A comparison of zero and minimal intelligence agendas in majority-rule voting models

Paul Brewer · Jeremy Juybari ·

Raymond Moberly

October 12 2022

Abstract Emergent behavior in repeated collective decisions of minimally intelligent agents – who at each step in time invoke majority rule to choose between a status quo and a random challenge – can manifest through the long-term stationary probability distributions of a Markov Chain. We use this known technique to compare two kinds of voting agendas: a zero-intelligence

Paul Brewer

Economic and Financial Technology Consulting LLC, Ooltewah TN 37363, USA

E-mail: drpaulbrewer@eaftc.com; corresponding author

Jeremy Juybari

Faster Logic LLC, San Diego CA 92104, USA and Department of Electrical and Computer Engineering, University of Maine, Orono, ME, 04469.

E-mail: jeremy.juybari@maine.edu

Raymond Moberly (deceased)

Faster Logic LLC, San Diego CA 92104, USA

agenda that chooses the challenger uniformly at random, and a minimally-intelligent agenda that chooses the challenger from the union of the status quo and the set of winning challengers. We use Google Co-Lab’s GPU accelerated computing environment to compute stationary distributions for some simple examples from spatial-voting and budget-allocation scenarios. We find that the voting model using the zero-intelligence agenda converges more slowly, but in some cases to better outcomes.

Received: date / Accepted: date

Keywords Emergent Behavior; Zero and Minimally Intelligent Agents; Majority Rule; Spatial Voting; Budget Allocation; Limits of Artificial Intelligence; Markov Chain Applications

JEL classifications:

C63 - Computational Techniques and Simulation Modelling

D71 - Social Choice

Acknowledgements The authors thank [Redacted List] and anonymous [redacted], and also the participants at the Monash University Behavioral Economic Theory seminar and Yale’s Second Conference on Zero and Minimal Intelligence Agents. Errors and opinions are our own. Finally, we acknowledge the untimely death of our co-author Raymond Moberly (July 2021). He was a mentor to Jeremy Juybari for four years, and a lifelong friend of Paul Brewer. Jeremy would like to thank Raymond for all the opportunities and guidance he provided. He will be missed.

Declarations

Preprint

This is a preprint. In general, the results reported in preprints are considered preliminary. Also, a small portion of the content is redacted.

Funding

This project was entirely funded by the authors and their respective companies.

Conflicts of interest/Competing interests

The authors have no conflicts of interest to declare that are relevant to the content of this article.

Data availability

The data are located in a repository that will be made public after peer review.

Code availability

The code for this project consists of a portion that remains private, pending peer review, and a public python module. The public python module may be found at: <https://github.com/drpaulbrewer/gridvoting>

1 Introduction

Everyday, many times per day, robots, AIs, and algorithms (hereafter, simply “bots”) make decisions in place of humans, often faster than humans, and on decentralized information that may need to be aggregated without being passed on to humans. These decisions range across multiple domains: finance, healthcare, military, and advertising. While many paradigms are in use, this paper focuses on a single one – majority rule voting. Should we trust reasonable outcomes to emerge from groups of simple bots operating under majority rule – or should we expect the bots to run amok? What processes should be in place – should the bots consider all feasible choices or restrict consideration to speed up decision-making? Can we characterize the emerging bot behavior even in the absence of human notions of strategy, learning or equilibria?¹ These questions are no longer science fiction.² The answers, while not obvious and impossible to complete in a single study, involve issues at the forefront of modern ethical, economic, political, and regulatory controversies.³

¹ This question is inspired by the editors’ call for papers involving the “Emergence of Macro Phenomena from Interactions among Zero- and Minimal-Intelligence Agents.”

² See the next subsection below for some potential applications.

³ Concerns over the proper role of automation in society and the motives of the humans behind it are expressed through popular criticisms of newsworthy examples (O’Neil, 2016) and controversial regulations adopted by the European Union. General Data Protection Regulation (EU-Parliament, 2016) Article 13(2)(f) requires disclosures of “meaningful information about the logic involved” for profiling and automated decisions. GDPR Article 22(1) then provides EU residents a limited “...right not to be subject to a decision based

To begin, we incorporate into our modelling the following four properties:

(i) an impartial, non-voting, chair-bot determines a voting agenda over a sequence of alternative items; (ii) the alternative items form a finite set, possibly from a discrete space, but not a continuous space;⁴ (iii) bots can communicate and vote much faster than humans; (iv) bots faithfully execute their programming; they vote truthfully and do not exhibit strategic or deceptive behaviors (cryptography can help prevent unauthorized modifications).

Two random agendas are examined: Zero intelligence (ZI)⁵ and Minimal intelligence (MI).⁶ The ZI agenda avoids any element of strategic behavior and picks the next agenda item by uniform randomization over the set of all feasible

solely on automated processing...” A number of other nations are developing their own regulations.

⁴ Because bots involve computer software, natural data models are digitized, finite and pixelated in nature: planar data is stored as a discrete grid of values; real numbers are only approximated by finite-precision floating point numbers; etc. In spatial voting theory, pixelation violates the conditions of proofs that assume continuity or use differential calculus. The pixelated or grid models retain misbehavior such as cycling or spiraling. Also, across many fields of engineering, it is common practice to replace equations defined on continuous spaces with similar equations defined on grids for the purposes of computer simulation.

⁵ The term “zero intelligence” (ZI) originates with a series of papers by Gode and Sunder (1993a,b, 1997) that examined the market behavior of bots who act randomly but consistent with budget and self-interests. They found highly efficient market allocations and prices eventually similar to competitive equilibrium theory.

⁶ The term “minimal intelligence” (MI), as used in the online 2019-2020 ZI MI conferences sponsored by Yale University, describes a variety of strategies that do not apply optimization but instead can be described by simple rules-of-thumb – such as trying to win an auction by bidding the price up by \$1 on successive bids.

items. The MI agenda is slightly “smarter” than ZI in that it randomizes only over the current status quo and items that would immediately win versus that status quo. Unlike the ZI agenda, the MI agenda never picks alternatives that would be rejected by the voters. Both random agendas have the advantage of being impartial but the disadvantage of eventually including sequences that are well known to cause poor outcomes. Neither agenda is based on any historical trends or previous votes.

Next, we embed bots into three well known problematic, but minimal, sequential voting environments: (E1) the three item Condorcet cycle, (E2) spatial voting with Euclidean preferences, and (E3) divide-the-dollar budgeting. Each of these minimal environments requires only three voters. The intent is to take a first step in stress-testing any majority-rule bot-based system.⁷

The combination of the sequential voting environments – together with the ZI or MI random agendas – yield sequences of outcomes that are Markov chains. The random agendas we define are similar to classes of random challenger dynamics considered by Ferejohn et al. (1984, p.49). They were able to establish bounds on the probability of reaching extreme outcomes. Instead of rough bounds, we calculate stationary probability distributions on high-performance graphics processing units (GPUs).⁸ In this way we can explore

⁷ Precedents exist in the test-bedding of new kinds of smart markets (Plott, 1994) and in the random “chaos monkeys” created at Netflix Inc for anti-fragile systems management (Tseitlin, 2013).

⁸ We did find the limits of easily available hardware. Analyzing a spatial voting scenario on a 2-D integer grid $[-80, \dots, 80] \times [-80, \dots, 80]$ has grid size 161×161 or 25,921 individual states; and will involve constructing a $25,921 \times 25,921$ transition matrix totaling about

various aspects of long-run performance. We can also see that convergence to the “long-run” may be too long for evaluating human political systems that act on a scale of days to years, but perhaps still reasonable for bots that act on a time scale of minutes or seconds.

Will the bots run amok? In the short-run the answer is “yes - at times” (the possibility is embedded into the Markov chain’s transition matrix through the properties of the agendas) but the long-run answer appears to be “no.” The long-run average outcomes are not extreme. Extreme outcomes will occur but with low probability. In operation the ZI and MI agendas can yield different stationary distributions. The MI agenda can introduce additional variance, additional inequality or increase the likelihood of suboptimal outcomes when compared with the ZI agenda.

To provide expectations about this article and appeal to a diverse audience, we briefly discuss the use of mathematics and technical jargon. The technical sections are not optional for a full understanding, but the major points should be clear from a quick scan. We prove no new theorems and assume, in describing this research, an engineering-level undergraduate introduction to mathematical notation, finite sets, functions, limits, basic probability, vector and matrix algebra. For a first reading it is possible to skim or skip, e.g. skim-

671 million entries. Finding the stationary distribution will involve taking higher and higher powers of that transition matrix involving trillions of floating-point multiplications and then checking rows for convergence. Achieving this grid size in 2021 helps explain how Markov chain techniques are only now becoming practical for the analysis of moderately-sized spatial voting scenarios.

ming the mathematical modeling and procedures of Sections 2 and 3, skipping 3.3 (Validation), reading Section 4 on the Condorcet Cycle and the subsequent sections. We try not to assume prior training in voting theory, Markov chains, numerical techniques, computer programming, robots, algorithms, or artificial intelligence. However, some familiarity with the jargon will be helpful. It would be difficult to define or remove all the technical jargon of these various fields without evolving this article into a book. Readers who want to take on the challenges of new research can find our raw data and code in a public online repository.

1.1 Voting-like and Voting-based technologies

This subsection provides a few examples of multi-bot systems incorporating voting-like and voting-based technologies.

Voting-like technologies have been used for guidance control for the NASA space shuttle missions, through parallel outputs with pilot override (Kent and Williams, 1988) – and as part of fault tolerance in cloud computing clusters (Kubernetes blog contributors, 2016). These technologies include some of the steps of voting but actual decisions are made by others (the shuttle pilot) or by randomization and shared deterministic rules (when replacing faulty servers).

Voting-based technologies use votes to reach some kind of decision. For instance, ensemble learning is when multiple algorithm outputs are combined, with a voting strategy such as majority or plurality voting, to yield more accurate and robust predictions than a single algorithmic output (Raschka, 2015).

There are other approaches as well. Examples from the research literature include:

- ultrasonic sensors voting to create images (Barshan, 2001)
- bots voting to make financial credit risk decisions (Yu et al., 2009)
- a majority-voting classifier for improved cancer detection (Paul and Iba, 2009)
- bots voting on behalf of students to schedule courses (Aseere et al., 2010)
- codewords voting to summarize video content (Oikonomopoulos et al., 2010)
- a majority crossover operator for genetic algorithms (Friedrich et al., 2016)
- majority voting in a fleet of autonomous vehicles (Teixeira et al., 2018, 2020)
- voting-based multi-agent reinforcement learning (Xu et al., 2020)
- a faster method to search over abstract spaces (Farsa et al., 2021)

This list illustrates how many areas of active research and development are dependent on understanding the capabilities and limitations of bots operating under majority-rule.

Future applications could more closely follow spatial voting models with bots having private data or preferences over some space. These could range from groups of bots cooperating in 2-D or 3-D physical environments (exploration, anomaly detection) to abstract environments with many dimensions (e.g. credit card fraud detection; Turing-test sock-puppet detection in chat rooms). Such abstract environments are often reduced to a few dimensions through mathematical techniques such as principal component analysis (James et al., 2013, pp.228-241).

1.2 Overview of Voting Theory

The literature on voting theory is large. It encompasses not only political science but law, philosophy, history, economics, game theory and other subjects.⁹ It is also peppered with mathematical theorems and demonstrations concerning the fundamental capabilities of majority rule. We will cite a few of the major results to highlight a path.

Mathematical and procedural results have the advantage of an accepted internal consistency or proof existing outside of the human historical milieu and therefore constrain all systems utilizing voting, whether populated by humans or bots. Intriguing disappointments arise in the form of “paradoxes” (Condorcet et al., 1785) or intransitive cycles where voting picks A over B , B over C , but C over A ; and impossibility theorems (Arrow, 1950; Gibbard, 1973; Satterthwaite, 1975; Gibbard, 1977; Schwartz, 1982; Duggan and Schwartz, 2000) that demonstrate potential issues with all “reasonable” voting schemes. These issues include cycling, indecisiveness, sensitivity to “irrelevant” changes, vulnerability to strategic manipulation, and potential dictatorship. The implication is that both the aggregation of decentralized preference information through majority rule and resulting voting decisions can be problematic.¹⁰

⁹ While a full survey is beyond the scope of this paper, McKelvey et al. (1990) have a thorough review of early spatial voting models and experimental literature. Duggan (2005) reviews more recent models. Aldrich et al. (2007) contains a partial collection of McKelvey’s works mixed with contributed chapters from numerous colleagues.

¹⁰ In contrast, the economic literature strongly suggests that, ignoring phenomena like pollution or device compatibility, markets successfully process decentralized information

Two of the early results on majority rule are encouraging. In addition to the voting paradox, Condorcet et al. (1785) also provides a “jury theorem” demonstrating information aggregation and noise reduction when majority rule is applied to binary decisions (yes or no, guilty or not-guilty). Black (1948) provides a “median voter theorem,” establishing that a stable equilibrium (or *core*) exists at the median of voters’ ideal points when voters with single-peaked preferences make choices over a single dimension. These features of binary voting have become popularized and renamed to fit modern applications.¹¹

Attempts to extend the early successes encountered logical difficulties. Still, Tullock (1967, p.256) informally explores extension to two-dimensional continuous spaces. He suggests that the difficulties implied by Arrow’s theorem are “phantoms” that would disappear with a sufficiently large number of voters (999,999). He proposes that a central tendency would emerge and that cycling consistent with Arrow’s theorem would exist as a microscopic detail (p.260) unless the choice space is a grid (p.267) where cycles might be “apt to be important when they occur.”

through price-formation and trade (Hayek, 1945; Smith, 1962). Collections of simple ZI agents can do so (Gode and Sunder, 1993a). We do not further compare voting with trade, but trade is highly relevant to the early and continuing ZI/MI agent literature.

¹¹ Especially: *crowdsourcing*, *collective intelligence* or *the wisdom of the crowd*. The benefits of binary majority-rule are operative in popular public forums through up-votes and down-votes and in several of the voting-based technologies of the previous subsection.

Using differential calculus, Plott (1967) showed that extension to two or more dimensions would not yield a stable outcome unless a rare alignment occurs – the outcome must be located at some voter’s ideal point and a rare symmetry of voter utility gradients must exist pitting all remaining pairs of voters directly against each other. The likely absence of equilibrium led to further attempts to characterize the dynamics. Subsequently, (McKelvey, 1976; Schofield, 1978) proved that an agenda exists connecting any point to any other point.¹² These results are sometimes referred to as the *chaos theorems*, “despite [McKelvey’s] admonitions to the contrary” (Ordeshook, 2007, p.165) and other criticisms (Cox and Shepsle, 2007). McKelvey also co-authored Ferejohn et al. (1984), establishing a connection between voting cycles, instability, and the convergence of continuous-space Markov processes. This placed probabilistic bounds on the outcomes of cycling under random agendas.

Our paper is inspired by the intersection of this previous work with the results of Gode and Sunder (1993a,b, 1997) in market trading scenarios that show that seemingly over-simplified agents and processes partially reproduce complex, human, market-level behavior. In contrast, the voting literature con-

¹² McKelvey(1976)[section 3] is quite clear: “when transitivity breaks down, it completely breaks down, engulfing the whole space in a single cycle set”; “The existence of a single cycle set implies that it is possible for majority rule to wander anywhere in the space of alternatives.”; “The likelihood of this occurring is strongly dependent ...[on which rules]... generate the agenda”; “It follows that if any one voter, say the *Chairman*, has complete control of the agenda, that he can construct an agenda that will arrive at any point in space...” He also provides caveats and limitations.

tinues to produce models introducing additional complexities – such as voters who deviate from truthful voting to act on long-term strategies, and voters who randomize based on noisy observations. Like Gode and Sunder, we wonder if the overly-simplified approach might be fruitful in generating some aspects of emergent behavior. This is part of why we choose to investigate myopically truthful voter behavior with explicit, simple, preferences and finite sets (grids) of feasible alternatives.

A continuum of choices is not necessary because the outward spiraling behavior possible in 2-D spatial voting models involves a sequence of large discontinuous jumps supported by constantly shifting alliances. Such agendas can also exist on a grid.

Myopically truthful voters are very simple and – as Gode and Sunder have emphasized with their ZI market agents – do not engage in a list of more complex human behaviors such as optimizing, strategizing, utilizing history, or predicting others' action. Brewer et al. (2002, p.205) point out that such simplifications result in models that conveniently lack complicating parameters. Additionally, the myopically truthful voting agents do not randomize, or even make proposals. In contrast, when Kalandrakis (2004) and Penn (2009) assume agents make optimized self-interested budget proposals which are then voted on by the group, agents can sometimes simply take turns proposing their individual ideal points (Kalandrakis) or minimal-coalition splits (Penn).

Dougherty and Edward (2012) has several similarities to our paper (randomness, sequential voting, two dimensions,¹³ examining the probability of obtaining an outcome in the Pareto set) and also credits partial inspiration to the work of Gode and Sunder. Their focus is on comparing unanimity and majority rule. The differences are numerous and include the research questions, agendas, types of environments, number of voters, and methodologies.

The remainder of the paper is organized as follows: Section 2 describes the voting model and two types of random agendas, ZI (Zero Intelligence) and MI (Minimal Intelligence). Section 3 describes Markov chains for our models and discusses the mathematical requirements and computational procedures to obtain a unique stationary probability distribution over outcomes. Section 4 discusses (E1) the three-item Condorcet cycle modelled as a Markov chain. Section 5 provides (E2) an example with outcomes on a spatial grid, where each of three voters have ideal points and vote to minimize the distance of the outcome state to their ideal point. Section 6 provides (E3) a spatial grid example where the three voters divide a fixed budget. The conclusions are in section 7.

2 Models, transitions, and labelling

Assume a finite set of voters $\mathbf{V}$ who have utility functions over a finite set of feasible alternatives $\mathbf{F}$. Utility functions exist for each voter to indicate

¹³ An earlier paper (Dougherty and Edward, 2005) considers one dimension. We thank Keith Dougherty for bringing this research to our attention.

preferences. Ties may exist in a voter's preferences. Voters are myopic and vote truthfully.

At each step in time $t = 0, 1, 2, \dots$ a transition occurs. To describe transitions, it is useful to let i and j be indexes of the feasible set $\mathbf{F}$. There is a vote between the status-quo i in effect at time t , $i = f[t] \in \mathbf{F}$, and a challenger $j = c[t] \in \mathbf{F}$. The transition results in either $f[t+1] = f[t] = i$ or $f[t+1] = c[t] = j$. The seemingly trivial “no challenge” case (where $j = c[t] = f[t] = i$) is possible, maintains the status quo, contributes to the diagonal of the transition matrix, and helps guarantee convergence of the Markov chains defined in the next section.

The winner set $\mathbf{W}(i)$ is the set of challengers that would strictly beat i under majority rule. By definition, $i \notin \mathbf{W}(i)$. Let $||\mathbf{F}||$ be the number of feasible alternatives. Similarly, $||\mathbf{W}(i)||$ is the size of the winner set at i .

Next we will define two different kinds of random agendas and provide their transition matrices.

The *ZI agenda* chooses $c[t]$ at random from the uniform distribution on $\mathbf{F}$. The transition matrix $P_{\text{ZI}} = (p_{i,j})$ will be zero where j loses to i ; $1/||\mathbf{F}||$ where j beats i ; and $1 - ||\mathbf{W}(i)||/||\mathbf{F}||$ on the diagonal. Note that each row i is a discrete probability distribution. Therefore for each i , $\sum_j p_{i,j} = 1$.

When *zero intelligence* is used to choose the challenger, the probability that the status quo will be maintained is higher relative to the next agenda we introduce immediately below.

The *MI agenda* chooses $c[t]$ at random from $i \cup \mathbf{W}(i)$. Thus, the minimal intelligence agenda can preserve the status quo, but is more likely to yield a change when many possible winning challengers exist.

Here the transition matrix $P_{\text{MI}} = (p_{i,j})$ will be zero where j loses to i and $1/||1 + \mathbf{W}(i)||$ where j beats i and also along the diagonal. Again, each row sum is $\sum_j p_{i,j} = 1$.

We also note that the status-quo preserving diagonal elements for the two agendas have different ranges in the absence of a core. The minimum diagonal value for either transition matrix is $1/||\mathbf{F}||$ in the case of a dominated alternative. When a core exists the maximum diagonal value is 1 for either agenda. In the absence of a core the maximum diagonal would be the case of an alternative that loses to only one other alternative. For P_{ZI} this is $1 - 1/||\mathbf{F}||$ and is close to 1. However, for P_{MI} that diagonal value is $1/2$.

A *coordinate labelling* of the alternatives is used to embed $\mathbf{F}$ in some larger space such as the integer grid $\mathbf{Z}^2$ or real plane $\mathbf{R}^2$ and perform relevant transformations for spatial models. For example, to embed into two dimensions we can define two functions $x(f), y(f)$ that provide the x and y coordinates of each feasible alternative. Implicit transformations then occur when considering voters' utility functions over such a space. This allows some methodological separation, as preferences on a space become mapped to preferences on the finite set $\mathbf{F}$. We choose Python's convention that a list of N items is indexed from 0 to $N - 1$. We label grids in typewriter convention with alternative 0 at the top left.

3 The Finite Markov Chain methodology

The goal of this section is to provide concepts and methodological details necessary for readers to evaluate or replicate the research. We state some of the results of the mathematics of Markov chains and apply it to the model of the previous section.

While lacking the rigor and theorem proving that some readers might prefer, we provide informal exposition hoping to make these techniques accessible to a broader interdisciplinary audience. A Markov chain is characterized completely by a set of states, a current state and a transition matrix that provides the probability of changing states. Based on this information, some Markov chains yield stationary distributions that can describe long-term expected outcomes and averages irregardless of the initial state. This could allow estimating the average climate from observed frequencies of weather transitions (a common pedagogic example), or estimating the viewership of web pages from their links (Google PageRank).

3.1 Theory of Operation

The notation and terminology in this section adapts those from the modern graduate textbook by Levin et al. (2009) on Markov chains and connects some relevant ideas. We also found Häggström et al. (2002) and the open source textbook by Grinstead et al. (2006, pp. 438-439, 447-451) useful. In particular, this article borrows the notation for distributions π , transition matrices $P(i, j)$,

time evolution through left-multiplication πP , and a convergence theorem. We use square brackets for time, e.g. an alternative at time t , $f[t]$; or a probability over future outcomes at time t , $\pi[t]$.

In our model, the sequence of feasible alternatives $f[0], f[1], f[2], \dots$ is a finite Markov chain as the evolution of the sequence depends only upon the initial $f[0]$ and the transition matrix P ; and P is completely determined by the agenda (ZI or MI) and voter preferences.

Predictions about a Markov chain's future behavior can be determined solely from a probability distribution $\pi[0](f)$ over the initial value $f[0]$, and the transition matrix P . The ex-ante ($t = 0$) probability vector over f at $t = 1$ is determined by multiplying the $t = 0$ probability vector over f by the transition matrix P . For example: $\pi[1] = \pi[0]P$; $\pi[2] = \pi[1]P = \pi[0]P^2$; $\pi[3] = \pi[0]P^3$; (Levin et al., p.5, *t-step transition probability*).

We use the term *ex-ante* to emphasize that $\pi[t]$ is best interpreted as a prediction – a future probability based on the probability $\pi[0]$ reflecting present $t = 0$ conditions. When the process evolves from $t = 0$ to $t = 1$, the voters will have voted and chosen a specific alternative. The various short-term predictions based on the previous state are no longer valid but must be recalculated given the new state. Further evolution occurs randomly via the transition matrix P – and specifically, the row for the current status quo. However, a long-term prediction does converge and decouple from its past.

Under the conditions Levin et al., p.7 call *irreducibility* and *aperiodicity*, Theorem 4.9 (pp. 52-53) states that every row of P^t converges to π^* as $t \rightarrow \infty$.

This occurs with an exponential decay in the sum of absolute differences – the L1 norm they call total variation. This can be restated as the distributions $\pi[1], \pi[2], \dots$ converging (Häggström et al., 2002, Theorem 5.2) such that

$$\lim_{t \rightarrow \infty} \pi[t] = \lim_{t \rightarrow \infty} \pi[0]P^t = \pi^*$$

Because every row of P^t converges, π^* is insensitive to changes in $\pi[0]$.

π^* is called the *stationary distribution* of the Markov chain. π^* uniquely satisfies $\pi^* = \pi^*P$ (Levin et al., p.13, 17) and is therefore unaltered or *stationary* with respect to time evolution via the transition matrix P .

Irreducibility is satisfied when, for any pair i, j of alternatives there exists a power $k > 0$ such that $P^k(i, j) > 0$. In other words, every feasible alternative must be reachable from every other in some number of steps.

The McKelvey and Schofield global cycling results for 2-D spatial voting – that in the absence of a voting equilibrium, some majority rule agenda connects any feasible alternative to any other alternative – is also a claim of irreducibility. But those previous results were proven on a continuum.

Our models exist on finite integer 2-D grids. With an odd number of voters, there are no tie votes. It is then difficult to see how majority rule, in the absence of a core, could result in a Markov chain consisting of isolated independent subchains. We therefore posit a single irreducible cycle of alternatives satisfying the irreducibility condition.

Alternatively, irreducibility is not satisfied when there is a stable voting equilibrium or core. Recall that an alternative in the core beats all non-core alternatives. The Markov chain will therefore get “stuck” in the core set, con-

sistent with voting theory. All non-core outcomes are *transient* in the sense that the predicted future probability should shrink to zero. We can detect core-points in the Markov chain by looking for diagonal elements of P equal to 1.

When there is no core, there may still be transient alternatives under majority rule that only lose votes to other alternatives and never win over any other alternatives. These transients do not affect the stationary distribution and can be ignored.

Aperiodicity is satisfied when the greatest common denominator of the cycle lengths in the chain is 1. This can be trivially satisfied if all diagonal elements of the transition matrix are strictly positive, i.e. the agenda must always provide a positive one-step probability of remaining at the status quo. Aperiodic chains can be created as *lazy versions* of periodic chains by simply adding a probabilistic do-nothing step that maintains the status quo (Levin et al., p.8). The MI agenda is a lazy version of a rule that always randomly yields some winning challenger. Both the ZI and MI model agendas have positive diagonal elements in the transition matrices and satisfy aperiodicity.

3.2 Calculation

Various methods exist for practical calculation of Markov chain stationary distributions π^* . To easily replicate calculations in Google Co-Lab Pro,¹⁴

¹⁴ Google Co-Lab Pro is a low-cost (US \$10/mo) service for running analysis notebooks in Python. Servers are shared, have time limits, and include a GPU (e.g. Nvidia K80, P100,

(Redacted) implemented the two methods below in a public Python module *gridvoting* (Redacted Author Name, 2021) by calling on the matrix-math modules *cupy* (Preferred Networks Inc, 2015) and *numpy* (Numpy Developers, 2006).

Method 1: Raise P to larger and larger powers $N = 2^K$ by squaring the matrix at each step K . Watch for the rows of P^N to converge. Our convergence criteria examines the difference in the first row and middle row. We expect these two rows to be initially very different, and then converge in the sum of absolute differences (i.e, the L1-norm) to a tolerance value less than 10^{-10} . This is an arbitrary choice of tolerance, but is reasonable given that 64-bit floating-point provides a mantissa resolution of 16 decimal digits. Once that occurs, method 1 calculates π^* as the average of all rows of P^N , or

$$\pi_j^* = \frac{\sum_i (P^N)_{i,j}}{||\mathbf{F}||} \quad (1)$$

Method 2: Solve $\pi^* = \pi^* P$ by solving the transposed equations

$$(P^\top - I)\pi^* = 0$$

together with the scale condition

$$\sum_{f \in \mathbf{F}} \pi_f^* = 1$$

T4 with 12-16 GB of GPU memory). Co-Lab’s homogeneous architecture increases the likelihood of successful replication by others. Larger GPUs (e.g. 80GB Nvidia A100) were available elsewhere as hourly rentals.

This leads to $\|\mathbf{F}\| + 1$ conditions on $\|\mathbf{F}\|$ unknowns. The scale insensitivity of the first set of conditions suggests eliminating a superfluous degree of freedom. A well-posed problem can be constructed via various ad-hoc techniques (Grinstead et al. (2006, pp.436-437) contains two such examples). Method 2 replaces of the first row of $(P^\top - I)$ with the scale condition, a row of all ones, obtaining a new matrix Q . The scale condition's right hand side is a 1, and the remaining right hand sides are all 0, yielding a basis vector e_0 for the right hand side. A matrix inversion may then yield a result.

$$\pi^* = Q^{-1}e_0 \quad (2)$$

Employing both calculations (1) and (2) can help identify programming bugs or numerical issues. With large grids it will also exhaust GPU memory. In such cases we resort to the power method alone.

3.3 Validation

Here we provide details of how calculations reported in the subsequent sections are double checked for salient errors. It seems apropos to discuss such validation here, at the general level, instead of within each subsequent example.

The same diagnostic expressions are calculated for each environment we examine, and can be found in Table 1 and Table 2.

Table 1 Spatial Voting: Diagnostics for calculation of π^*

grid	agenda	$ \mathbf{F} $	power	$ \pi - \pi P _1$	$\sum \pi - 1$	$ \pi_{\text{method 1}} - \pi_{\text{method 2}} _1$
80	ZI	25921	4096	2.861e-15	-1.386e-11	NA
80	MI	25921	64	1.752e-14	-1.987e-14	NA
60	ZI	14641	4096	9.399e-16	-3.735e-12	3.783e-12
60	MI	14641	64	8.549e-15	-5.551e-16	2.698e-14
40	ZI	6561	1024	1.113e-14	-6.815e-13	7.59e-13
40	MI	6561	64	3.663e-15	0	6.942e-15
20	ZI	1681	256	3.072e-14	-1.109e-13	2.967e-13
20	MI	1681	32	4.892e-16	-2.22e-16	2.358e-15

Table 2 Budget Voting: Diagnostics for calculation of π^*

agenda	$ \mathbf{F} $	power	$ \pi - \pi P _1$	$\sum \pi - 1$	$ \pi_{\text{method 1}} - \pi_{\text{method 2}} _1$
ZI	5151	64	1.35e-14	-2.56e-13	2.56e-13
MI	5151	32	1.75e-15	2.22e-16	2.53e-15

As the tolerance for the Method 1 convergence criteria was 10^{-10} we would like for diagnostics involving summed differences to also be at or below this level.

Each diagnostic table reports model particulars in the leftmost columns: the random agenda (ZI or MI), and, for the spatial voting models, grid size – together with the number of feasible alternatives $||\mathbf{F}||$). Successive columns report the diagnostics.

The *power* is the power N where P^N first satisfied the convergence criteria of calculation Method 1. In our context, the transition matrix to the power N itself describes a transition. It is the transition to N steps in the future, calculated from all possible agendas of length N . Note that in section 2 we showed that the Markov chains using the ZI agenda have higher transition matrix diagonals than their MI-agenda counterparts and therefore should require higher powers for convergence.

Different diagnostics expose different deviations from an ideal. Each is a sum or a sum of absolute values of differences and should, ideally, be zero.

Values around 10^{-16} may represent single errors at the 16-digit limits of 64-bit floating-point numbers of order 1. Higher numbers represent accumulations of such errors over large numbers of operations. For example, values around $\|\mathbf{F}\| * 10^{-16} \approx 10^{-13} - 10^{-12}$ could represent 16th digit errors in each alternative summed over the number of alternatives. In our opinion, neither these levels nor levels a bit higher indicate cause for alarm.

$\|\pi - \pi P\|_1$ examines how well the calculated stationary distribution satisfies the fundamental stationary equation it is supposed to solve. This diagnostic is the L1-norm, or sum of component absolute values, of the difference between the calculated stationary distribution π and its one-step future evolution πP . The values range from about 10^{-16} to 10^{-14} .

$\sum \pi - 1$ examines if the calculated stationary distribution is a proper probability distribution that sums to 1. The discrepancies range from zero to about

10^{-11} . Renormalization to exactly one is not possible because of the limited precision of 64-bit floating-point numbers. Finally:

$||\pi_{\text{method 1}} - \pi_{\text{method 2}}||_1$ examines the difference in stationary distributions that result when using the two different methods of the previous subsection: method 1 (the power method) vs. method 2 (the linear algebra method). Again the L1 norm is used. None of the values are particularly alarming. The values for the first two rows of Table 1 are not reported because method 2 exhausted available GPU memory in Google CoLab Pro. Values were obtained in additional replications on other, larger GPUs (Nvidia A100 on Google Cloud Platform; A6000 on Lambda Labs). These values were $6.6 * 10^{-12}$ and $1.4 * 10^{-11}$ for (row 1) 80 ZI and $2.7 - 2.9 * 10^{-14}$ for (row 2) 80 MI.

4 Example: Condorcet Cycle

In this section, we will demonstrate the application of our ZI and MI random agenda models in the relatively simple Condorcet Cycle environment. In this case, the outcomes are identical stationary distributions and the only difference is speed of convergence. This simple example may prove useful to interdisciplinary researchers who may not have previously encountered it, and provides an opportunity to gain intuition with transition matrices and solution techniques.

4.1 Environment E1

Table 3 Individual preferences that yield a cycle

Voter	A	B	C
1	best	middle	worst
2	worst	best	middle
3	middle	worst	best

The profile in Table 3 is well known to lead to an intransitive majority rule result $A > B; B > C; C > A$. This follows from examining each two-voter coalition: Voters 1 and 3 prefer A to B. Voters 1 and 2 prefer B to C. But voters 2 and 3 prefer C to A.

4.2 Transition Matrices

Recall from section 2 that the ZI agenda chooses a challenger uniformly at random from the three feasible alternatives A, B, C . Since only one of these can win, the probability of transitioning to the winning challenger is $1/3$ and the probability of maintaining the status quo is $2/3$.

$$P_{ZI} = \begin{bmatrix} \frac{2}{3} & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{2}{3} & 0 \\ 0 & \frac{1}{3} & \frac{2}{3} \end{bmatrix}$$

In contrast, the MI agenda chooses a challenger uniformly at random from the status quo and the (single) winning challenger. Therefore the probabil-

ity of transitioning to the winning challenger is $1/2$ and the probability of maintaining the status quo is $1/2$.

$$P_{\text{MI}} = \begin{bmatrix} \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

4.3 Stationary Distributions

Here we briefly review methods for calculating the stationary distributions.

4.3.1 *Guess and Verify*

The simplicity and symmetry of the Condorcet example suggests guessing a uniform distribution as the stationary distribution, i.e. $\pi_{\text{ZI}}^* = \pi_{\text{MI}}^* = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$.

To verify, it is necessary and sufficient that π^* satisfies $\pi^* = \pi^* P$. For P_{ZI} and P_{MI} this is straightforward. The matrix multiplication of $[\frac{1}{3}, \frac{1}{3}, \frac{1}{3}] * P$ averages the three row vectors comprising any 3×3 P -matrix. But each column of either P_{ZI} or P_{MI} sums to one ($\frac{2}{3} + \frac{1}{3} + 0$ for ZI, or $\frac{1}{2} + \frac{1}{2} + 0$ for MI), so the resulting row vector is $[\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$.

4.3.2 *Method 1: Powers*

Method 1 for finding a stationary distribution is to take higher and higher powers of the transition matrix until a convergence criteria between rows is satisfied. To balance exposition vs space, we briefly examine P^8 (rounded to 3 digits):

$$P_{ZI}^8 = \begin{bmatrix} 0.329 & 0.342 & 0.329 \\ 0.329 & 0.329 & 0.342 \\ 0.342 & 0.329 & 0.329 \end{bmatrix} \quad \text{and} \quad P_{MI}^8 = \begin{bmatrix} 0.332 & 0.332 & 0.336 \\ 0.336 & 0.332 & 0.332 \\ 0.332 & 0.336 & 0.332 \end{bmatrix}$$

The rows of P_{MI}^8 are more similar than the rows of P_{ZI}^8 and closer to $[1/3, 1/3, 1/3]$. This is due to faster convergence for the expectations under the MI agenda as compared to the voting outcome distribution under the ZI agenda.

4.3.3 Method 2: Linear Algebra

In method 2, Q matrices are calculated by transposing the relevant P matrix, subtracting the identity matrix I, and then replacing the first row with ones:

$$Q_{ZI} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & -\frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & -\frac{1}{3} \end{bmatrix} \quad \text{and} \quad Q_{MI} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & -\frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & 0 & -\frac{1}{2} \end{bmatrix}$$

We calculate the matrix inverses in python and then obtain stationary distributions from the formula (2):

$$\pi_{ZI}^* = Q_{ZI}^{-1} e_0 = \begin{bmatrix} \frac{1}{3} & 1 & 2 \\ \frac{1}{3} & -2 & -1 \\ \frac{1}{3} & 1 & -1 \end{bmatrix} * \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{bmatrix}$$

and

$$\pi_{\text{MI}}^* = Q_{\text{MI}}^{-1}e_0 = \begin{bmatrix} \frac{1}{3} & \frac{2}{3} & \frac{4}{3} \\ \frac{1}{3} & -\frac{4}{3} & -\frac{2}{3} \\ \frac{1}{3} & \frac{2}{3} & -\frac{2}{3} \end{bmatrix} * \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{bmatrix}$$

4.4 Consequences of ZI and MI Agendas

In this example, the MI voting agenda and ZI voting agenda have the same stationary distribution. They differ in speed of convergence. Because of the higher probability of retaining the status quo, the ZI agenda is slower to converge to the stationary distribution.

5 Example: Spatial voting with a triangle of ideal points

In this section, we will demonstrate the application of our ZI and MI random agenda models in a somewhat more complex environment.

As in the previous section, the ZI agenda will show slower convergence.

Unlike the previous section, the choice of a ZI or MI agenda will yield differences in the stationary distribution. The most obvious difference will be an increase in the expected variance of outcomes, with MI having higher variance than ZI at their respective stationary distributions.

The voting environment will involve a series of 2-D (x, y) points chosen from a finite integer grid. In tables we will report results for grids of *grid limit* 20, 40, 60, and 80 determining the maximum values of $|x|$ and $|y|$. In figures, we only provide results for grid limit 80, as the smaller grids (except for 20) are similar in most respects.

Bots in this environment could be searching for an item and voting for the next location to search, or voting for the placement of additional resources, or some other 2D activity where they have private information and/or preferences.

5.1 Environment E2

Fig. 1 Euclidean utility functions for voters 1,2,3 with ideal points at $(x=-15,y=-9)$, $(x=0,y=17)$, and $(x=15,y=9)$ (± 80 grid)

Here, each voter has their own ideal point located on a 2-D integer grid: voter 1's ideal point is at $(x = -15, y = -9)$, voter 2's ideal point is at $(x = 0, y = 17)$ and voter 3's ideal point is at $(x = 15, y = -9)$. Each voter prefers points that are closer to their ideal point to those that are further away. This results in circular indifference curves for each voter, based on Euclidean distance from the ideal point, such as shown in Figure 1.

Figure 2 provides an additional polar coordinate system (r, θ) based on the center of the triangle. Here, $r = 0$ denotes the centroid of the triangle at $(x = 0, y = -1/3)$. The angle θ , in radians, is measured at the centroid and

Fig. 2 Triangle of ideal points in cartesian and polar coordinate systems. The center of the polar coordinates ($r = 0$) is placed at the triangle centroid ($x = 0, y = -1/3$).

is zero in the upward direction along the y-axis and at voter 2's ideal point $(x=0, y=17)$ and approximately $\pm 2\pi/3$ at the other voters' ideal points.

The voters' ideal points are on a near-equilateral triangle in order to make voters' situations as equivalent as possible while maintaining separation. Unfortunately, perfectly equilateral triangles cannot be constructed on 2D square

grids.¹⁵ The provided triangle has a slight asymmetry that can be found by comparing the triangle's horizontal bottom (length 30) and the sloping sides (length $\sqrt{901} \approx 30.017$).

As mentioned in the introduction, the possibility of outward-spirals on the grid as well as cycles is well-established in previous work. To see how an outward-spiral might begin in the current environment, notice that the voters at the bottom of the triangle, voters 1 and 3, both have ideal points where $y = -9$ and thus both prefer alternatives $(x = 0, y = -1) \dots (x = 0, y = -17)$ to the grid center alternative $(0, 0)$. Further steps in a spiral would entice the other two-voter coalitions.

5.2 Transition Matrices

The transition matrices for the ZI and MI agendas are constructed as described in section 2. For the “ ± 80 ” grid, the grid includes $161 \times 161 = 25,921$ alternatives. The resulting $25,921 \times 25,921$ transition matrices cannot be compactly visualized in a way that relates to coordinates x and y on the grid.

However, the diagonal terms are simple to interpret and visualize. They provide the retention probabilities for each grid point: the probability of remaining at that point when it is the status quo instead of moving to another alternative.

¹⁵ Each angle in an equilateral triangle is 60° or $\pi/3$ radians. Constructions on grids fail due to the irrationality of $\tan(\pi/3) = \sqrt{3}$. For further discussion, see Parsons and Truran (1970) and Beeson (1992).

Fig. 3 $\log_{10}$ diagonal elements of P_{ZI} and P_{MI} vs radial distance (± 80 grid). At each r , the probability of remaining at the status quo is much greater with ZI than with MI.

Figure 3 shows the diagonal terms of the transition matrices as a function of radius. The thickness of the plot represents the variability at different angles. Differences in transition matrix diagonal values are clear, both in the values and rates of decline as one moves further from the center.

The ZI agenda transition matrix diagonal values generally exceed those from the MI agenda. Diagonal values of P_{ZI} range from near $10^0 = 1.0$ at the center down to around $10^{-2.5}$ at the grid corners. Diagonal elements of P_{MI} range from a high value below 10^{-2} at the center to a low value around $10^{-4.4}$ at the corners.

The shape of the radial plot of $\log_{10} P_{ZI}(i, i)$ is concave while the radial plot of $\log_{10} P_{MI}(i, i)$ is convex. Diagonals of P_{ZI} initially decline very slowly

with radius, then increase the rate of decline above $r > 80$, while those of P_{MI} decline by 10^{-1} by $r = 20$ and at $r > 40$ slows its decline.

5.3 Results

5.3.1 Summary Statistics

Summary properties of the stationary distributions π_{ZI}^*, π_{MI}^* for different grid sizes are provided in Tables 4 and 5.

The reported expectations $E(x), E(y)$ and standard deviations $SD(x), SD(y)$ are calculated using the relevant stationary distribution and the coordinate mapping described at the end of Section 2. For ZI:

$$E_{ZI}(x) = \sum_{f \in \mathbf{F}} x(f) \pi_{ZI}^*(f)$$

$$SD_{ZI}(x) = \sqrt{E_{ZI}(x^2) - (E_{ZI}(x))^2}$$

$p_{within\ triangle}$ gives the probability, summed from the relevant stationary distribution, of a result within the triangular region defined by the ideal points. This is also the Pareto-optimal set for this environment, the set of outcomes where no voter can be made better off without making a voter worse off.

$p_{grid\ boundary}$ gives the probability, summed from the relevant stationary distribution, of a result on the square outer boundary of the grid.

p_{min} gives the minimum value of the stationary distribution. This always occurs at one of the corners of the grid.

p_{max} gives the maximum of the stationary distribution. This always occurs at the grid point $(x = 0, y = -1)$.

Table 4 Spatial Voting: Long-term Means and SDs (values below 10^{-16} truncated)

grid	agenda	$E(x) \pm SD(x)$	$E(y) \pm SD(y)$
80	ZI	0 ± 7.425	-0.3431 ± 7.424
80	MI	0 ± 10.82	-0.342 ± 10.81
60	ZI	0 ± 7.425	-0.3431 ± 7.424
60	MI	0 ± 10.82	-0.3419 ± 10.81
40	ZI	0 ± 7.421	-0.3428 ± 7.424
40	MI	0 ± 10.78	-0.3373 ± 10.77
20	ZI	0 ± 6.985	-0.2937 ± 7.123
20	MI	0 ± 9.032	-0.1452 ± 9.023

Table 5 Spatial Voting: Summary

grid	agenda	$p_{\text{within triangle}}$	$p_{\text{grid boundary}}$	$p_{\text{min corners}}$	$p_{\text{max (0,-1)}}$	H (bits)	$\log_2 \mathbf{F} $
80	ZI	0.675	1.64e-11	1.33e-21	0.00495	9.82	14.66
80	MI	0.396	6.54e-10	6.79e-20	0.00147	10.93	14.66
60	ZI	0.675	2.82e-08	5.55e-16	0.00495	9.82	13.84
60	MI	0.396	6.33e-07	1.60e-14	0.00147	10.93	13.84
40	ZI	0.675	2.55e-05	1.82e-10	0.00495	9.82	12.68
40	MI	0.396	0.000254	2.32e-09	0.00147	10.92	12.68
20	ZI	0.680	0.00864	4.66e-06	0.00494	9.68	10.72
20	MI	0.458	0.0242	1.69e-05	0.00171	10.32	10.72

Shannon's entropy (Shannon, 1948, p. 390) may be the least familiar:

$$\text{Shannon's Entropy } H = - \sum_{f \in \mathbf{F}} \pi_f \log_2 \pi_f$$

H measures the amount of information (bits) a draw from the stationary distribution would require when compressed under an optimal encoding scheme. The “bits” of a $\log_2$ based entropy are the same unit used for computer memory. The minimum possible entropy $H = 0$ occurs when a single point has probability one. $H = 1$ occurs when the distribution consists of two points, each with probability $1/2$. The maximum possible entropy involves a uniform distribution over $\mathbf{F}$, where $H = \log_2 ||\mathbf{F}||$. But the reported entropy of the ZI and MI stationary distributions never reaches that maximum.

Tables 4 and 5 show that grid limits of 40 or higher have very similar summary statistics when considering a single agenda, either ZI or MI. The long-term mean grid coordinate $(E(x), E(y))$ is very close to the ideal point triangle's centroid $(0, -1/3)$. The entropy H appears to stop increasing above size 40 (due to rounding). Clearly the stationary distribution fills the entire space available but has lower and lower outer values as the grid is expanded. The grid size of 20 would place the square grid boundary very close to the edges of the triangle of voter ideal points, and so behaves a bit differently.

Contrasting ZI and MI, the distribution resulting from the ZI-agenda is more compact than that resulting from the MI-agenda, as seen from the standard deviations SD reported in Table 4 and the higher probability-within-triangle and lower entropy (H) columns in Table 5.

5.3.2 Stationary Distributions

Fig. 4 log₁₀ contour and radial plots of stationary distributions π_{ZI}^* and π_{MI}^* . The ideal points $(-15, -9)$, $(0, 17)$, $(15, -9)$ are marked for reference. The contours show a 3-fold flattening consistent but opposite to the triangular configuration of ideal points. Examination indicates lower probabilities along directions through the bots' individual ideal points and higher probability along directions through two-bot midpoints. The radial plots show that MI has a higher probability of greater radial distances than ZI.

The stationary distributions arising from the ZI and MI agendas are shown in Figure 4. Notice that the radial plots of both $\log_{10} \pi_{\text{ZI}}^*$ and $\log_{10} \pi_{\text{MI}}^*$ are primarily linear with radial distance – not quadratic; thus they are not multivariate normals. The ZI agenda concentrates more probability at smaller r . There is extra sharpness in the shapes of the MI-agenda 2-D contours. The thickness of the radial plot lines stem from a range of values found at each radius.

Fig. 5 $\log_{10}$ angular plot of stationary distributions π_{ZI}^* and π_{MI}^* vs θ for $20 \leq r < 21$.

The points are from the stationary distributions on the grid, and the solid curve is an approximation $\text{const} - \cos(3\theta)$. Minima correspond to rays through the ideal points at $\theta = \{0, \pm 2\pi/3\}$. Maxima correspond to rays through 2-voter midpoints.

Further investigation of a narrow ring $20 \leq r < 21$ reveals (as shown in Figure 5) that the stationary distribution probabilities have lows in the directions of the ideal points ($\theta = \{0, \pm 2\pi/3\}$) and highs at the directions of

2-voter coalition midpoints. This suggests a dependence on $-\cos(3\theta)$.¹⁶ The range of 0.4-0.5 in $\log_{10}\pi^*$ is not small, but represents a multiplicative increase in probability of 2.6-3.2. Outliers involve a combination of effects, including the variety of r values found in the ring. From the radial plot in Fig 4, the effect of changing r by 1 ($\Delta\pi_{ZI}^*/\Delta r \approx -0.15$) is about 1/3 of the min-max variation in Fig 5.

Continuing with this idea, we model $\log_{10}\pi^*$ to include a scaled trigonometric portion. For each agenda (ZI,MI) and integer radius r we construct a rough model as follows.

Let the approximation for $\log_{10}\pi^*(f)$, $f \in \mathbf{F}$ be:

$$\text{pred}(f; r) = C(r) - A(r) \cos(3\theta(f))$$

Using the coordinate mapping $r(f)$, the required subset of grid data is:

$$\text{ring}(r) = \{f \in \mathbf{F}; r \leq r(f) < (r + 1)\}$$

From the ring data we calculate:

$$A(r) = \frac{\max_{f \in \text{ring}(r)} \log_{10}\pi^*(f) - \min_{f \in \text{ring}(r)} \log_{10}\pi^*(f)}{2}$$

$$C(r) = \frac{\max_{f \in \text{ring}(r)} \log_{10}\pi^*(f) + \min_{f \in \text{ring}(r)} \log_{10}\pi^*(f)}{2}$$

¹⁶ Various exploratory techniques were attempted but are not reported here, including fitting $\log_{10}\pi^*$ (as the dependent variable) to various trig functions and powers of r . This seemed overly complex, risks overfitting, and likely to become mired in statistical details. Identifying primary frequencies with a fast Fourier transform (FFT) should be possible on a polar grid with equally-spaced θ values. The square grids we use are equally-spaced on x and y only.

and

$$RMSE(r) = \sqrt{\frac{\sum_{f \in ring(r)} [\log_{10} \pi^*(f) - pred(f; r)]^2}{||ring(r)||}}$$

Fig. 6 Radial plots of the Amplitude (left) and Root mean squared error (right) of the cosine waves found in the stationary distributions π_{ZI}^* and π_{MI}^* in unit rings with inner radius r . Generally, the MI Amplitude and RMSE lie slightly below those for ZI. The loss of fit for $r > 80$ occurs because of the grid boundaries at ± 80 .

Figure 6 shows the resulting plots of $A(r)$ and $RMSE(r)$ vs. r . We omitted $C(r)$ because it replicates existing radial plots. Generally, the Amplitude A appears to be an increasing function of r , with a change in slope near 20, perhaps due to the ideal points at $(r \approx 17.33)$. The RMSE increases slowly, with large increases after $r > 80$ due to the corners of the ± 80 grid.

One last mystery to unravel involves the crisper shape of the contour plots for π_{MI}^* as compared to π_{ZI}^* in the upper and middle panes of Figure 4. The 3-fold symmetry of the environment matches the symmetry of the contours of the stationary distributions, and the symmetry of $\cos(3\theta)$. So why is MI

crisper than ZI? The Amplitudes for ZI and MI seem nearly identical, with the ZI Amplitude being slightly higher. The reason that the ZI stationary distribution does not have the crisper contour plot may lie in the fact that RMSE is also higher for ZI, especially within $r < 20$. To capture a kind of signal-to-noise ratio, define $Prominence = A/RMSE$ in Figure 7.

Fig. 7 Radial plot of the prominence ($A/RMSE$) of the cosine waves found in the stationary distributions π_{ZI}^* and π_{MI}^* in unit rings with inner radius r . Prominence is higher for the MI agenda. Both appear to peak around the radius for the ideal points ($r \approx 17$).

Figure 7 shows a much higher prominence ratio for MI than for ZI in the range $5 < r < 25$, similar to the range at which the (Figure 4) contour plots for MI appear sharper than those than for ZI. Apparently the functional forms

for the angular dependence are very similar, but MI is simultaneously “louder” and less “noisy” in the central region towards which the viewers eye is drawn.

5.4 Consequences of ZI and MI Agendas

We can see the benefits of the ZI agenda played out as a more compact stationary distribution concentrated closer to the center of the grid.

Specifically, (a) standard deviations in Table 4 are always lower for the ZI agenda (SD 7.4) than the MI agenda (SD 10.8); and from Table 5 we see (b) the entropy of the stationary distribution is lower for the ZI agenda (9.8) than for the MI agenda (10.9); (c) the probability of an outcome in the Pareto-set, which is equivalent to the set of alternatives within the triangle of ideal points, is higher for the ZI agenda (0.67) than for the MI agenda (0.39); and (d) the probability of being on the grid boundary is always lower for the ZI agenda.

Therefore, in this environment, the ZI agenda outperforms the MI agenda if we judge the agendas by their resulting stationary distributions. The drawback of the ZI agenda is longer convergence times (higher powers in Table 1), compared with the MI agenda.

6 Example: Voting to divide a budget

In this section, we will demonstrate the application of our ZI and MI random agenda models to budget allocation.

As in the previous sections, the ZI agenda will show slower convergence. Long-term behavioral differences will arise as differences in the stationary dis-

tributions. Not only will MI agenda be seen to exhibit higher variance than the ZI agenda, but the maximum-likelihood feasible alternative will be different.

Like the previous section, the voting environment will involve a series of 2-D (x, y) points chosen from a finite integer domain. However, the square grid of the previous section will be replaced by a triangular grid naturally occurring in budget allocation.

6.1 Environment E3

Consider three bots who must divide a budget of 100 units of some good. The good might be crypto-currency, electrical power, time, computer memory, human attention, or something else. Each bot prefers their allocation to be more, not less, and is indifferent about the allocations of other bots.

A well-known 2-D model for this environment is to let the feasible set $\mathbf{F}$ be pairs of integers (x, y) satisfying:

$$0 \leq x \leq 100$$

$$0 \leq y \leq 100 \text{ and}$$

$$0 \leq (x + y) \leq 100$$

This set forms a triangle (including the interior).

At a feasible alternative $f = (x, y)$, voter 1 receives x units and so has utility function $u(x, y) = x$; voter 2 receives y units and has utility function $u(x, y) = y$; and voter 3 receives the remainder, $100 - x - y$ units, and has utility function $u(x, y) = 100 - x - y$. These utility functions are shown in Figure 8.

Fig. 8 Utility functions for Voters 1,2, and 3 dividing a total budget of 100

Cycling in this environment occurs because at every alternative, winning challengers can be constructed by reducing some voter's allocation for redistribution to the remaining two voters.

There are three transient alternatives – at $(0,0)$ and $(100,0)$ and $(0,100)$ – where all interior challengers win in a pairwise vote.

6.2 Transition Matrices

As explained in the previous example, the transition matrices cannot be visualized directly. However, as before, we provide contour plots of the diagonal terms in Figure 9. While the contours appear similar, the diagonal values for P_{MI} are about 10^{-3} lower in magnitude than those of P_{ZI} .

Fig. 9 Diagonal elements of Markov Chain transition matrices for budget voting. Note the P_{ZI} diagonals approach 1.0 whereas the P_{MI} diagonals are all below 0.001

6.3 Results

6.3.1 Summary Statistics

Similar to the previous spatial voting example, moving from the ZI agenda to the MI agenda increases the entropy and variance of outcomes. The entropy values of 12.16 and 12.25 are also very close to the maximum possible entropy of $\log_2 5151 = 12.33$ that would exist for a uniform distribution on all alternatives. The increase in variance is seen in the SD entries in Table 7. These change from 19.6 to 21.1 for both x and y . These differences in ZI and MI also include a change in the alternative having maximum probability.

Table 6 Budget Voting: Summary

agenda	$p_{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}}$	$p_{0, \frac{1}{2}, \frac{1}{2}}$	p_{boundary}	p_{min}	H	$\log_2 \mathbf{F} $
ZI	9.86e-04	6.91e-04	0.0346	1.35e-94	12.16	12.33
MI	7.28e-04	7.72e-04	0.0443	1.14e-122	12.25	12.33

Table 7 Budget Voting: Long-term Means and SDs

agenda	$E(x) \pm SD(x)$	$E(y) \pm SD(y)$
ZI	33.3 ± 19.6	33.3 ± 19.6
MI	33.3 ± 21.1	33.3 ± 21.1

In what is the most striking difference in the two stationary distributions, we see the replacement of a maximum (under π_{ZI}^*) near the “equal” point $(33, 33)$ with a saddle and a new maxima at the two-way splits $(50, 50)$, $(0, 50)$, $(50, 0)$ (under π_{MI}^*). While this can be determined from the first three columns of Table 5, it is also visible in the upcoming contour plots.

6.3.2 Stationary Distributions

Fig. 10 Markov Chain stationary distributions for budget voting. Note that π_{ZI}^* has maxima at $(33, 33)$, $(33, 34)$, $(34, 33)$ whereas π_{MI}^* has a saddle point near $(33, 33)$ and constrained maxima at $(0, 50)$, $(50, 50)$, and $(50, 0)$

The stationary distributions are shown in Fig 10 above. Various summary statistics are given in Table 6 and 7.

Both stationary distributions feature near-zeroes near at the three corners of the triangle where a single voter receives all 100 units of the good. These in fact should be zero, but due the nature of the calculations show up at probabilities around 10^{-100} in the prob_min entries in Table 6.

Permuting voters 1, 2, and 3, or exchanging x with y or $z = (100 - x - y)$ are cases of a special symmetry inherent in the budget allocation scenario that will leave the stationary distribution unchanged. An example of this can be

seen at most levels of π^* , where the same level of probability is seen at the six alternatives $\{(x, y), (y, x), (x, z), (z, y), (z, x), (y, z)\}$.

Fig. 11 Payoff distributions for budget voting. The frequency of mid-range payoffs is higher under the ZI-agenda.

The long-term payoff distributions under the ZI and MI agendas are given in Figure 11. Although the result is stated for agent one, it applies to the others (though not simultaneously as there is a finite sum and interdependency). The

probability of a zero payoff with the ZI-agenda is approximately 1.1%, but it increases to almost 1.5% with the MI agenda. This is consistent with Table 6 where $p_{boundary}$ can be divided by 3 to obtain the zero-payoff probability here.

6.4 Consequences of ZI and MI Agendas

In essence, the MI agenda – compared against the ZI agenda – provides each voter a higher chance of receiving nothing or a very low payoff and reduced chances of mid-range payoffs in exchange for a “lottery” that increases variance and increases the probability of higher values. In each case the mean is near the common value of 33.3.

The most likely outcome under the MI agenda is a two way 50-50 split, whereas under the ZI agenda the most likely outcome is a permutation of 33-33-34. The most likely outcomes are not very likely ($p \leq 0.001$) because the stationary distribution is well-spread over the 5,150 other possibilities.

Convergence to the stationary distribution under the transition matrix for the ZI agenda is about 1/2 the speed of the MI agenda (Table 2, power column). This is a narrower ratio than achieved with the spatial voting examples of the previous section.

7 Conclusions

In this section, we provide conclusions in the context of our opening questions and the larger literature.

7.1 Answers to the motivating questions

In the introduction, we consider three questions about bots acting under majority rule:

Q1: *Should we trust reasonable outcomes to emerge from groups of simple bots operating under majority rule – or should we expect the bots to run amok?*

The answer to this question depends on the criteria applied and is unlikely to satisfy everyone.

If the criteria relates to total chaos, the kind that Ordeshook (2007, p.165) criticizes as carrying “the implication that even probabilistic predictions were problematical,” then our results show none and instead are closer in spirit to the central tendency of Tullock (1967) but with cycling present and only three voters. Probabilistic predictions are possible in the long-term via the stationary distributions. These stationary distributions exist and their entropy stabilizes as grid size is increased in E2. They have zones of concentration. In all environments we study, the stationary distributions contain symmetries associated with features of the environment.

We claimed in the introduction that the bots run amok in the short run but do not exhibit extreme behavior in the long run. When the short-run consists of the behavior created by the transition matrix, and the long-run consists of the various long-run averages implied by the stationary distributions, the fact that the long-run is less variable can help sustain an opinion that the robots do not run amok in the long-run. But this view is subjective.

From a perfectionist engineering perspective, it might seem that the bots almost always run amok and there is a potential solution in an alternative mechanism. After all, a perfectionist can argue that in E2 the solution should be at the center, perhaps through averaging the ideal points, but the ZI agenda puts it adjacent to the exact center with $p < 1\%$ (p_{max} , Table 5). In E3 the same perfectionist can argue that the solution should be the almost equal split, which occurs with $p < 0.001$. Benevolent dictatorship can impose these outcomes. Besides being a bit beyond the initial scope raised in the introduction, this approach potentially requires human analysis of each scenario. Switching from sequential voting to collection of other information is certainly possible in an environment of faithful truth-telling bots whose programming is secured by modern cryptography. But ad-hoc dictatorial or heuristic approaches will have their own failure modes and are still subject to the voting and mechanism design impossibility theorems. Too much human meddling also violates the introductory expectation that bots “make decisions in place of humans ... on decentralized information that may need to be aggregated without being passed on to humans.”

Examining the earlier social choice literature resets our general expectations lower and encourages quite a bit more tolerance. A variety of outcomes should emerge, not simply those that are most “reasonable” or “fair” by some arbitrary standard. If such a variety or particularly-shaped distribution is also desirable from a systems engineering perspective, then the choice of the ZI or MI agenda, or some similar approach, may be warranted. Giving the bots

a substantial probability of searching outside their Pareto-set / ideal point triangle might slow down some activities but also avoid bad outcomes where the bots are collectively wrong. Actual engineering is not the province of perfectionists but that of savvy traders balancing these kinds of tradeoffs.

Dynamics exist that include running amok. Bots can eventually reach outcomes that would be defeated in a comparison with an earlier state. Maintaining an extreme outcome is improbable under either the MI or ZI agenda, as shown by the diagonal elements of the transition matrix. In this sense, the long-term behavior does not run amok. Earlier or more central states are always potential competitors to later outcomes. In this sense, the potential for cycling back to an earlier state can be a benefit. However, to support the conclusion that this dynamic is “reasonable” overall requires a movement of the goal-posts away from the settings of a perfectionist.

Q2: What processes should be in place – should the bots consider all feasible choices or restrict consideration to speed up decision-making?

There is a tradeoff between speed of convergence (preferring MI) and the tendency to retain well-accepted status quos (ZI).

The ZI agenda considers all feasible choices, but as a result has a tendency to retain the status quo when there are few successful challengers and many unsuccessful challengers. In contrast, the MI agenda speeds up change by giving the status quo much less weight and is almost in constant motion in E2 and E3. The series of outcomes resulting from the ZI agenda will appear to stall

or slow down at well-accepted outcomes. Both agendas will quickly replace highly dominated outcomes.

Using the MI agenda made these features worse:

- **extreme outcomes:** occur more often for the MI agenda. See Tables 5 6, column: `prob_on_boundary`. In the budget allocation scenario, these are outcomes where a bot receives zero of something and, for physical bots, that might mean failure or being turned off.
- **variance of outcomes:** higher for the MI agenda. See Tables 4 7, columns $SD(x)$, $SD(y)$
- **likelihood of Pareto-set outcome:** lower for the MI agenda (39% MI vs 67% ZI). See Table 5, column: `prob_within_triangle`.

The result that a speedier, impartial, “fair” and voter-conscious smart-government – aware of the alternatives voters would approve – could perform objectively worse than a slower, zero intelligence random government may surprise some readers. Notice that the slow-to-change aspect does not involve an unwillingness to consider radical alternatives, but rather is because radical alternatives crowd out passable alternatives for opportunities to be the next challenger. There is a constant stream of radical challengers as the ZI agenda becomes more stable, and they are all rejected in voting. Perhaps the surprise here is that the low-probability “irrelevant outcomes” at the limits of the grid can turn out to have stabilizing effects in the center.

The ZI agenda’s Markov chain does take longer – sometimes many times longer – to converge (Tables 1 and 2). This is a tradeoff for more stable out-

comes. The ZI agenda's consideration of often-called "irrelevant" alternatives that cannot possibly win is part of what creates more stability at better outcomes as well as longer convergence time.

Q3: Can we characterize the emerging bot behavior even in the absence of human notions of strategy, learning or equilibria?

This research demonstrates one approach: characterization of emerging system behavior as a Markov chain.

The stationary distribution is not a strategy, and it does not reflect any learning or change in the bot behaviors. It is useful for calculating long-term averages of functions that can be expressed in terms of the x and y variables on the grid. However, it is not an equilibrium. It is the present expectation about the far future state. The Markov chain describes short term behavior that does not converge. This short-term behavior is always described by the current state and the state-dependent row of the transition matrix.

It may be possible to reuse the grid and matrix-based approach to study emergence in other Markov chain scenarios, such as markets or games in which stochastic ZI or MI agents participate. But this currently excludes ZI-populated double auction markets – only because the reduction of the state space to a reasonably sized grid seems a fruitless challenge. Traditional iteration is still a more practical approach for markets and other complex scenarios.

7.2 Implications for voting-based technologies and political science

The primary engineering implication for voting-based technologies is that tools and prototype-testing concepts are possible and necessary. While prototyping often includes application-specific training and testing regimens (especially in machine learning and data science), our approach is different. We did not begin with large external datasets. Instead we identified test environments where majority rule may perform poorly. The techniques and concepts are transferable to other appropriate endeavors. Proper engineering choices involving the behavior of bots acting under majority rule (or other mechanisms) generally require testing the impact of seemingly minor procedural details. Issues surrounding agendas, who sets them or how they are set are significant and affect outcomes. Our software provides a practical example for jump-starting more exploration. While we have shown how majority rule can function in some environments thought to be problematic, we also noted that it was fairly easy to imagine other rules as possibly superior. Testing is crucial, and system designers should take care to evaluate a variety of mechanisms for their applications.

The implication for political science is that there may be new customers with smaller, more tractable problems: the bots and the firms that create them. We have placed cooperative, majority-voting bot swarms firmly within the framework for spatial voting, which has typically concentrated on abstract, mathematical voters or agents – usually hyper-rational ones (although we take the opposite approach for simplicity) – and not necessarily on groups of actual humans with all of their complexity and fallibility (empirical and

experimental political science collects data examining actual human behavior). We've mentioned differences (speed, discrete grids vs. continuous spaces, faithful simplicity vs. strategic rationality) that in practice did not prevent modelling, testing, or analysis.

Possible future work and open questions include exploring what other scenarios in political science can be appropriately modelled with simple ZI /MI agents, the replication and durability of our results on polar and other grids, the practical roles for seemingly irrelevant alternatives, the consideration of other voting mechanisms such as Borda count or single-transferable-vote, and better ways to characterize or visualize outcomes and the transition matrices. (Each row of the transition matrix defines a probability distribution. Does the entropy of these row-based distributions play a role?).

Symmetry of our environments (E1, E2, E3) is impressed into the transition matrices and results. The importance of symmetry across numerous sciences and engineering is well known and difficult to ignore. In the early 1900s, physicists began connecting symmetry to the fundamental conservation laws and the properties of elementary fields and particles (Gross, 1996). What are the opportunities for symmetry to expand understanding of the social sciences? We found symmetry and order in our results, not total chaos. But various forms of mathematical chaos exist that are not total chaos and are compatible with symmetry, resulting in fractals and other consequences. Chaos theory has ergodic theorems yielding stationary distributions. Does that kind of chaos

and related symmetry phenomena exist in the spatial voting behavior of simple bots or more complex artificial intelligence?

Opportunities for scaling up exist. In particular, there is virtually unlimited scale available for adding more voters on the same spatial voting grid. It does not change the size of the transition matrix.¹⁷ With many voters on the same grid, it simply takes a little longer to calculate the transition matrix the first time. Subsequent steps (powers, matrix algebra) are mostly unaffected. There are no computational issues with multiple voters having the same ideal point, so one could numerically recreate the Tullock (1967) suggestion of a grid containing millions of voters. The consideration of a number of alternative voting procedures may also be possible.

Limitations involve grid size and the number of dimensions which in turn affect the size of the transition matrix. In the 2-D case, a grid limit of $|x|, |y| \leq N$ implies transition matrices of dimension $O(N^2 \times N^2)$ requiring $O(N^4)$ memory. Going from a 16 GB Co-Lab GPU to an expensive 80 GB GPU, a five fold increase in memory, suggests only a 1.5 fold increase in 2-D grid limits. Research on larger grids may need to wait – either for hardware or effective software for utilizing multiple GPUs. Future innovations should slowly increase capabilities and drive down the cost of this methodology.

We see this paper as a contribution to both the understanding of emergent ZI/MI behavior, particularly within majority rule, and to *political engineer-*

¹⁷ However, the budget voting environment E3 is an exception. Adding more changes the dimensional model and vastly enlarges the grid and transition matrix.

ing, a field suggested by Ordeshook (1993) that is gradually emerging due to viable problem-solving opportunities.^{18,19} There is both opportunity and obvious danger with the adoption and growth of swarms of autonomous bots, aerial drones, and similar technologies – whether based in very simple ZI/MI behavior or sophisticated artificial intelligence. We believe the social sciences has a role to play in the analysis, design and engineering of such artificial societies, and will continue to adapt to meet these future challenges.

¹⁸ Ordeshook was writing partially in jest, criticizing the mathematically complex and often unrealistic assumptions required to obtain provable theoretical results, the lengthy time required for any kind of analysis, and the growing divide between the state of theory and common demands for answers. He claims that the field needs to learn to distinguish science from engineering. He envisions science as patiently searching for general laws and models and proving theorems, while engineering gets the large impractical tasks (unified models of entire national and international political apparatus, crisis management, democratic reform) and must attract the necessary interdisciplinary teams to solve the practical hurdles. There is further discussion, intriguing but beyond our scope.

¹⁹ According to Ordeshook, political engineering would combine computational and institutional knowledge with practical demand. With the advantage of the intervening decades, we would nominate a few small-scale examples: Levine and Plott (1977), perhaps the earliest, for engineering the outcome of a flying club's vote (involving the purchase of a fleet of aircraft) through agenda manipulation and design; Conitzer and Sandholm (2003) for using computational complexity to engineer difficulties for strategic insincere voting; and Avin et al. (2019) for estimating elite influence on majority opinion in large online communities.

References

- Aldrich JH, Alt JE, Lupia A (eds) (2007) Positive changes in political science: The legacy of Richard D. McKelvey's most influential writings. University of Michigan Press
- Arrow KJ (1950) A difficulty in the concept of social welfare. *Journal of Political Economy* 58(4):328–346
- Aseere AM, Gerding EH, Millard DE (2010) A voting-based agent system for course selection in e-learning. In: 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, IEEE, vol 2, pp 303–310
- Avin C, Lotker Z, Mizrahi A, Peleg D (2019) Majority vote and monopolies in social networks. In: Proceedings of the 20th International Conference on Distributed Computing and Networking, pp 342–351
- Barshan B (2001) Ultrasonic surface profile determination by spatial voting. In: IMTC 2001. Proceedings of the 18th IEEE Instrumentation and Measurement Technology Conference. Rediscovering Measurement in the Age of Informatics (Cat. No.01CH 37188), vol 1, pp 583–588 vol.1, DOI 10.1109/IMTC.2001.928885
- Beeson MJ (1992) Triangles with vertices on lattice points. *The American mathematical monthly* 99(3):243–252, DOI <https://doi.org/10.2307/2325060>
- Black D (1948) On the rationale of group decision-making. *Journal of Political Economy* 56(1):23–34

- Brewer PJ, Huang M, Nelson B, Plott CR (2002) On the behavioral foundations of the law of supply and demand: Human convergence and robot randomness. *Experimental economics* 5(3):179–208
- Condorcet JAMNC, et al. (1785) Essai sur l'application de l'analyse à la probabilité des décisions, rendues à la pluralité des voix/— c par m. le marquis de condorcet
- Conitzer V, Sandholm T (2003) Universal voting protocol tweaks to make manipulation hard. DOI 10.48550/ARXIV.CS/0307018, URL <https://arxiv.org/abs/cs/0307018>
- Cox GW, Shepsle KA (2007) Majority Cycling and Agenda Manipulation: Richard McKelvey's Contributions and Legacy. In: Aldrich JH, Alt JE, Lupia A (eds) Positive changes in political science: The legacy of Richard D. McKelvey's most influential writings, University of Michigan Press, chap 3
- Dougherty KL, Edward J (2005) A nonequilibrium analysis of unanimity rule, majority rule, and pareto. *Economic Inquiry* 43(4):855–864
- Dougherty KL, Edward J (2012) Voting for pareto optimality: a multidimensional analysis. *Public Choice* 151(3):655–678
- Duggan J (2005) A survey of equilibrium analysis in spatial models of elections
- Duggan J, Schwartz T (2000) Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare* 17(1):85–93, URL <http://www.jstor.org/stable/41106341>
- EU-Parliament (2016) Regulation 2016/679, of the European Parliament and of the Council of 27 April 2016 on the Protection of

- Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), 2016 O.J. (L 119) 1. URL <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:02016R0679-20160504&from=EN#tocId31>
- Farsa DZ, Bidgoli AA, Rokhsat-Yazdi E, Rahnamayan S (2021) Population-based coordinate descent algorithm with majority voting. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion, Association for Computing Machinery, New York, NY, USA, GECCO '21, p 1283–1289, DOI 10.1145/3449726.3463186, URL <https://doi.org/10.1145/3449726.3463186>
- Ferejohn JA, McKelvey RD, Packel EW (1984) Limiting distributions for continuous state markov voting models. *Social Choice and Welfare* 1(1):45–67
- Friedrich T, Kötzing T, Krejca MS, Nallaperuma S, Neumann F, Schirneck M (2016) Fast building block assembly by majority vote crossover. In: Proceedings of the Genetic and Evolutionary Computation Conference 2016, pp 661–668
- Gibbard A (1973) Manipulation of voting schemes: A general result. *Econometrica* 41(4):587–601, URL <http://www.jstor.org/stable/1914083>
- Gibbard A (1977) Manipulation of schemes that mix voting with chance. *Econometrica* 45(3):665–681, URL <http://www.jstor.org/stable/1911681>

- Gode DK, Sunder S (1993a) Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* 101(1):119–137
- Gode DK, Sunder S (1993b) Lower bounds for efficiency of surplus extraction in double auctions. In: *The Double Auction Market Institutions, Theories, and Evidence, Proceedings of the Workshop on Double Auction Markets held June 1991 in Santa Fe, New Mexico*, pp 199–219
- Gode DK, Sunder S (1997) What makes markets allocationally efficient? *The Quarterly Journal of Economics* 112(2):603–630
- Grinstead C, Snell J, Doyle P (2006) Introduction to probability. URL <https://math.dartmouth.edu/~prob/prob/prob.pdf>
- Gross DJ (1996) The role of symmetry in fundamental physics. *Proceedings of the National Academy of Sciences* 93(25):14256–14259
- Häggström O, et al. (2002) *Finite Markov chains and algorithmic applications*, vol 52. Cambridge University Press
- Hayek F (1945) The use of knowledge in society. *The American Economic Review* 35(4):519–530
- James G, Witten D, Hastie T, Tibshirani R (2013) *An introduction to statistical learning*. Springer
- Kalandrakis A (2004) A three-player dynamic majoritarian bargaining game. *Journal of Economic Theory* 116(2):294–314, DOI [doi.org/10.1016/S0022-0531\(03\)00259-X](https://doi.org/10.1016/S0022-0531(03)00259-X)

Kent A, Williams J (1988) Computers in the space shuttle avionics system.

URL <https://history.nasa.gov/computers/Ch4-4.html>

Kubernetes blog contributors (2016) Simple leader election with ku-

bernetes and docker. URL <https://kubernetes.io/blog/2016/01/simple-leader-election-with-kubernetes/>

Levin DA, Peres Y, Wilmer E (2009) Markov chains and mixing times. American Mathematical Soc, Providence

Levine ME, Plott CR (1977) Agenda influence and its implications. Virginia Law Review pp 561–604

McKelvey RD (1976) Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic theory* 12(3):472–482

McKelvey RD, Ordeshook PC, et al. (1990) A decade of experimental research on spatial models of elections and committees. *Advances in the spatial theory of voting* pp 99–144, DOI 10.1017/CBO9780511896606.007

Numpy Developers (2006) Numpy: The fundamental package for scientific computing with python (version 1.19.5). URL <https://pypi.org/project/numpy/1.19.5/>, see also <https://numpy.org/>

Oikonomopoulos A, Patras I, Pantic M (2010) Discriminative space-time voting for joint recognition and localization of actions. In: *Proceedings of the 2nd International Workshop on Social Signal Processing*, Association for Computing Machinery, New York, NY, USA, SSPW '10, p 11–16, DOI 10.1145/1878116.1878122, URL <https://doi.org/10.1145/>

1878116.1878122

- O'Neil C (2016) Weapons of math destruction: How big data increases inequality and threatens democracy. Broadway books
- Ordeshook P (1993) The development of contemporary political theory. In: Barnett WA, Hinich MJ, Schofield N (eds) Political economy: Institutions, competition, and representation. Proceedings of the seventh international symposium in economic theory and econometrics, chap 4
- Ordeshook P (2007) The competitive solution revisited. In: Aldrich JH, Alt JE, Lupia A (eds) Positive changes in political science: The legacy of Richard D. McKelvey's most influential writings, University of Michigan Press, chap 9
- Parsons R, Truran J (1970) 206. equilateral triangles on geoboards. The Mathematical Gazette 54(387):53–54, DOI <https://doi.org/10.2307/3613163>
- Paul TK, Iba H (2009) Prediction of cancer class with majority voting genetic programming classifier using gene expression data. IEEE/ACM Transactions on Computational Biology and Bioinformatics 6(2):353–367, DOI 10.1109/TCBB.2007.70245
- Penn EM (2009) A model of farsighted voting. American Journal of Political Science 53(1):36–54
- Plott C (1967) A notion of equilibrium and its possibility under majority rule. The American Economic Review 57(4):787–806
- Plott CR (1994) Symposium: Market architectures, institutional landscapes and testbed experiments: Introduction. Economic Theory 4(1):3–10, URL <http://www.jstor.org/stable/25054746>

Preferred Networks Inc (2015) Cupy: Numpy and scipy for gpu (version 9.1.0).

URL <https://pypi.org/project/cupy/9.1.0/>, see also <https://cupy.dev/>

Raschka S (2015) Python machine learning. Packt publishing ltd

Redacted Author Name (2021) Python module “gridvoting” (version 0.9.3).

URL <https://pypi.org/project/gridvoting/0.9.3/>

Satterthwaite MA (1975) Strategy-proofness and arrow’s conditions: Existence

and correspondence theorems for voting procedures and social welfare func-

tions. *Journal of Economic Theory* 10(2):187–217, DOI [https://doi.org/](https://doi.org/10.1016/0022-0531(75)90050-2)

10.1016/0022-0531(75)90050-2, URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/0022053175900502)

[science/article/pii/0022053175900502](https://www.sciencedirect.com/science/article/pii/0022053175900502)

Schofield N (1978) Instability of Simple Dynamic Games. *The Review*

of Economic Studies 45(3):575–594, DOI 10.2307/2297259, URL

<https://doi.org/10.2307/2297259>, [https://academic.oup.com/](https://academic.oup.com/restud/article-pdf/45/3/575/4475802/45-3-575.pdf)

[restud/article-pdf/45/3/575/4475802/45-3-575.pdf](https://academic.oup.com/restud/article-pdf/45/3/575/4475802/45-3-575.pdf)

Schwartz T (1982) No minimally reasonable collective-choice process can be

strategy-proof. *Mathematical Social Sciences* 3(1):57–72, DOI [https://doi.](https://doi.org/10.1016/0165-4896(82)90006-3)

[org/10.1016/0165-4896\(82\)90006-3](https://doi.org/10.1016/0165-4896(82)90006-3), URL [https://www.sciencedirect.](https://www.sciencedirect.com/science/article/pii/0165489682900063)

[com/science/article/pii/0165489682900063](https://www.sciencedirect.com/science/article/pii/0165489682900063)

Shannon CE (1948) A mathematical theory of communication. *The Bell sys-*

tem technical journal 27(3):379–423

Smith VL (1962) An experimental study of competitive market behavior. *Jour-*

nal of political economy 70(2):111–137

- Teixeira M, d'Orey PM, Kokkinogenis Z (2018) Autonomous vehicles coordination through voting-based decision-making. In: International Conference on Agreement Technologies, Springer, pp 199–207
- Teixeira M, d'Orey PM, Kokkinogenis Z (2020) Simulating collective decision-making for autonomous vehicles coordination enabled by vehicular networks: A computational social choice perspective. *Simulation Modelling Practice and Theory* 98:101983
- Tseitlin A (2013) The antifragile organization. *Commun ACM* 56(8):40–44, DOI 10.1145/2492007.2492022, URL <https://doi.org/10.1145/2492007.2492022>
- Tullock G (1967) The general irrelevance of the general impossibility theorem. *The Quarterly Journal of Economics* 81(2):256–270
- Xu Y, Deng Z, Wang M, Xu W, So AMC, Cui S (2020) Voting-based multiagent reinforcement learning for intelligent iot. *IEEE Internet of Things Journal* 8(4):2681–2693
- Yu L, Wang S, Lai KK (2009) An intelligent-agent-based fuzzy group decision making model for financial multicriteria decision support: The case of credit scoring. *European journal of operational research* 195(3):942–959