{x}^{(N,r)}(t,p) - \bar{x}^{(N,r)}(0,p) - \int_0^t f^{(N,r)}(\bar{x}^{(N,r)}(s,p))ds| < \varepsilon$, then for all $t \in [0,T]$,

$$\begin{aligned}\left|y^{(N,r)}(t,p) - \bar{x}^{(N,r)}(t,p)\right| &= \left|y^{(N,r)}(t,p) - y^{(N,r)}(0,p) - \bar{x}^{(N,r)}(t,p) + \bar{x}^{(N,r)}(0,p)\right| \\ &< \varepsilon + \int_0^t \left|f^{(N,r)}\left(y^{(N,r)}(s,p)\right) - f^{(N,r)}\left(\bar{x}^{(N,r)}(s,p)\right)\right| ds \\ &\le \varepsilon + C\int_0^t \left|y^{(N,r)}(s,p) - \bar{x}^{(N,r)}(s,p)\right| ds\end{aligned}$$

The first inequality follows from the assumption about $\bar{x}^{(N,r)}$ and the fact that $y^{(N,r)}$ satisfies (14). The second inequality uses the Lipschitz-continuity of $f^{(N,r)}$. We apply Grönwall's inequality to conclude that

$$\sup_{t\le T}\left|y^{(N,r)}(t,p) - \bar{x}^{(N,r)}(t,p)\right| < \varepsilon e^{CT} < \delta$$

for $\varepsilon > 0$ small enough. Note that the choice of $\varepsilon$ does not depend on $p = \bar{x}^{(N,r)}(0,p)$.

By (18), for every $\rho > 0$, we can find $N_0 \in \mathbb{N}$ such that for every $N \ge N_0$,

$$\sup_{p\in\Delta^{(N)}} \mathbb{P}\left(\sup_{t\le T}\left|\bar{x}^{(N,r)}(t,p) - \bar{x}^{(N,r)}(0,p) - \int_0^t f^{(N,r)}\left(\bar{x}^{(N,r)}(s,p)\right) ds\right| \ge \varepsilon\right) < \rho$$

Hence, for all $N \ge N_0$,

$$\sup_{p\in\Delta^{(N)}} \mathbb{P}\left(\sup_{t\le T}|y^{(N,r)}(t,p) - \bar{x}^{(N,r)}(t,p)| \ge \delta\right) < \rho$$

Since $\rho$ was arbitrary, (19) follows. □

The last tool, Lemma 9, is in essence a one-sided strong law of large numbers for indicator random variables. Instead of the usual assumption of i.i.d. random variables, it only assumes that the probability of each variable being 1 is conditionally upper bounded. The proof uses the following auxiliary lemma about binomial distributions.

**Lemma 8.** *Let $\alpha \in [0,1]$ and $\{B_n : n \in \mathbb{N}_0\}$ be binomial distributions with $B_n \sim B(n,\alpha)$. Then, for any $\varepsilon > 0$, $\sum_{n\in\mathbb{N}_0} \mathbb{P}\left(\frac{B_n}{n} > \alpha + \varepsilon\right) < \infty$.*

*Proof.* We use the following tail bound for binomial distributions, which can be obtained from the Chernoff bound (see, e.g., Arratia and Gordon, 1989).

$$\mathbb{P}\left(\frac{B_n}{n} \geq \alpha + \varepsilon\right) \leq \exp\left(-nD\left(1-\alpha-\varepsilon \mid 1-\alpha\right)\right)$$

where $D(\beta \mid \gamma) = \beta \log \frac{\beta}{\gamma} + (1-\beta)\log\frac{1-\beta}{1-\gamma}$ is the relative entropy of indicator random variables with success probabilities $\beta$ and $\gamma$. Since the lemma obviously holds if $\alpha \in \{0,1\}$, we may assume that $0 < \alpha < \alpha + \varepsilon < 1$. Let $C(\varepsilon) = D\left(1-\alpha-\varepsilon \mid 1-\alpha\right) \in (0,\infty)$. Then,

$$\sum_{n\in\mathbb{N}_0} \mathbb{P}\left(\frac{B_n}{n} \geq \alpha + \varepsilon\right) \leq \sum_{n\in\mathbb{N}_0} \exp\left(-nC(\varepsilon)\right) < \infty$$

which is what we needed to show. □

**Lemma 9.** *Let $\alpha \in [0,1]$. Let $\{Z_n : n \in \mathbb{N}_0\}$ be indicator random variables and, for $n \geq 1$, $S_n = \sum_{k=1}^n Z_k$. If $\mathbb{P}(Z_1 = 1) \leq \alpha$ and for all $n \geq 2$, $\mathbb{P}(Z_n = 1 \mid S_{n-1}) \leq \alpha$, then*

$$\mathbb{P}\left(\limsup_{n\to\infty} \frac{S_n}{n} > \alpha\right) = 0$$

*Proof.* Let $\{B_n : n \in \mathbb{N}_0\}$ be binomial distributions with $B_n \sim B(n,\alpha)$. We first prove that for all $n \geq 1$ and $l \in \{0,\ldots,n\}$,

$$\mathbb{P}\left(S_n \geq l\right) \leq \mathbb{P}\left(B_n \geq l\right) \tag{20}$$

We proceed by induction over $n$. For $n = 1$, (20) follows from the assumption that $\mathbb{P}(Z_1 = 1) \leq \alpha$. Now let $n \geq 2$ and assume the statement holds for smaller values of $n$. Let $l \in \{0,\ldots,n\}$. The case $l = 0$ is trivial since then both probabilities are 1. So assume that $l \geq 1$. We distinguish two cases.

*Case* 1. If $\mathbb{P}(S_{n-1} = l-1) \leq \mathbb{P}(B_{n-1} = l-1)$, we get that

$$\begin{aligned}
\mathbb{P}\left(S_n \geq l\right) &= \mathbb{P}\left(S_{n-1} \geq l\right) + \mathbb{P}\left(S_{n-1} = l-1\right)\mathbb{P}\left(Z_n = 1 \mid S_{n-1} = l-1\right) \\
&\leq \mathbb{P}\left(B_{n-1} \geq l\right) + \mathbb{P}\left(B_{n-1} = l-1\right)\alpha \\
&= \mathbb{P}\left(B_n \geq l\right)
\end{aligned}$$

For the inequality, the bound on the first term follows from the induction hypothesis if $l \leq n-1$ and is trivial if $l = n$ since $\mathbb{P}(S_{n-1} \geq n) = \mathbb{P}(B_{n-1} \geq n) = 0$. For the second term, we use the assumption of the present case for the first factor and the hypothesis of the lemma for the second factor.

*Case* 2. If $\mathbb{P}(S_{n-1} = l-1) \geq \mathbb{P}(B_{n-1} = l-1)$, we get that

$$\begin{aligned}
\mathbb{P}\left(S_n \geq l\right) &= \mathbb{P}\left(S_{n-1} \geq l-1\right) - \mathbb{P}\left(S_{n-1} = l-1\right)\left(1 - \mathbb{P}\left(Z_n = 1 \mid S_{n-1} = l-1\right)\right) \\
&\leq \mathbb{P}\left(B_{n-1} \geq l-1\right) - \mathbb{P}\left(B_{n-1} = l-1\right)(1-\alpha) \\
&= \mathbb{P}\left(B_n \geq l\right)
\end{aligned}$$

The inequality holds by the same arguments as in the first case. The only difference is that we now apply the induction hypothesis at $l-1$.

This proves (20). We use Lemma 8 to conclude that for every $\varepsilon > 0$,

$$\sum_{n\in\mathbb{N}_0} \mathbb{P}\left(\frac{S_n}{n} \geq \alpha + \varepsilon\right) \leq \sum_{n\in\mathbb{N}_0} \mathbb{P}\left(\frac{B_n}{n} \geq \alpha + \varepsilon\right) < \infty$$

Thus, it follows from the Borel-Cantelli Lemma that $\mathbb{P}\left(\limsup_{n\to\infty} \frac{S_n}{n} \geq \alpha + \varepsilon\right) = 0$. Using that $\mathbb{P}(\cdot)$ is countably additive, this gives

$$\mathbb{P}\left(\limsup_{n\to\infty} \frac{S_n}{n} > \alpha\right) \leq \sum_{m\in\mathbb{N}_0} \mathbb{P}\left(\limsup_{n\to\infty} \frac{S_n}{n} \geq \alpha + \frac{1}{m}\right) = 0$$

as desired. □

Putting together Lemma 4, Lemma 7, and Lemma 9, we show that $\bar{x}^{(N,r)}$ is almost surely close to $p^{(N,r)}$ most of the time for large enough $N$. More precisely, for $\delta > 0$, we consider the fraction of time up to time $T$ that $\bar{x}^{(N,r)}$ is further than $\delta$ away from $p^{(N,r)}$. For $N$ large enough, the probability that this fraction is larger than $\delta$ goes to 0 as $T$ goes to $\infty$.

**Lemma 10.** *Let $\delta, \tau > 0$ and $r > 0$. Then, there is $N_0$ such that for all $N \geq N_0$ and $p_0 \in \Delta^{(N)}$,*

$$\mathbb{P}\left(\liminf_{T\to\infty} \frac{1}{T}\lambda\left\{t \in [0,T] \colon x^{(N,r)}(t, p_0) \in B_\delta(p^{(N,r)})\right\} \geq 1 - \tau\right) = 1 \tag{21}$$

*where $\lambda$ is the Lebesgue measure.*

*Proof.* By Lemma 5, we can find $\eta > 0$ such that

$$\sup\left\{\left|y^{(N,r)}(t,p) - p^{(N,r)}\right| : t \geq 0,\, p \in B_\eta(p^{(N,r)})\right\} < \frac{\delta}{2}$$

By Lemma 4, we can find $T_1 > 0$ such that for all $T \geq T_1$,

$$\sup\left\{\left|y^{(N,r)}(T,p) - p^{(N,r)}\right| : p \in \Delta\right\} < \eta$$

Let $T_0 = \frac{2}{\tau}T_1$. Note that $y^{(N,r)}$ is time-invariant, that is, $y^{(N,r)}(t,p) = y^{(N,r)}(t - t_0, y^{(N,r)}(t_0,p))$ for all $t \geq t_0 \geq 0$. Combining these facts, it follows that for every $p \in \Delta$, the measure of $t \in [t_0, t_0 + T_0]$ for which $y^{(N,r)}(t,p)$ is in an $\frac{\delta}{2}$-ball around $p^{(N,r)}$ is at least $(1 - \frac{\tau}{2})T_0$. We may assume that $T_0$ is integral.

By Lemma 7, there is $N_0 \in \mathbb{N}$ such that for all $N \geq N_0$,

$$\sup_{p\in\Delta^{(N)}} \mathbb{P}\left(\sup_{0\leq t\leq T_0} \left|y^{(N,r)}(t,p) - \bar{x}^{(N,r)}(t,p)\right| \geq \frac{\delta}{2}\right) < \frac{\tau}{2}$$

Now fix $N \geq N_0$ and $p \in \Delta^{(N)}$. We upper bound the fraction of time $\bar{x}^{(N,r)}$ is further than $\delta$ away from $p^{(N,r)}$. To simplify notation, let $t_k = kT_0$ and $\bar{x}_k = \bar{x}^{(N,r)}(t_k, p)$.

For $n \geq 1$, we calculate the expected number of intervals $[t_{k-1}, t_k]$, $1 \leq k \leq n$ so that $|\bar{x}(t,p) - p^{(N,r)}| \geq \delta$ for some $t \in [t_{k-1}, t_k]$. Let $Z_k^{(N,r)} = \chi(\sup_{t_{k-1} \leq t \leq t_k} |\bar{x}^{(N,r)}(t,p) - y^{(N,r)}(t - t_{k-1}, \bar{x}_{k-1})| \geq \frac{\delta}{2})$ be the indicator variable for the event that $\bar{x}^{(N,r)}(t,p)$ and $y^{(N,r)}(t - t_{k-1}, \bar{x}_{k-1})$ differ by at least $\frac{\delta}{2}$ on the time interval $[t_{k-1}, t_k]$ given that both start at the point $\bar{x}_{k-1}$ at time $t_{k-1}$. Notice that $\{Z_k^{(N,r)} \colon k \in \mathbb{N}_0\}$ satisfies the hypothesis of Lemma 9 with $\alpha = \frac{\tau}{2}$.[8]

If $Z_k^{(N,r)} = 0$, then

$$\begin{aligned}
&\lambda\left\{t \in [t_{k-1}, t_k] \colon \left|\bar{x}^{(N,r)}(t,p) - p^{(N,r)}\right| \geq \delta\right\} \\
&\leq \lambda\left\{t \in [t_{k-1}, t_k] \colon \left|\bar{x}^{(N,r)}(t,p) - y^{(N,r)}(t - t_{k-1}, \bar{x}_{k-1})\right| \geq \frac{\delta}{2}\right\} \\
&+ \lambda\left\{t \in [t_{k-1}, t_k] \colon \left|y^{(N,r)}(t - t_{k-1}, \bar{x}_{k-1}) - p^{(N,r)}\right| \geq \frac{\delta}{2}\right\} \\
&\leq \frac{\tau}{2} T_0
\end{aligned}$$

It follows that

$$\begin{aligned}
&\frac{1}{nT_0} \sum_{k \in [n]} \lambda\left\{t \in [t_{k-1}, t_k] \colon \left|\bar{x}^{(N,r)}(t,p) - p^{(N,r)}\right| \geq \delta\right\} \\
&= \sum_{\substack{k \in [n] \\ Z_k^{(N,r)} = 0}} \frac{1}{nT_0} \lambda\left\{t \in [t_{k-1}, t_k] \colon \left|\bar{x}^{(N,r)}(t,p) - p^{(N,r)}\right| \geq \delta\right\} \\
&+ \sum_{\substack{k \in [n] \\ Z_k^{(N,r)} = 1}} \frac{1}{nT_0} \lambda\left\{t \in [t_{k-1}, t_k] \colon \left|\bar{x}^{(N,r)}(t,p) - p^{(N,r)}\right| \geq \delta\right\} \\
&\leq \frac{\tau}{2} + \frac{1}{n} \sum_{k=1}^{n} Z_k^{(N,r)}
\end{aligned}$$

Applying Lemma 9 to $\{Z_k^{(N,r)} \colon k \in \mathbb{N}_0\}$ with $\alpha = \frac{\tau}{2}$ gives

$$\mathbb{P}\left(\limsup_{n \to \infty} \frac{1}{n} \sum_{k=1}^{n} Z_k^{(N,r)} \geq \frac{\tau}{2}\right) = 0$$

Hence, with the preceding inequality we get

$$\mathbb{P}\left(\limsup_{T \to \infty} \frac{1}{T} \lambda\left\{t \in [0,T] \colon \bar{x}^{(N,r)}(t,p) \notin B_\delta(p^{(N,r)})\right\} \geq \tau\right) = 0$$

This is a restatement of (21). □

[8] While the probability that $Z_n^{(N,r)}$ equals 1 may depend on $Z_k^{(N,r)}$ for $k < n$, the bound of $\frac{\tau}{2}$ holds independently of the $Z_k^{(N,r)}$ since the bound obtained in Lemma 7 is uniform in the initial state $p$.

**Theorem 1.** *Let $\delta, \tau > 0$. Then, there are $N_0 \in \mathbb{N}$ and $r > 0$ such that for all $N \geq N_0$ and $s_0 \in S^{(N)}$,*

$$\mathbb{P}\left(\liminf_{n\to\infty} \frac{1}{n} \# \left\{k \in [n] \colon \frac{1}{N} X^{(N,r)}(k, s_0) \in B_\delta\left(ML(R)\right)\right\} \geq 1 - \tau\right) = 1$$

*Proof.* By definition of $\bar{x}^{(N,r)}$, we can rewrite the conclusion of the theorem as

$$\mathbb{P}\left(\liminf_{T\to\infty} \frac{1}{T} \lambda \left\{t \in [0, T] \colon \bar{x}^{(N,r)}(t, p_0) \in B_\delta\left(ML(R)\right)\right\} \geq 1 - \tau\right) = 1 \qquad (22)$$

where $p_0 = \frac{s_0}{N}$.

By Lemma 1, we can choose $r > 0$ so that $p^{(N,r)} \in B_{\frac{\delta}{2}}(ML(R))$. Applying Lemma 10 to $\frac{\delta}{2}, \tau$, and $r$, we get $N_0$ such that (21) holds (with $\frac{\delta}{2}$ in place of $\delta$). Combining these two facts gives (22). $\square$