

Iterated Majority Voting

Stéphane Airiau and Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam
{s.airiau,u.endriss}@uva.nl

Abstract. We study a model in which a group of agents make a sequence of collective decisions on whether to remain in the current state of the system or switch to an alternative state, as proposed by one of them. Examples for instantiations of this model include the step-wise refinement of a bill of law by means of amendments to be voted on, as well as resource allocation problems, where agents successively alter the current allocation by means of a sequence of deals. We specifically focus on cases where the majority rule is used to make each of the collective decisions, as well as variations of the majority rule where different quotas need to be met to get a proposal accepted. In addition, we allow for cases in which the same proposal may be made more than once. As this can lead to infinite sequences, we investigate the effects of introducing a deadline bounding the number of proposals that can be made. We use both analytical and experimental means to characterise situations in which we can expect to see a convergence effect, in the sense that the expected payoff of each agent will become independent from the initial state of the system, as long as the deadline is chosen large enough.

1 Introduction

We consider the very general problem where a finite set of *agents* must choose one *alternative* among many, and we are interested in *decentralised* solutions. The alternatives may represent different policies, world states, or allocations of resources, etc. One simple idea is to start with an initial current alternative, let a *random* agent propose a different alternative, and organize a vote between these two alternatives. The process *iterates* using the winner of the election as current alternative. Unfortunately, with no restriction on the agents' preferences, this simple process may iterate forever. In this paper, we investigate the possibility of using a *bound* on the number of iterations. In particular, we investigate the problem of determining whether the choice of the bound can guarantee that no agent benefits unduly from the choice of the initial alternative.

An example from political science that fits our generic model is the step-wise refinement of a bill of law by means of amendments to be voted on (although in this case, repeating the same proposal twice may not be allowed). Another example concerns multiagent resource allocation problems: some work in multiagent systems has focussed on negotiation scenarios where agents approach a solution in small steps rather than computing the best solution in one go [1, 2].

In this case, every possible allocation of resources constitutes an alternative and proposing a new alternative means proposing a deal regarding the reallocation of some items. Here, the voting rule used would typically require each agent affected by the proposal to give their consent. Finally, recent work in computational social choice has shown that decomposing combinatorial voting problems into a sequence of smaller elections has a number of advantages [3].

In general, there are many possible choices of a voting rule. In this paper, we focus on the *majority rule*, which specifies that a proposal is accepted if at least half of the concerned agents (those that are not indifferent) vote in favour. This is clearly a very natural choice, and it is the only rule that is anonymous, neutral, and monotone (May’s Theorem [4]). We also consider generalisations of the majority rule, where a quota different from 50% may be needed to accept a proposal. For example, for some important elections, a higher proportion of votes in favour of an alternative is needed, e.g., a two-thirds majority in parliament is needed to change the constitution. As we allow agents to make the same proposal over and over, there is the possibility for cycles. This phenomenon is linked to the fact that the preference relation we obtain when several individual preferences are aggregated by means of the majority rule need not be transitive. In social choice theory, one approach to address this problem has been to restrict the range of allowed preferences, e.g., to single-peaked preferences or preferences meeting Sen’s triple-wise value restriction [5]. When we have no control over the agents’ preferences, we need to modify the protocol to induce the agents to choose a good social alternative. A simple solution to the problem of cycles is to introduce a deadline that limits the number of iterations.

We assume that the agents’ preferences are common knowledge and that agents are strategic: they will make proposals and vote in elections so as to maximise their expected payoff. What would be a good choice of deadline under these circumstances? If it is small, the choice of initial state will play an important role for the final outcome. If it is large, as we shall see, it is sometimes the case that the expected payoff of an agent becomes independent of the initial state, which provides some level of fairness. Our aim in this paper is to get a clearer understanding of such convergence effects.

The remainder of this paper is as follows. Section 2 further motivates and defines our model. In particular, we detail how strategic agents can compute their best moves using backward induction. Section 3 formally defines the notion of convergence as used here and establishes sufficient conditions on a game for being convergent. Then, Section 4 takes this analysis further by mapping out the convergence behaviour for a wider class of games by means of an experimental study. We conclude with a discussion about related work and present future axes of research.

2 The Model

We study games consisting of a finite set N of n *agents* and a finite set X of m *alternatives* or *states*. Each agent $i \in N$ has a *utility* function $u_i : X \rightarrow \mathbb{R}$, which

is assumed to be common knowledge. Each agent is rational, i.e., it maximizes its expected utility. The utility functions are represented together as an $m \times n$ matrix U_0 , with $U_0(x, i) = u_i(x)$. We do *not* assume that utility is transferable: the utility of two agents may not be comparable (e.g., agents can use different currencies). However, we do assume that agents will use their knowledge of the utility of other agents for predicting their behaviour.

A game proceeds in successive iterations. At each iteration t , there is a current alternative $x(t)$ (a given allocation, a current bill). One agent is randomly selected, with equal probability, to propose a new alternative x^* to be considered (e.g., a new allocation, an amendment to the current bill). An agent may propose an alternative that was already proposed in the past; and it may propose to maintain the *status quo* by proposing $x^* = x(t)$. The agents vote between x^* and $x(t)$. If the proposed alternative wins the election, and we will present the criterion to win an election next, it replaces the current alternative for the next iteration. Else, the current alternative remains in place for the next iteration.

Elections are decided using a *quota system* for some fixed quota q : a proposal will be declared the winner iff it receives at least q percent of the votes. More precisely, if n_{\oplus} agents are voting in favour and n_{\ominus} against a proposal (and some agents may abstain), then the proposal is accepted if $n_{\oplus} > q \cdot (n_{\oplus} + n_{\ominus})$. The standard majority rule is the quota system with $q = 50\%$. When the majority rule is used, cycles may occur. The same is true for quota systems with $q \neq 50\%$. In the presence of a cycle, the sequence of elections could be infinite. In order to force the eventual choice of an alternative, we propose the use of a *deadline* limiting the number of iterations to be played. The following definition summarises the components that make up a game:

Definition 1 (Game). *A game is a quadruple $\langle N, X, U_0, q \rangle$, where N with $n = |N|$ is a finite set of agents, X with $m = |X|$ is a finite set of states (or alternatives), U_0 is an $m \times n$ matrix defining the utility each agent assigns to each state, and $q \in [0, 1]$ is a quota (typically expressed in percent).*

Playing a game requires us to also specify a *deadline*, i.e., the number of iterations to be played, and an *initial state* from X .

2.1 Backward induction

Agents are assumed to be expected-utility maximisers, i.e., the goal of a proposal or a vote is to maximise expected utility in the final state. We now discuss how to perform this strategic vote and strategic choice of proposals.

Let the matrix W_t of size $m \times m$ specify the transition between the alternatives at iteration t . An entry $W_t(x, y)$ of W_t is the probability to have alternative y become the current alternative at the next iteration when alternative x was current at iteration t . A row of W_t is a probability distribution over the alternatives that can become current at the next step. Hence, the sum of the entries of a row is equal to 1, i.e., W_t is a stochastic matrix. Furthermore, let the matrix U_t of size $m \times n$ contain the expected utility of all agents for all alternatives,

i.e., $U_t(x, i)$ is the expected payoff of agent i for alternative x at iteration t . We have $U_{t+1} = W_{t+1} \cdot U_t$, and therefore:

$$U_{t+1} = \left(\prod_{\tau=t+1}^1 W_\tau \right) \cdot U_0 = W_{t+1} \cdot W_t \cdot W_{t-1} \cdot \dots \cdot W_1 \cdot U_0$$

Next, we discuss how to compute W_{t+1} from U_t . Let us assume that the agents know what to propose and how to vote during the t^{th} iteration for any alternative $x \in X$. Because of the common knowledge assumption on the utility functions, the possible proposals and votes of any agents are also known, hence the matrices W_t, W_{t-1}, \dots, W_1 are known. What should an agent do during iteration $t + 1$?

How to vote? The decision depends on the comparison of the expected utility of the current alternative $x(t)$ with the one of the proposed alternative x^* , i.e., agent i will vote in favour of the proposal when $U_t(x^*, i) > U_t(x(t), i)$ and against when $U_t(x^*, i) < U_t(x(t), i)$. Note that the agent does not vote in case it is indifferent between the two alternatives.

What to propose? First, the agent needs to compute the outcome of the vote between the current state $x(t)$ and each possible alternative $x' \in X$. Let $X^w \subseteq X$ denote the set of winning alternatives against $x(t)$. For agent i , the set of best proposals is $P_i = \text{argmax}_{x' \in X^w} U_t(x', i)$. If the expected utility of the alternatives in P_i is greater than $U_t(x(t), i)$ (the expected utility of the current alternative), then agent i proposes with equal-probability one of the alternatives in P_i . Else, agent i is content with the current alternative and proposes maintaining the *status quo* (there is no decision to make since $x^* = x(t)$). Since each agent has an equal probability to be selected to make a proposal, we can compute the probability of any alternative to be proposed. And since any alternative that is proposed is winning against the current alternative, we can compute the probability of an alternative to become current at the next iteration.

To summarise, a game $\langle N, X, U_0, q \rangle$ induces a sequence of $m \times m$ transition matrices W_1, W_2, \dots , as described above, as well as a sequence of $m \times n$ matrices U_1, U_2, \dots , fixing the expected payoffs for each agent at each iteration, with $U_{t+1} = W_{t+1} \cdot U_t$. Here, “iteration 1” is the final iteration/step in a play of the game, “iteration 2” is the penultimate iteration, and so forth. These matrices allow us to study the game for all possible choices of initial current alternative.

2.2 Example: a cycle with majority voting

Consider the following problem with 3 agents and 3 states. The utility vectors are: $\langle 4, 1, 2 \rangle$ for state a , $\langle 2, 4, 1 \rangle$ for state b and $\langle 1, 3, 4 \rangle$ for state c . The corresponding matrix U_0 is shown in Table 1. With no deadline, the agents would be stuck in a cycle, as a majority of agents would prefer to move to a different state ($a \rightarrow b \rightarrow c \rightarrow a \rightarrow \dots$). Using a deadline and backward induction, the agents can break the cycle. If the agents were voting sincerely instead of strategically in this example, the final state would be entirely determined by the initial state (there is a unique cycle). First, let us explain the computation of the expected payoffs; then we will describe the properties of the outcome.

$$\begin{aligned}
U_0 &= \begin{bmatrix} 4 & 1 & 2 \\ 2 & 4 & 1 \\ 1 & 3 & 4 \end{bmatrix} & W_1 &= \begin{bmatrix} 1/3 & 0 & 2/3 \\ 2/3 & 1/3 & 0 \\ 0 & 2/3 & 1/3 \end{bmatrix} \\
U_1 &= \frac{1}{3} \begin{bmatrix} 6 & 7 & 10 \\ 10 & 6 & 5 \\ 5 & 11 & 6 \end{bmatrix} & W_2 &= \begin{bmatrix} 1 & 0 & 0 \\ 1/3 & 1/3 & 1/3 \\ 2/3 & 0 & 1/3 \end{bmatrix} & U_\infty &= \begin{bmatrix} 2.0795 & 2.6549 & 2.7560 \\ 2.0795 & 2.6549 & 2.7560 \\ 2.0795 & 2.6549 & 2.7560 \end{bmatrix} \\
U_2 &= \frac{1}{9} \begin{bmatrix} 18 & 21 & 30 \\ 21 & 24 & 21 \\ 17 & 25 & 26 \end{bmatrix} & W_3 &= \begin{bmatrix} 1/3 & 2/3 & 0 \\ 0 & 1/3 & 2/3 \\ 2/3 & 0 & 1/3 \end{bmatrix}
\end{aligned}$$

Table 1: Breaking a cycle with a deadline

State a would lose an election against state c , and win an election against state b . If the current alternative is state a one step before the deadline, the second and third agents should propose state c , the first agent should propose the *status quo*. As the agents are chosen to make a proposal with equi-probability, the probability to stay in a is $\frac{1}{3}$, the probability to move to c is $\frac{2}{3}$, and the probability to move to b is zero. This provides the first row of the matrix W_1 in Table 1. We carry this reasoning to complete W_1 . The expected utility of the agents one step before the deadline is $U_1 = W_1 U_0$. We iterate the reasoning to obtain the matrices W_2, U_2 , etc. We note that, in this example, W_1, W_2 and W_3 are different. We have implemented the iterative algorithm for computing the matrices W_t and U_t . For large values of t , U_t converges to a particular matrix (see U_∞ in Table 1). That is, for t large enough, the expected utility of each agent is independent of the initial state (e.g., agent 1’s expected utility approaches 2.0795). Hence, if the deadline is far enough, no agent can take advantage of the initial choice of alternative. Finally, note that, in the limit, the expected utility of an agent is not simply the average utility over the 3 alternatives.

3 Convergence

The example given in the previous section shows that there are instances of games where we will observe some kind of *convergence*. For the example in question, we have seen that the expected payoff of any given agent for any given initial state will converge to a certain value as we increase the deadline (we call this *intra-state convergence*). We have also observed that the expected payoff will become less and less dependent on the initial state as we increase the deadline (we call this *inter-state convergence*). In this section, we will define these notions of convergence formally and identify some classes of games for which convergence can be guaranteed.

3.1 Types of convergence

We define a game as being intra-state convergent if, for any given agent and any initial state, the difference in expected payoff for small changes in the deadline will become arbitrarily small as deadlines become larger:

Definition 2 (Intra-state convergence). A game $\langle N, X, U_0, q \rangle$ is said to be *intra-state convergent* if, for any agent $i \in N$ and any state $x \in X$, $\lim_{t \rightarrow \infty} [U_t(x, i) - U_{t+1}(x, i)] = 0$.

Next, we define a game as being inter-state convergent if, for any agent and any two states, the difference in expected payoff for making either one of these states the initial state can be made arbitrarily small when we increase the deadline:

Definition 3 (Inter-state convergence). A game $\langle N, X, U_0, q \rangle$ is said to be *inter-state convergent* if, for any agent $i \in N$ and any two states $x, x' \in X$, $\lim_{t \rightarrow \infty} [U_t(x, i) - U_t(x', i)] = 0$.

Inter-state convergence provides a level of *fairness*: as long as we choose a sufficiently large deadline, the choice of initial state will not affect the expected payoff of the individual agents. This does not mean that all agents can be expected to do equally well, but it does mean that one important parameter that determines how a game is played (the initial state) does not influence the (expected) outcome. Intra-state convergence does not directly affect fairness, but offers some level of *robustness* of the mechanism: expected payoffs will not depend on the exact deadline chosen (for instance, we would avoid situations whereby an agent's expected payoff could depend on whether the chosen deadline is odd or even).

Our example did exhibit both types of convergence. Both are entailed by a third notion of convergence, expressed in terms of the transition matrices W_t induced by a game.

Definition 4 (Fundamental convergence). A game $\langle N, X, U_0, q \rangle$ is said to be *fundamentally convergent* if the limit of the product of its transition matrices

$$W = \lim_{t \rightarrow \infty} \prod_{\tau=t}^1 W_\tau$$

is a matrix in which all row vectors are identical.

Definition 4 is reminiscent of the Fundamental Limit Theorem for regular Markov chains, which says that if P is the transition matrix of a regular Markov chain, then $\lim_{n \rightarrow \infty} P^n = W$ for some matrix W with identical row vectors [6]. Recall that a stochastic matrix P is a matrix defining a regular Markov chain if there exists a k such that P^k has only non-zero elements. In particular, the theorem applies when P itself only has non-zero elements. Inspection of the standard proof of the Fundamental Limit Theorem shows that the same is true for the product of several *different* matrices, provided that an infinite number of them are zero-free. While there are similarities to our scenario, we stress that the Fundamental Limit Theorem does *not* apply here, because the transition matrices generated by our games need neither be zero-free nor regular.

Before linking the three definitions of convergence, we state a simple property of stochastic matrices. The proof is standard and omitted for lack of space.

Lemma 1. Let A and B be $m \times m$ stochastic matrices. If B has all row vectors identical, then $A \cdot B = B$.

Proposition 1. If a game is fundamentally convergent, then it is also inter-state convergent and intra-state convergent.

Proof. Let $W = \lim_{t \rightarrow \infty} \prod_{\tau=t}^1 W_\tau$ for the game under consideration. Suppose the game is fundamentally convergent, i.e., W is a matrix of identical row vectors. Then the game is also inter-state convergent: if we multiply W with the payoffs for agent i , then we get the same expected payoff for any initial state.

Next, we show that intra-convergence also follows. By Lemma 1, as $\prod_{\tau=t}^1 W_\tau$ converges to a matrix with identical rows, $W_{t+1} \cdot \prod_{\tau=t}^1 W_\tau$ converges to that very same matrix. The former determines $U_t(x, i)$, while the latter determines $U_{t+1}(x, i)$. Hence, their difference must converge to 0. \square

Furthermore, it is not hard to verify that inter-state convergence implies intra-state convergence. Intra-state convergence is weaker than the other two forms of convergence. There are games that are intra-state convergent but not inter-state convergent. A simple example would be a game with 2 agents, 2 states, and $q = 50\%$, where agent 1 prefers state a and agent 2 prefers state b . Then, if the initial state is a , this will remain the *status quo*, independently of the deadline (and analogously for the case where b is the initial state). The transition matrices for this game are all equal to the identity matrix. Hence, inter-state convergence is not satisfied (your expected payoff is equal to the utility you assign to the initial state), while intra-state convergence is (your expected payoff remains constant when we vary the deadline).

3.2 Sufficient conditions

We will now identify sufficient conditions for a game to be convergent. There are some clear-cut cases, when the quota q takes extreme values. First, if $q = 0\%$, i.e., when a single agent in favour is sufficient for a proposal to be accepted, then all forms of convergence are satisfied. In such a game, whichever agent is chosen to make a proposal in the last iteration will propose their favourite state, and that motion will carry—independently from the current state. Hence, W_1 (the transition matrix for the last iteration) will be a matrix with all rows equal. Therefore, by Lemma 1, also the product W of all transition matrices will be such a matrix, which means that the game will satisfy the convergence condition of Definition 4, and by Proposition 1 also all other notions of convergence.

Second, if $q = 100\% - \epsilon$, i.e., when only proposals opposed by no agent are accepted, then a proposal is accepted iff it represents a Pareto improvement. Hence, the final state will be Pareto optimal—provided the deadline is chosen large enough. But if there is more than one Pareto optimal state, then the expected payoff can be different for an agent if either one of these Pareto optimal states is selected as the initial state. Therefore, inter-state convergence is not generally satisfied (and neither is fundamental convergence). Only in very special cases, such as when all agents are indifferent between all states, would be obtain inter-state convergence.

In addition, note that the presence of a Condorcet winner is not sufficient to guarantee convergence to that state. For example, it is possible that no agent proposes the Condorcet winner. In the example in Table 2 we observe inter-state convergence and there is a unique Condorcet winner (the state with payoffs

$$U_0 = \begin{bmatrix} 4 & 4 & 4 \\ 6 & 2 & 3 \\ 2 & 6 & 2 \\ 3 & 0 & 6 \end{bmatrix} \quad W_1 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 2/3 & 1/3 & 0 & 0 \\ 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 0 & 1/3 \end{bmatrix} \quad U_\infty = \begin{bmatrix} 4.3676 & 3.0288 & 3.9973 \\ 4.3676 & 3.0288 & 3.9973 \\ 4.3676 & 3.0288 & 3.9973 \\ 4.3676 & 3.0288 & 3.9973 \end{bmatrix}$$

Table 2: Example with a unique Condorcet winner: the final outcome may not be the Condorcet winner.

$\langle 4, 4, 4 \rangle$), but the expected payoffs in the limit are not the payoffs of the Condorcet winner. This is because from the state with payoff $\langle 2, 6, 2 \rangle$, no agent has an incentive to propose the Condorcet winner.

We now analyse the case of games with two states in detail.

Proposition 2. *Any game with two states is intra-state convergent.*

Proof. In games $\langle N, X, U_0, q \rangle$ with $|X| = 2$, each transition matrix is of the following form:

$$W_t = \begin{pmatrix} p & 1-p \\ 1-q & q \end{pmatrix}$$

Two special cases are $p = q = 1$ (the identity matrix) and $p = q = 0$ (the “switch matrix”). We first claim that no transition matrix can be such a switch matrix: if all agents are indifferent between both states in iteration t , then W_t is the identity matrix. Otherwise, w.l.o.g., assume agent 1 prefers state 1 in iteration t . Then agent 1 will propose the *status quo* when in state 1, so $p \geq \frac{1}{n}$. This proves our claim.

So each transition matrix will either be the identity matrix or different from both the identity matrix and the switch matrix. We now distinguish two cases:

(i) There exists a transition matrix W_t that is the identity matrix. Then $U_{t+1} = U_t$, and W_{t+1} and all subsequent transition matrices are also the identity matrix. Hence, from point t on the expected payoff for any given state will not change anymore and we have intra-state convergence as required.

(ii) There does not exist a transition matrix that is the identity matrix. We shall prove inter-state convergence in this case (which entails intra-state convergence). W.l.o.g. we analyse the expected payoff of agent 1. Suppose x is its expected payoff for state 1 and y for state 2, at some step of the process. The expected payoffs for the next step are computed as follows:

$$\begin{pmatrix} p & 1-p \\ 1-q & q \end{pmatrix} \cdot \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} y + p \cdot (x - y) \\ x + q \cdot (y - x) \end{pmatrix}$$

Consider the difference $|x - y|$. In the next step, this difference becomes $|(y + p \cdot (x - y)) - (x + q \cdot (y - x))| = |p + q - 1| \cdot |x - y|$. The factor of change $|p + q - 1|$ is of course ≤ 1 . But we can give a better bound. We know that neither $p = q = 0$ nor $p = q = 1$. We also know that p and q must be multiples of $\frac{1}{n}$, where n is the number of agents. (This follows from the rules of the game: each of the n agents has the same likelihood of being the proposer, and each agent will either

propose the *status quo* or the other state with certainty—for the special case of two states no agent will ever randomise between several top proposals). Hence, $|p + q - 1| \leq 1 - \frac{1}{n}$. Now, if x and y are the actual utilities for the two states, t steps before the deadline the difference in expected payoffs can be at most $(1 - \frac{1}{n})^t \cdot |x - y|$. Therefore, as t goes to infinity, this difference must go to 0. This is the case for both agents, meaning that inter-state convergence holds as claimed. \square

Proposition 3. *Any game with two states and a quota of $q < 50\%$ is inter-state convergent.*

Proof. The proof of Proposition 2 shows that, first, no transition matrix can be the “switch matrix”; and second, *if* no transition matrix is the identity matrix (or the switch matrix), then inter-state convergence is satisfied. The only remaining possibility is when there exists a iteration t such that W_t is the identity matrix. We distinguish two cases:

(i) First, assume all agents are indifferent between states 1 and 2 in iteration t . Then we clearly have convergence.

(ii) Otherwise, w.l.o.g., assume agent 1 strictly prefers state 1 in iteration t . Then the only explanation for agent 1 not proposing state 1 when in state 2 is that such a proposal would not make the quota. Hence, as $q < 50\%$ by assumption, at least 50% $(1 - q)$ of the concerned agents (those not indifferent between state 1 and state 2 in iteration t) must prefer state 2. But then, when in state 1, each of these agents would propose to move to state 2 and any such proposal would get accepted, so the probability of moving from state 1 to 2 in iteration t must be > 0 . Hence, W_t cannot be the identity, and we have a contradiction. \square

The bound on the quota in Proposition 3 is tight: the example sketched after Proposition 1 demonstrates that inter-state convergence cannot be guaranteed anymore for quotas $q \geq 50\%$. In the next section, we will analyse our games experimentally, to see to what extent the trends reflected by Propositions 2 and 3 extend to larger games.

4 Experimental analysis

The aim of the experiments is to investigate what parameters affect convergence. In all the experiments we performed, we have always observed intra-state convergence. In the following, we will therefore only report on inter-state convergence (and hence, we write simply “convergence” to denote inter-state convergence).

4.1 Varying the utility range

The three obvious parameters of the experimental study are the number of agents, the number of alternatives, and the quota. Another important parameter concerns the generation of utility functions. Each utility value could be drawn from a continuous or a discrete distribution (over either the interval $[0, 1]$ or a set

of integers $\{0, 1, \dots, u_{max}\}$. Here we focus on uniform distributions (other cases are also interesting and left for future work). The goal of our the first experiment is to answer the question: how does the set of possible utility values affect the rate of convergence?

When a continuous uniform distribution is used to generate the utility values, the preference order on the set of alternatives is strict with probability one. When a discrete distribution is used, however, the agent may have a weak preference order as some alternatives may receive the same values. For a fixed number of alternatives, the smaller u_{max} , the more likely the agents have a weak order. In the extreme case of $u_{max} = 1$, agents have dichotomous preferences (i.e., either they approve or disapprove the alternative).

We fixed the number of states to 15 and varied the number of agents in $\{2 \dots 24\}$. We generated 1000 utility matrices U_0 for each number of agents and checked for convergence for a quota of 50%. The results are provided in Figure 1.

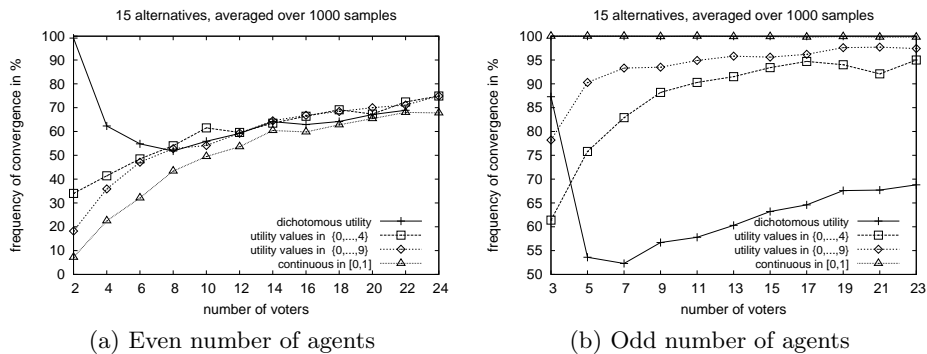


Fig. 1: Frequency of convergence for $q = 50\%$

First, we do not always observe convergence. Convergence is more frequently observed when the number of agents is large. Then, the frequency of convergence is higher when the number of agents is odd than when it is even. For an even number (Figure 1a), there is no significant difference between the different generators, except for very small numbers of agents.¹ For an odd number of agents (Figure 1b), the convergence occurred for all the tested cases with a continuous distribution. Finally, convergence is less frequent when u_{max} is small.

These observations suggest that convergence is more frequent when there are fewer ties between alternatives. When ties are extremely unlikely, we observe high frequency of convergence (i.e., when the number of agents is odd and utilities are

¹ For small numbers of agents and dichotomous utility functions, it is very likely (99% for 2 agents and 15 states) that at least one alternative y is preferred by all agents, which guarantees convergence.

drawn from a continuous distribution; or when the number of agents is large, as exact ties would be required). When ties are likely, however, we observe a lower frequency of convergence (i.e., when the number of agents is even and small and the utility values are drawn from a continuous distribution; or when utilities are drawn from a discrete distribution and u_{max} is small). Experiments with different numbers of alternatives did not alter our conclusions.

4.2 Varying the quota

For two-state games, we have seen that inter-state convergence is guaranteed for quotas $< 50\%$ and can fail for greater quotas. Our second experiment is aimed at checking whether the same trend can be observed for larger games, and at getting an understanding of the likelihood of convergence, when it cannot be guaranteed. We have used 15 alternatives and a population of 100 agents and we have randomly generated 1000 matrices U_0 . The results are shown in Figure 2.

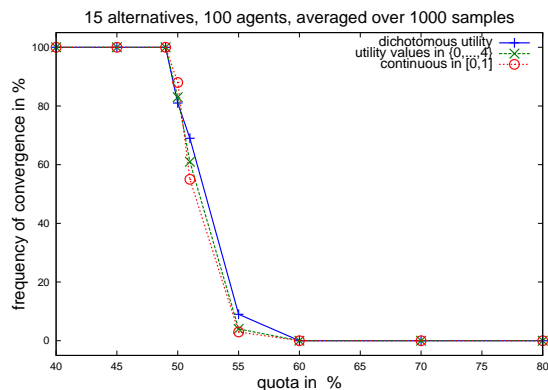


Fig. 2: Frequency of convergence with different quotas

For $q < 50\%$, we always observe convergence, which leads us to conjecture that Proposition 3 generalises to games with any number of states. For higher quotas, we do not always observe convergence (about 80% of the time for $q = 50\%$); and for quotas $q \geq 60\%$, we have never observed convergence in our experiments. Still, even for $q = 100\%$, it is clear that cases satisfying convergence *do* exist (e.g., if all agents are indifferent between all states); such cases are just very unlikely to occur, certainly for our method of data generation.

Maybe the most striking effect we can observe in Figure 2 is the sudden and sharp decrease in the frequency of convergence at the 50% mark. This clearly singles out the majority rule as having a special status within the set of all quota rules, thus suggesting an interesting characterisation of this rule above and beyond the characterisations given by May's Theorem [4] and the Condorcet Jury Theorem.

5 Conclusion

We have proposed a model for iterated voting, where a group of agents make a social choice by implementing a sequence of binary decisions between the *status quo* and an alternative proposal. Each decision is made using a quota rule.

Our model is related to the study of *tournaments* [7] as we can think about our generic approach as a walk in the majority graph. The Markov solution is related to our work as it considers a random walk in the majority graph, and an outcome is in the solution set when it has a positive probability of being the current outcome in the limit. Our model differs by allowing strategic behaviour (*i*) in the choice of the proposed alternative and (*ii*) in the vote (an agent may vote in favor of a less preferred outcome in the short run if this promises a better outcome in the long run). Models closer to ours have been studied in political science, e.g., by Baron [8], although in that model each voter receives a payoff at every time step, while we only ascribe utility to the final state. In the work by Penn [9] another difference is that the challengers are drawn from a given probability density rather than proposed by the agents.

For the case of games with just two states, we have shown that the expected payoff of each agent converges as the deadline increases, when the initial state is fixed. For games with two states and a quota of less than 50%, we have furthermore shown that the expected payoff is independent of the initial state, which offers a level of fairness. Our experimental study shows that this trend generalises also to larger games. For the majority rule, corresponding to a quota of exactly 50%, we have seen that convergence is frequent, but cannot be guaranteed. We have also illustrated how the range of possible utility values an agent may assign to an alternative, and thereby the likelihood of ties, affect convergence.

Future work should be directed towards formulating further conditions under which convergence can be guaranteed, and prediction of a bound guaranteeing that no agent benefits from the choice of the initial state.

References

1. Rosenschein, J.S., Zlotkin, G.: Rules of Encounter. MIT Press (1994)
2. Endriss, U., Maudet, N., Sadri, F., Toni, F.: Negotiating socially optimal allocations of resources. *Journal of Artif. Intell. Res.* **25** (2006) 315–348
3. Lang, J., Xia, L.: Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences* **57**(3) (2009) 304–324
4. May, K.O.: A set of independent necessary and sufficient conditions for simple majority decisions. *Econometrica* **20** (1952) 680–684
5. Sen, A.K.: Collective Choice and Social Welfare. Holden-Day (1970)
6. Grinstead, C.M., Snell, J.L.: Introduction to Probability. 2nd edn. American Mathematical Society (1997)
7. Laslier, J.F.: Tournament Solutions and Majority Voting. Springer (1997)
8. Baron, D.P.: A dynamic theory of collective goods programs. *The American Political Science Review* **90** (1996) 316–330
9. Penn, E.M.: A model of farsighted voting. *American Journal of Political Science* **53** (2009)