

# Improving Scalability and Dependability of Stochastic MAS Analyses

Logan Brooks\*  
Computer Science Department  
Carnegie Mellon University  
Pittsburgh, PA, USA  
lcbrooks@cs.cmu.edu

Sandip Sen  
Tandy School of Computer Science  
University of Tulsa  
Tulsa, OK, USA  
sandip-sen@utulsa.edu

## ABSTRACT

When studying stochastic multiagent system models, it is currently most common to perform analysis using Monte Carlo simulations. However, this approach can be prohibitively expensive in certain settings. When system behavior is highly variable, a large number of simulations is needed to understand its behavior for even a single parameter configuration. Simulation performance can also scale poorly with certain parameter values, such as the number of agents and time steps. Working over a parameter space increases costs even further. We present an analytical approach to characterizing stochastic multiagent systems with (a) runtime independent of system variability, logarithmic in the number of time steps, and dependent on interaction size rather than population size; and (b) hard, non-probabilistic, error bounds. This method is applied in a sociodynamics setting, illustrating how an analytical approach can produce exact predictions quickly when a simulation approach could perform poorly. We also demonstrate how to characterize behavior across a parameter space and perform approximate inference and optimization tasks when using this analytical approach.

## Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

## Keywords

Stochastic multiagent systems; opinion dynamics; analysis

## 1. INTRODUCTION

Stochastic multiagent systems represent the behavior of populations of interacting agents, incorporating elements of randomness and uncertainty. This class of systems is useful for investigating phenomena in many fields, including social sciences, economics, biology, and physics.

Familiar examples from multiagent system literature include models of opinion dynamics in human societies, which

---

\*This paper represents work done while L. Brooks was a student at the University of Tulsa.

**Appears in:** *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*  
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

define rules for selection and effects of interaction between individuals [15]. In particular, the Bounded Confidence model is frequently referenced [2, 5]. Another commonly studied stochastic MAS is policy evolution of players in a population over time as they repeatedly interact in Prisoner's Dilemma settings [13]. Typically, this analysis relies on mean field theory. However, this category includes systems from other fields as well, such as population dynamics and predator-prey models [4, 11], and electromagnetics [16].

## 1.1 Current simulation approaches

Existing studies tend to focus on empirical analysis, examining the behavior of a system by inspecting aggregate data from a large number of simulations. This approach presents a number of problems. The accuracy and reliability of the predictions is often unknown; measures of uncertainty, such as confidence intervals and probably-approximately-correct bounds, are difficult to develop for complicated systems. Even for simple systems, though, the number of samples needed for a certain level of confidence increases exponentially with the desired degree of accuracy. Characterizing the behavior of a system across a range of parameter values increases the cost even further, making it difficult to see trends in data, tune parameters to match real data, infer details about unobserved quantities, or choose parameter values to promote certain types of system behavior.

## 1.2 Current analytical approaches

Current analytical approaches, on the other hand, often focus on studying a very small subclass of possible problems, and are not usually applied to newly proposed stochastic MAS models. The Kalman filter [1, 14], for example, is well-studied, but focuses on a specific scenario of a Markov process with a linear, deterministic state transition, an additive deterministic control component, and additive Gaussian noise. Generalization is performed manually: for instance, the extended Kalman filter [8] and unscented Kalman filter [10] have been developed separately to handle nonlinear state transitions. However, these approaches do not handle nondeterministic state transitions and other types of process noise. Mean field analysis typically is even more focused on specific tasks, such as prediction of phase transitions [7], behavior of electromagnetic systems, and specific formulations of the Bounded Confidence model [3]. Often, results are presented in the form of limits as a population's size approaches infinity [6, 7, 9]. Other analytical approaches that are more easily generalized, while still focusing on specific problem instances, often neglect much of the variation between possible

population configurations, resulting in suboptimal accuracy and a lack of proven error bounds. An approach that is both rigorous and fairly general would be preferred.

## 2. MOTIVATING EXAMPLE

Consider the following situation: a single agent flips an unbiased coin exactly once. If the result is heads, then the agent receives \$1; otherwise, it receives nothing. Suppose that we are interested in the expected monetary payoff of the agent, as well as the variability in the payoff, as expressed by higher-order moments. Analytically, we can find the moments immediately: the  $j$ th raw moment is

$$\mu'_j = (0.5)(\$1)^j + (0.5)(\$0)^j = 0.5 \text{ dollar}^j. \quad (1)$$

However, finding just the mean payoff using simulations can be inefficient.

The procedure for analyzing the accuracy and reliability of the empirical approach very closely resembles that used for constructing confidence intervals for the coin's bias. Let  $K$  be the number of heads obtained given  $n_{\text{sims}}$  simulated coin flips;  $K$  is drawn from the Bernoulli distribution  $B(n_{\text{sims}}, 0.5)$ . The probability of obtaining  $k$  heads in a randomly sampled  $n_{\text{sims}}$ -simulation batch is

$$P(K = k) = \binom{n_{\text{sims}}}{k} (0.5)^k (1 - 0.5)^{n_{\text{sims}} - k} \quad (2)$$

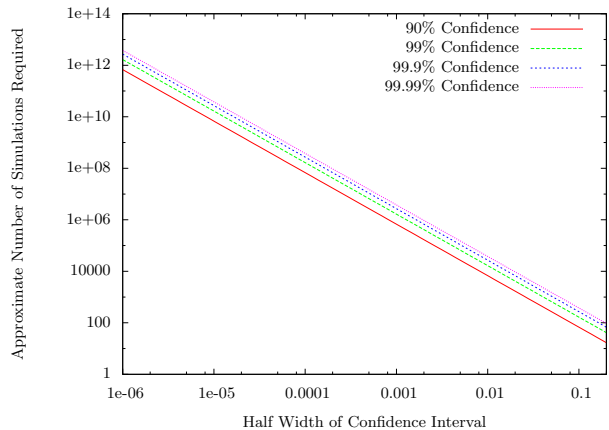
$$= \binom{n_{\text{sims}}}{k} (0.5)^{n_{\text{sims}}}, \quad (3)$$

and that of observing a sample mean payoff within  $e$  dollars of the true mean payoff, for  $0 \leq e \leq 0.5$ , is

$$P(\lfloor (0.5 - e)n_{\text{sims}} \rfloor \leq K \leq \lfloor (0.5 + e)n_{\text{sims}} \rfloor) = (0.5)^{n_{\text{sims}}} \sum_{k=\lfloor (0.5 - e)n_{\text{sims}} \rfloor}^{\lfloor (0.5 + e)n_{\text{sims}} \rfloor} \binom{n_{\text{sims}}}{k}. \quad (4)$$

Applying this formula, we can find, for example, that the probability that a random 1,000-simulation batch will have a sample mean between 0.49 and 0.51 (i.e., yield a relative error of 2% or less) is only about 49.3%. For large  $n$ , we can approximate the sample mean's distribution by  $\frac{1}{2} + \frac{1}{2\sqrt{n_{\text{sims}}}}Z$ , where  $Z \sim \mathcal{N}(0, 1)$ . The probability that a randomly generated sample mean is within  $e = \frac{z}{2\sqrt{n_{\text{sims}}}}$  of the true mean  $p$  is approximated by  $\Phi(z) - \Phi(-z)$ , where  $\Phi(\cdot)$  is the cumulative distribution function of the standard normal distribution. Thus, for a given level of confidence, i.e., a fixed  $z$  value, the required number of simulations to ensure an error bound of  $e$  is  $\frac{z^2}{4e^2}$ . Figure 1 shows the number of simulations required to build various confidence intervals for the mean payoff for the approximate distribution (note the logarithmic scale). To gain an order of magnitude in accuracy, the sample size must increase by two orders of magnitude. This is decidedly suboptimal, since the true mean value can be instantly obtained from the description of the problem. In more complicated scenarios, the accuracy and reliability of simulations is usually not analyzed, which can lead to problems such as

- overlooking low-probability events due to encountering several outcomes with similar results, or



**Figure 1: Approximate number of simulations required to build various confidence intervals for the mean payoff in the coin flip scenario**

- unreliable predictions on some parameter configurations, due to tuning the number of simulations to give consistent results for another parameter configuration, which happens to have more consistent behavior in general.

## 3. BOUNDS ARITHMETIC

Rather than manually performing theoretical analysis on every system of interest, it would be preferable to use an automated framework. Since quantities with error bounds may be introduced in several places throughout the analysis, some basic operations must be redefined to work on sets of possible values instead of just individual values. This section discusses one possible choice for representing and working with error bounds.

Suppose some quantity  $x$  lies on domain  $X$ ; in general, sets of valid values for  $x$  will lie on  $2^X$ . However, even for discrete domains  $X$ , it can be infeasible to store and work with these subsets directly; instead, some subset  $\mathcal{E} \subseteq 2^X$  should be selected with elements that are easier to represent. Then, for every basic function  $f : X \rightarrow Y_X$  involving  $X$ , a “superimage”  $\mathbf{g} : \mathcal{E} \rightarrow Y_{\mathcal{E}}$  of the function should be defined with the following property:

$$x \in E \implies f(x) \in \mathbf{g}(E), \forall x \in X, \forall E \in \mathcal{E}. \quad (5)$$

Note the one-sided implication; essentially, superimages of an operation are allowed to decrease the specificity of an answer, but not to increase it, as this could lead to incorrect error bounds.

We discuss the specific case where  $X$  is an ordered field,<sup>1</sup> and  $\mathcal{E}$  is the set of intervals over  $X$ ; extending this definition to handle vectors spaces over the reals is relatively straightforward. Intervals can be straightforwardly represented as two endpoints from  $X$ , each tagged as open (exclusive) or closed (inclusive). Addition is exactly what would be expected: lower bounds are summed together, and upper

<sup>1</sup>Ordered fields, also known as formally real fields, are algebraic structures formalizing certain properties about addition, negation, multiplication, inversion, and ordering of numbers. One important instance of an ordered field is the real numbers.

bounds are summed together; the type of a bound is inclusive iff both of the associated summands corresponded to inclusive bounds. The symbols for common operations will be overloaded to operate between element pairs of scalar, bound, or mixed types. Here are some examples:

$$[a, b] + [c, d] := [a + c, b + d]; \quad (6)$$

$$[a, b) + [c, d] := [a + c, b + d); \quad (7)$$

$$(a, b] + [c, d) := (a + c, b + d); \text{ and} \quad (8)$$

$$(a, b) + (c, d) := (a + c, b + d). \quad (9)$$

An interpretation of the second result is that, for any  $x \in [a, b)$  and  $y \in [c, d]$ , it is guaranteed that  $(x + y) \in [a + c, b + d)$ . Negation flips the sign and order of the bounds:

$$-[a, b] := [-b, -a]; \quad (10)$$

$$-[a, b) := (-b, -a]; \quad (11)$$

$$-(a, b] := [-b, -a); \text{ and} \quad (12)$$

$$-(a, b) := (-b, -a). \quad (13)$$

Multiplication is more complicated than the previous operations. Here is an example for closed intervals:<sup>2</sup>

$$[a, b] \cdot [c, d] := [\min\{ac, ad, bc, bd\}, \max\{ac, ad, bc, bd\}]. \quad (14)$$

Bounds functions for other operations can be formed by composing the above operations. However, the resulting functions are not necessarily tight. For example, consider the function  $f(x) = x - x$ . It is clear that, for all  $x$ ,  $f(x) = 0$ , and so the tightest bounds function  $\mathfrak{g}(E)$  is equal to  $\{0\}$  for all  $E$ . Suppose that it is known that  $x \in [a, b]$ ; the interval subtraction operation defined using the above operations produces the bounds  $[a, b] - [a, b] := [a - b, b - a]$ , rather than the optimal result  $[0, 0]$ . It is important to try to limit the number and effects of these bounds-widening operations.

## 4. OPINION DYNAMICS MODEL

To illustrate the proposed analysis techniques, we will focus on the Bounded Confidence model of opinion dynamics with a confidence interval width of  $\varepsilon = 1$ . In this model, an agent  $i$ 's opinion  $X_i^t$  on some subject at time  $t$  is modeled as a real number on the interval  $[0, 1]$ . Whenever two agents  $i$  and  $j$  interact, their opinions become more similar:

$$X_i^{t+1} = X_i^t + \mu(X_j^t - X_i^t), \quad X_j^{t+1} = X_j^t + \mu(X_i^t - X_j^t). \quad (15)$$


An individual's opinion only changes through interaction with others. We use the fairly common interaction scheme in which two distinct agents, chosen randomly with uniform probability from all possible pairs, interact at each time step (with all others retaining their old opinions). Half of the  $n_a$  agents in the society will be initialized with an opinion of  $o \in [0, 1]$ , and the others will start with an opinion of 0; it is assumed that  $n_a$  is even.

Let  $\mathbf{X}^t = (X_i^t)_{i=1}^{n_a}$  be the vector of all agent opinions at time  $t$ . Let

$$\mathcal{M}_j^{t'} = \frac{1}{n_a} \sum_{i=1}^{n_a} (X_i^t)^j \quad (16)$$

be the  $j$ th raw moment of opinions in a population at time  $t$ ; for example,  $\mathcal{M}_1^{t'}$  is the average opinion of the society at

<sup>2</sup>Intervals including open endpoints require some additional checks to ensure that valid results are returned, but are otherwise addressed in a similar manner.



**Figure 2: The agent opinions over time for one opinion dynamics simulation run,  $n_a = 10$ ,  $\mu = 0.1$ ,  $o = 1$**

time  $t$ . Usually, moments are simply constants; however, in this case, each  $\mathcal{M}_j^{t'}$  is a random variable, since it can vary across different runs of a simulation. We use terms such as “population” moments to refer to expectations across agents in a population, such as  $\mathcal{M}_j^{t'}$ , and “run” moments to refer to expectations across different runs of a simulation. Some properties of this system are that:

- $\mathbf{X}^t \xrightarrow{\text{a.s.}} (\mathcal{M}_1^{t'})_{i=1}^{n_a} = (0.5)_{i=1}^{n_a}$  as  $t \rightarrow \infty$  — the population approaches consensus almost surely, and
- $\mathcal{M}_1^{t'} = \mathcal{M}_1^{t'0} = 0.5$  for all  $t$  — opinion changes from interactions do not change the population mean.

Figure 2 shows how agent opinions change over time during a single simulation run.

## 5. ANALYTICAL PREDICTION

Let  $\mathcal{M}^{t'} = (\mathcal{M}_j^{t'})_{j=0}^{\infty}$  be a zero-indexed vector of all raw population moments at time  $t$ . Since the agents' opinions have bounded support, these moments are guaranteed to exist; more specifically, they are supported on the interval  $[0, 1]$ . The zeroth order moment  $\mathcal{M}_0^{t'}$ , which is always 1, is included as a check on the validity of predictions. Recall that the population moments are random variables, since they can take on different values in different runs. We are interested in the probability distribution across runs of these population moments, which can be described by the moments of  $\mathcal{M}^{t'}$ . These population-run moments will be referenced with multi-indices<sup>3</sup>: let

$$\mathfrak{m}_\alpha^{t'} = \mathbb{E}_{\text{runs}} [(\mathcal{M}^{t'})^\alpha] = \mathbb{E}_{\text{runs}} \left[ \prod_j (\mathcal{M}_j^{t'})^{\alpha_j} \right] \quad (17)$$

be the  $\alpha$ th moment across runs of the population moments for (finite-dimensional) multi-index  $\alpha$ , where  $\mathbb{E}_{\text{runs}}$  is the expectation operator across the probability space of theoretical runs of the system. For example, at time  $t$ :

<sup>3</sup>An  $n$ -dimensional multi-index  $\alpha \in \mathbb{N}_0^n$  generalizes the concept of a power to vectors  $\mathbf{x} \in \mathbb{R}^n$ . It has order  $|\alpha| = \sum_i \alpha_i$ , factorial  $\alpha! = \prod_i \alpha_i!$ , and power operation  $\mathbf{x}^\alpha = \prod_i x_i^{\alpha_i}$ . We use zero-indexing of multi-index components. Multi-indices are partially ordered:  $\alpha \leq \beta \Leftrightarrow \bigwedge_{i=0}^{n-1} (\alpha_i \leq \beta_i)$ .

- $\mathbf{m}'_{(k)} = \mathbb{E}_{\text{runs}} [(\mathcal{M}'_0)^k] = 1$  is the  $k$ th moment across runs of the zeroth population moment, which is always 1
- $\mathbf{m}'_{(0,1)} = \mathbb{E}_{\text{runs}} [(\mathcal{M}'_1)^1] = 0.5$  is the mean across runs of the population mean, which is always 0.5
- $\mathbf{m}'_{(0,2)} = \mathbb{E}_{\text{runs}} [(\mathcal{M}'_1)^2] = 0.25$  is the raw variance across runs of the population mean, which is always 0.25
- $\mathbf{m}'_{(0,0,1)} = \mathbb{E}_{\text{runs}} [(\mathcal{M}'_2)^1]$  is the mean across runs of the raw population variance, which can be loosely interpreted as the expected level of disagreement in a population
- $\mathbf{m}'_{(0,1,1)} = \mathbb{E}_{\text{runs}} [(\mathcal{M}'_1)^1(\mathcal{M}'_2)^1]$  is the raw covariance of the population mean and raw population variance, linked to relationships between the average opinion in a society and the level of disagreement.

Note that  $\mathbf{m}'_\alpha$  is not a random variable; it does not vary from agent to agent, nor from run to run. Thus, there is no need to rely on confidence intervals or similar “probabilistic” error bounds to quantify error in  $\mathbf{m}'_\alpha$ : simple  $L_1$  error bounds will be sufficient. The goal of the proposed analytical method is to predict  $\mathbf{m}'_\alpha$ 's of interest with these hard error bounds.

Recall that the population is initialized with the same configuration every run, in which half of the agents hold opinion  $o$  (where  $o \in [0, 1]$ ), and the other half hold opinion 0. Thus, we can find simple equations for the deterministic values  $\mu'^0$  of the population moments at time 0:

$$\mathcal{M}'_j{}^0 = \mu'^0 = \frac{1}{2}(o)^j + \frac{1}{2}(0)^j = \frac{1}{2}(o)^j. \quad (18)$$

Since these population moments are nonrandom, the population-run moments can be written without an expectation:

$$\mathbf{m}'_\alpha{}^0 = \mathbb{E}_{\text{runs}} [(\mathcal{M}'^0)^\alpha] = (\mu'^0)^\alpha \quad (19)$$

$$= \prod_j (\mu'^0)^{\alpha_j} = \prod_j \left[ \frac{1}{2}(o)^j \right]^{\alpha_j} = \frac{(o)^{\sum_j j \alpha_j}}{2^{|\alpha|}}. \quad (20)$$

Of course, the starting population configuration could also be sampled from a distribution.

Let us now examine how to handle transitions in this case. Let  $B_1^t \sim_{\text{pop}} \mathbf{X}^t$  be the opinion of the first agent selected in the interacting pair at time  $t$ . Since the selection is uniform and not influenced by the agent’s profile, it follows the same distribution as the population, and has the same moments,  $\mathcal{M}'^t$ :

$$\mu'_{j, B_1^t} = \mathcal{M}'_j{}^t, \quad (21)$$

for all  $j$  such that the  $\mathcal{M}'_j{}^t$  exists. However, the same is not true for the bias,  $B_2^t$ , of the second agent selected, because the first agent is disallowed from interacting with itself. Instead, the  $j$ th moment is given by

$$\mu'_{j, B_2^t} = \frac{n_a \mathcal{M}'_j{}^t - (B_1^t)^j}{n_a - 1}. \quad (22)$$

In the Bounded Confidence model, interactions between agents are deterministic, and for the special case  $\varepsilon = 1$ , the new values  $B_1^{t'}$  and  $B_2^{t'}$  can be written in a simpler form:

$$B_1^{t'} = B_1^t + \mu(B_2^t - B_1^t) = (1 - \mu)B_1^t + \mu B_2^t, \quad (23)$$

$$\text{and } B_2^{t'} = B_2^t + \mu(B_1^t - B_2^t) = (1 - \mu)B_2^t + \mu B_1^t. \quad (24)$$

The  $j$ th population moment at the next time step can then be expressed by

$$\mathcal{M}'_j{}^{t+1} = \frac{n_a \mathcal{M}'_j{}^t - (B_1^t)^j - (B_2^t)^j + (B_1^{t'})^j + (B_2^{t'})^j}{n_a} \quad (25)$$

$$= \mathcal{M}'_j{}^t + \frac{1}{n_a} B_{\Delta, j}^t, \quad (26)$$

where

$$B_{\Delta, j}^t = -(B_1^t)^j - (B_2^t)^j + \sum_{i=0}^j \binom{j}{i} (1 - \mu)^i \mu^{j-i} [(B_1^t)^i (B_2^t)^{j-i} + (B_2^t)^i (B_1^t)^{j-i}]. \quad (27)$$

Note that this expression for  $\mathcal{M}'_j{}^{t+1}$  is symmetric — simultaneously substituting  $B_1^t$  for  $B_2^t$ , and  $B_2^t$  for  $B_1^t$ , produces an equivalent expression. In this case, the “roles” of the two agents in the interaction can be said to be equivalent. For interactions between more than two agents, identifying such symmetry can enable significant improvements in memory use and computation time.

Applying this formula, we can find, for example, that the mean opinion value in a population remains unchanged by this interaction:

$$\begin{aligned} \mathcal{M}'_1{}^{t+1} &= \mathcal{M}'_1{}^t + \frac{1}{n_a} \left\{ -B_1^t - B_2^t \right. \\ &\quad \left. + \left[ \binom{1}{0} \mu (B_2^t + B_1^t) + \binom{1}{1} (1 - \mu) (B_1^t + B_2^t) \right] \right\} \quad (28) \end{aligned}$$

$$= \mathcal{M}'_1{}^t + \frac{-B_1^t - B_2^t + [\mu + (1 - \mu)](B_1^t + B_2^t)}{n_a} \quad (29)$$

$$= \mathcal{M}'_1{}^t. \quad (30)$$

However, the analogous expression for the second population moment maintains a dependency on  $B_1^t$  and  $B_2^t$ :

$$\begin{aligned} \mathcal{M}'_2{}^{t+1} &= \mathcal{M}'_2{}^t + \frac{1}{n_a} \left\{ -(B_1^t)^2 - (B_2^t)^2 \right. \\ &\quad \left. + \binom{2}{0} \mu^2 [(B_1^t)^2 + (B_2^t)^2] \right. \\ &\quad \left. + \binom{2}{1} \mu(1 - \mu) [B_1^t B_2^t + B_2^t B_1^t] \right. \\ &\quad \left. + \binom{2}{2} (1 - \mu)^2 [(B_2^t)^2 + (B_1^t)^2] \right\}. \quad (31) \end{aligned}$$

Combining coefficients of monomials of  $B_1^t$  and  $B_2^t$  yields:

$$\begin{aligned} \mathcal{M}'_2{}^{t+1} &= \mathcal{M}'_2{}^t + \frac{1}{n_a} \left\{ 4\mu(1 - \mu) B_1^t B_2^t \right. \\ &\quad \left. + [\mu^2 + (1 - \mu)^2 - 1] [(B_1^t)^2 + (B_2^t)^2] \right\}. \quad (32) \end{aligned}$$

Expanding polynomials of  $\mu$  and then factoring gives a much simpler form:

$$\mathcal{M}'_2{}^{t+1} = \mathcal{M}'_2{}^t - \frac{2\mu(1 - \mu)}{n_a} (B_1^t - B_2^t)^2 \quad (33)$$

$$= \mathcal{M}'_2{}^t - \frac{2\mu(1 - \mu)}{n_a} (B_2^t - B_1^t)^2. \quad (34)$$

Since the values of  $B_1^t$  and  $B_2^t$  do not directly affect the behavior of the system at subsequent time steps (implied by

the fact that the system is Markovian), and the predictions we would like to make do not involve these values either, these dependencies can be safely marginalized out of representations of  $\mathcal{M}'^{t+1}$  with no effect on accuracy. This gives a simpler form for the  $\alpha$ th population-run moment:

$$\mathbf{m}'_{\alpha}{}^{t+1} = \mathbb{E}_{\text{runs}} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ (\mathcal{M}'^{t+1})^{\alpha} \right] \quad (35)$$

$$= \mathbb{E}_{\text{runs}} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ \prod_j (\mathcal{M}'_j{}^{t+1})^{\alpha_j} \right] \quad (36)$$

$$= \mathbb{E}_{\text{runs}} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ \prod_j \left( \mathcal{M}'_j{}^t + \frac{1}{n_a} B_{\Delta,j}^t \right)^{\alpha_j} \right] \quad (37)$$

$$= \mathbb{E}_{\text{runs}} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ \sum_{0 \leq |\beta| \leq |\alpha|} (\mathcal{M}'^t)^{\beta} \left( \frac{1}{n_a} B_{\Delta}^t \right)^{\alpha - \beta} \right] \quad (38)$$

$$= \sum_{0 \leq \beta \leq \alpha} \frac{\mathbb{E}_{\text{runs}} \left[ (\mathcal{M}'^t)^{\beta} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ (B_{\Delta}^t)^{\alpha - \beta} \right] \right]}{n_a^{|\alpha| - |\beta|}}, \quad (39)$$

where  $B_{\Delta}^t$  is a vector of the monomials  $B_{\Delta,j}^t$  for all  $j$  of interest. The expression  $(B_{\Delta}^t)^{\alpha - \beta}$  can be expanded further into a polynomial of  $B_1^t$  and  $B_2^t$  (with degree  $\sum_j j(\alpha_j - \beta_j)$ ), allowing each  $\mathbf{m}'_{\alpha}{}^{t+1}$  to be expressed as a sum over expectations of monomials of  $\mathcal{M}'^t$ ,  $B_1^t$ , and  $B_2^t$ . While the general expression for  $\mathbf{m}'_{\alpha}{}^{t+1}$  is quite complicated, computing it for particular multi-indices  $\alpha$  is straightforward using a basic computer algebra system. For purposes of applying the expectations  $\mathbb{E}_{B_1^t}$  and  $\mathbb{E}_{B_2^t}$ , the moments  $\mathcal{M}'^t$  can be considered a constant — these expectations produce results containing  $\mathcal{M}'_j{}^t$  terms and factors, but the outcomes of the underlying probability space do not determine any  $\mathcal{M}'_j{}^t$ . Applying the expectation  $\mathbb{E}_{\text{runs}}$ , we obtain an expression for  $\mathbf{m}'_{\alpha}{}^{t+1}$  in terms of elements from  $\mathbf{m}'^t$ .

To illustrate the above process, let us find an expression for  $\mathbf{m}'_{(0,0,1)}{}^{t+1}$ , the mean-over-runs of the uncentralized variance-over-agents of opinion values:

$$\mathbf{m}'_{(0,0,1)}{}^{t+1} = \sum_{(0,0,0) \leq \beta \leq (0,0,1)} \frac{\mathbb{E}_{\text{runs}} \left[ (\mathcal{M}'^t)^{\beta} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ (B_{\Delta}^t)^{(0,0,1) - \beta} \right] \right]}{n_a^{1 - |\beta|}} \quad (40)$$


$$= \frac{\mathbb{E}_{\text{runs}} \left[ (\mathcal{M}'^t)^{(0,0,0)} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ (B_{\Delta}^t)^{(0,0,1) - (0,0,0)} \right] \right]}{n_a^{1 - (0)}} + \frac{\mathbb{E}_{\text{runs}} \left[ (\mathcal{M}'^t)^{(0,0,1)} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ (B_{\Delta}^t)^{(0,0,1) - (0,0,1)} \right] \right]}{n_a^{1 - (1)}} \quad (41)$$

$$= \frac{-2\mu(1 - \mu)}{n_a} \mathbb{E}_{\text{runs}} \mathbb{E}_{B_1^t} \mathbb{E}_{B_2^t} \left[ (B_1^t)^2 - 2B_1^t B_2^t + (B_2^t)^2 \right] + \mathbb{E}_{\text{runs}} [\mathcal{M}'_2{}^t] \quad (42)$$

$$= \frac{-2\mu(1 - \mu)}{n_a} \mathbb{E}_{\text{runs}} \mathbb{E}_{B_1^t} \left[ (B_1^t)^2 - 2B_1^t \frac{n_a \mathcal{M}'_1{}^t - (B_1^t)^1}{n_a - 1} + \frac{n_a \mathcal{M}'_2{}^t - (B_1^t)^2}{n_a - 1} \right] + \mathbb{E}_{\text{runs}} [\mathcal{M}'_2{}^t] \quad (43)$$

$$= \frac{-4\mu(1 - \mu)}{n_a - 1} \mathbb{E}_{\text{runs}} \left[ \mathcal{M}'_2{}^t - (\mathcal{M}'_1{}^t)^2 \right] + \mathbb{E}_{\text{runs}} [\mathcal{M}'_2{}^t] \quad (44)$$

$$= \frac{-4\mu(1 - \mu)}{n_a - 1} (\mathbf{m}'_{(0,0,1)}{}^t - \mathbf{m}'_{(0,2,0)}{}^t) + \mathbf{m}'_{(0,0,1)}{}^t. \quad (45)$$




**Figure 3: Simulated raw variance for 150 individual simulations, and predicted mean raw variance; constant half-0-half-1 population initialization,  $\varepsilon = 1$ ,  $\mu = \frac{1}{2}$ ,  $n_a = 100$ .**

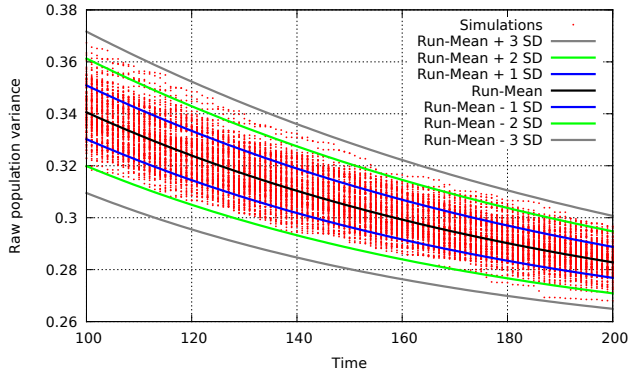
All of the (marginalized) population-run moments at time  $t + 1$  can be found in terms of population-run moments at time  $t$  using similar computations. Note that the expression for  $\mathbf{m}'_{\alpha}{}^{t+1}$  is a weighted sum over  $\mathbf{m}'_{\beta}{}^t$  terms where  $\sum_j j\beta_j = \sum_j j\alpha_j$ . The form of the solution as a weighted sum over previous population-run moments follows from the linearity of the expectation operations and form of their arguments as polynomials over population moments and agent profiles — this holds in many, but not all, cases. The pattern in index restrictions can be explained by the “unit compatibility” of the system, and can be used to guarantee better computation time and space than in the general case.

While the above prediction for  $\mathbf{m}'_{(0,0,1)}{}^{t+1}$  is exact — it has no error, and no associated confidence level or “probability” of being correct, a brief comparison to simulated data may ease any doubts about its fidelity. Figure 3 shows the simulated raw variance in opinion values across a population for 150 individual simulations, as well as the above prediction for the mean-across-runs of the raw population variance. The close (visual) correspondence between the mean of the simulation results and the predicted mean helps confirm that the analytical prediction is producing valid results. In this case, the raw population variance in the trials appears to be fairly consistent across runs, and it is reasonable to assume that the sample mean of the raw population variance would have tighter confidence interval bounds than the sample mean of the payoff in the coin flip scenario for the same confidence level and number of simulations. However, the deviations in the raw sample variance are still significant, and several additional orders of accuracy and reliability are needed for the purposes of making economic decisions or determining parameters for physical design based on this number. Performing comparative hypothesis tests of the mean of the raw population variance at different time steps will also generally require many more simulations. However, the central limit theorem also applies in this case, and the exponential batch-size requirement for degree of accuracy still holds.

We can use higher-order run-moments to gain additional information about the variability in system behavior. For example, the run standard deviation of the raw population



**Figure 4: Simulated raw variance for 150 individual simulations, and predicted mean raw variance plus or minus a specified number of run standard deviations; constant half-0-half-1 population initialization,  $\varepsilon = 1$ ,  $\mu = \frac{1}{2}$ ,  $n_a = 100$ .**



**Figure 5: Simulated raw variance for 150 individual simulations, and predicted mean raw variance plus or minus a specified number of run standard deviations; detail of inner region; constant half-0-half-1 population initialization,  $\varepsilon = 1$ ,  $\mu = \frac{1}{2}$ ,  $n_a = 100$ .**

variance is given by

$$\begin{aligned} & \sqrt{\mathbb{E}_{\text{runs}} [(\mathcal{M}_2^t - \mathbb{E}_{\text{runs}} [\mathcal{M}_2^t])^2]} \\ &= \sqrt{\mathbb{E}_{\text{runs}} [(\mathcal{M}_2^t)^2 - 2(\mathcal{M}_2^t) (\mathbf{m}_{(0,0,1)}^t) + (\mathbf{m}_{(0,0,1)}^t)^2]} \quad (46) \end{aligned}$$

$$= \sqrt{\mathbf{m}_{(0,0,2)}^t - (\mathbf{m}_{(0,0,1)}^t)^2}. \quad (47)$$

Figures 4 and 5 show the correspondence between the predictions for the run standard deviation of the raw population variance and the spread of the simulated raw population variance. Tools such as Chebyshev’s inequality give an idea of how much variability to expect:  $\geq 75\%$  of trials within two standard deviations of the mean,  $\geq 88\%$  within three, and so on. Edgeworth series provide a method to incorporate the information provided by higher-order moments. These tools deal with the true run distribution of the raw population variance, rather than the distribution of simulation results, though, so some violations are possible.

In some cases, we may be interested in high-order run-moments of low-order population moments, but not in low-order run-moments of high-order population moments. For example, the raw kurtosis of the population mean may be of interest while the mean of the raw population kurtosis is not. However, the latter quantity may influence the future values of the former. In this situation, there are two straightforward choices:

- track the influential, but “uninteresting”, population-run moments, then discard them; or
- approximate the effect of “uninteresting” population-run moments using lower-order ones, sacrificing some accuracy in exchange for lower computation costs.

For example, suppose that we want to track the second-order (co)moments-across-runs of the zeroth through second raw population moments. Performing analysis similar to that above, we find that the raw covariance of the population mean and raw population variance at time  $t + 1 \geq 1$  is

$$\mathbf{m}_{(0,1,1)}^{t+1} = \frac{1}{n_a - 1} \mathbf{m}_{(0,3,0)}^t + \frac{n_a - 2}{n_a - 1} \mathbf{m}_{(0,1,1)}^t \quad (48)$$

$$= \frac{1}{n_a - 1} \mathbb{E}_{\text{runs}} [(\mathcal{M}_1^t)^3 + (n_a - 2) (\mathcal{M}_1^t) (\mathcal{M}_2^t)]. \quad (49)$$

Suppose that  $\mathbf{m}_{(0,3,0)}^t$ , the raw skewness of the population mean at time  $t$ , is not of interest, and is not tracked for the sake of computational efficiency. Then the expectation  $\mathbb{E}_{\text{runs}} [(\mathcal{M}_1^t)^3]$  must be approximated using the other population-run moments that are available. We know that the support of  $\mathcal{M}_1^t$  is a subset of  $[0, 1]$ , since it is the mean of values that all lie on  $[0, 1]$ . Thus, one bounds function for  $(\mathcal{M}_1^t)^3$  is  $[0, 1] \cdot (\mathcal{M}_1^t)^2$ , and a bounds function for  $\mathbf{m}_{(0,1,1)}^{t+1}$  is

$$\frac{[0, 1]}{n_a - 1} \mathbf{m}_{(0,2,0)}^t + \frac{n_a - 2}{n_a - 1} \mathbf{m}_{(0,1,1)}^t. \quad (50)$$

For systems with many agents, the potential error introduced by this approximation at any individual time step will be small, but will grow exponentially over time. However, using bounds arithmetic to find support supersets actually preserves the exact answer in this case. Consider  $o = 1$ . The support of  $\mathcal{M}_0^0$  is  $R_{\mathcal{M}_0^0} = \{\frac{1}{2}\}$ , from the static initial population configuration. The support of the population mean at time  $t + 1$  can then be calculated recursively:

$$R_{\mathcal{M}_0^{t+1}} = R_{\mathcal{M}_0^t + \frac{1}{n_a} [-1 + \mu + (1 - \mu)(B_1^t + B_2^t)]} = R_{\mathcal{M}_0^t}. \quad (51)$$

In this case, the result is rather simple, due to the fact that the population mean  $\mathcal{M}_1^t$  is constant across all  $t$ , and was initialized to a constant value of  $\frac{1}{2}$ . (In fact, any run-moment  $\mathbf{m}_{(0,j)}^t$  of the population mean is also a constant.) Consequently, we find that  $R_{\mathcal{M}_0^t} = \{\frac{1}{2}\}$  for all  $t$ . In more complex cases, we can find supersets of supports (and in some cases, the supports themselves) using the error-bound arithmetic described earlier.

## 6. COMPLEXITY

### 6.1 Population size-independence

The space and time complexity of the proposed method are directly dependent on the number of agents in each interaction, rather than the number of agents in a population.



In systems such as the Bounded Confidence model and the repeated gameplay scenario, it takes the same amount of time (assuming floating-point math) to perform predictions for populations of size two as it does for those with a million agents or more, enabling prediction for larger systems than is feasible with simulations. Unlike typical mean field approaches, though, this independence is not achieved through approximating the population with the “infinite-agent” case. When agents change their profiles, migrate between populations, or are selected without replacement, the effects on population moments incorporate an appropriate dependency on the population size that generally takes a constant amount of time (e.g., floating-point multiplication).

## 6.2 Composing transitions

Note that the transitions of the moments in the system can be represented in matrix form. Let  $\mathbf{A}_t^{t+1}$  be such a transition matrix for the time step from  $t = 0$  to  $t = 1$ , with each row corresponding to a new moment  $\mathbf{m}'_\alpha^{t+1}$ , and each column corresponding to an old moment  $\mathbf{m}'_\alpha^t$ . Let us assume that the  $\alpha$ s of interest are the same at every time step, and that a consistent ordering over these  $\alpha$ s is used. The moments at time  $t = 1$  are given by the matrix-vector product

$$\mathbf{m}'^{t+1} = \mathbf{A}_t^{t+1} \mathbf{m}'^t. \quad (52)$$

If  $\mathbf{A}_t^{t+1}$  and  $\mathbf{m}'^t$  are dense, then evaluating this matrix-vector product takes  $O(|\mathcal{A}|^2)$  time, where  $\mathcal{A}$  is the set of  $\alpha$ s of interest. If the system is time-homogeneous, and there is an exact transition matrix  $\mathbf{A}$ , then the  $t$ th set of population-run moments is given by

$$\mathbf{m}'^t = (\mathbf{A})^t \mathbf{m}'^0, \quad (53)$$

which can be evaluated by either repeated matrix-vector products, taking  $O(|\mathcal{A}|^2 t)$  time, or a matrix power followed by a matrix-vector product, usually with a time complexity of

$$O(m(|\mathcal{A}|) \lg t + |\mathcal{A}|^2), \quad (54)$$

where  $m(|\mathcal{A}|)$  is the cost of multiplying two  $|\mathcal{A}| \times |\mathcal{A}|$  matrices. For naïve matrix multiplication,  $m(|\mathcal{A}|) = |\mathcal{A}|^3$ ; other options include Strassen’s algorithm [12] and any relevant sparse matrix multiplication algorithms. The  $\lg t$  term is the result of a common trick used to reduce the complexity of power operations. To illustrate, consider the task of finding  $\mathbf{A}^8$ ; this can either be computed as


$$\mathbf{A}^8 = ((((((\mathbf{A}\mathbf{A})\mathbf{A})\mathbf{A})\mathbf{A})\mathbf{A})\mathbf{A})\mathbf{A}), \quad (55)$$

using seven matrix multiplications, or as  $\mathbf{A}^8 = ((\mathbf{A}^2)^2)^2$ , using three. Eigendecomposition-based approaches offer additional candidates for changing the complexity of prediction.

## 7. INFERENCE AND OPTIMIZATION

### 7.1 Inferring moments of initial distribution

Given an estimate  $\hat{\mathbf{m}}'^t$  of the population-run moments  $\mathbf{m}'^t$  at time  $t$  (using Monte Carlo simulations or some other method), it is fairly simple to obtain an approximation of the initial population configuration, or run-moments  $\mathbf{m}'^0$  of the initial population configuration distribution. Since  $\mathbf{m}'^t = \mathbf{A}_0^t \mathbf{m}'^0$  for an appropriate transition matrix  $\mathbf{A}_0^t$ ,  $\mathbf{m}'^0 = (\mathbf{Z}_0^t) \mathbf{m}'^t$  for any  $\mathbf{Z}_0^t$  such that  $\mathbf{Z}_0^t \mathbf{A}_0^t = \mathbf{I}$ , and a



**Figure 6: Run mean of population raw variance vs.  $\mu$  for times  $0 \leq t \leq 10$ ,  $n_a = 10$ , constant half-0-half-1 initialization.**

rough approximation of a possible value of  $\mathbf{m}'^0$  is given by

$$\hat{\mathbf{m}}'^0 = \mathbf{Z}_0^t \hat{\mathbf{m}}'^t. \quad (56)$$

### 7.2 Inferring model parameters

We can perform approximate inference simply using equation solvers by incorporating model parameters (e.g.,  $\mathbf{m}'^0$  or  $\mu$  in the BC model) into predictions symbolically. For example, consider the BC model with  $\varepsilon = 1$ ,  $n_a = 10$ ,  $o = 1$ , and  $\mu$  unknown, with five agents initialized with opinion 0 and the rest with opinion 1. The run-mean of the raw population variance at time 1 is then

$$\mathbf{m}'_{(0,0,1)}^1 = \frac{1}{2} - \frac{1}{9}\mu + \frac{1}{9}\mu^2. \quad (57)$$

Given empirical data or simulation outputs, the above equation can be used to tune the parameter  $\mu$  so that the model predictions for  $\mathbf{m}'_{(0,0,1)}^1$  match the observations. For example, by substituting the estimate  $\hat{\mathbf{m}}'^1_{(0,0,1)} = 0.48224$  (which was obtained by averaging over the raw population variances at the first time step in  $10^5$  simulations) into the above equation for  $\mathbf{m}'_{(0,0,1)}^1$  and applying the quadratic formula, we obtain two potential (rounded) estimates for  $\mu$ :

$$\hat{\mu}_{1,2} = 0.1997, 0.8003. \quad (58)$$

In this case, the actual  $\mu$  value was  $\frac{1}{5}$ , and the true value for  $\mathbf{m}'_{(0,0,1)}^1$ , found using the proposed analytical techniques, is 0.482. Note that the behavior of the system is identical when  $\mu = \frac{4}{5}$ , so both estimates for  $\mu$  are very close to values that replicate the system’s behavior.

### 7.3 Studying effects of parameters

Once again, by performing prediction of some quantity of interest incorporating parameters symbolically, the effects of parameters can be studied using a single run of the theoretical model, rather than running batches of simulations for a large number of parameter configurations. Figure 6 plots the run mean of the population raw variance against the  $\mu$  parameter in the BC model. The initial value,  $\mathbf{m}'_{(0,0,1)}^0 = \frac{1}{2}$ , represents division of the population into two factions at the edges of the opinion space. For  $0 < \mu < 1$ , the population will almost surely converge to a configuration representing

consensus on the opinion value  $\frac{1}{2}$  as  $t \rightarrow \infty$ , and  $\mathbf{m}_{(0,0,1)}^t$  will approach  $\frac{1}{4}$ . From the graph, we can easily see that when  $\mu = 0$  or  $\mu = 1$ , no progress is made towards consensus of population opinions. The absolute minimum at  $\mu = \frac{1}{2}$ , noted earlier, is clearly seen in this graph, as well as the symmetry in system behavior about  $\mu = \frac{1}{2}$ .

## 7.4 Optimizing performance metrics

Suppose that we would like the run-mean raw population-variance to decrease as quickly as possible over the first time step, and have the ability to select any  $\mu$  value for the system on the interval  $[0, 1]$ . The optimal  $\mu$  value  $\mu^*$  can be calculated using a numerical optimization package:

$$\mu^* = \min_{\mu \in [0,1]} \mathbf{m}_{(0,0,1)}^1 = \min_{\mu \in [0,1]} \left( \frac{1}{2} - \frac{1}{9}\mu + \frac{1}{9}\mu^2 \right) = \frac{1}{2}. \quad (59)$$

Another example of the use of such optimization would be to find an optimal amount to penalize agents for defecting in the prisoner's dilemma game in order to maximize social welfare, minus the cost of enforcing the penalties.

## 8. EXTENSIONS

Suppose we are interested in more complex transitions involving functions involving division, roots, jump discontinuities, or other operations not well-represented by Taylor polynomials. The backbone of this method can remain the same: represent the function with some type of series, such as Laurent, Puiseux, or Fourier series (truncated to some number of terms). Raw moments will be replaced by expectations based on the type of series used; e.g., if function  $f(x)$  is modeled using a Laurent series  $\sum_{i \in \mathbb{Z}} c_i x^i$ , then we will track expectations from  $\{\mathbb{E}[X^i]\}_{i \in \mathbb{Z}}$ . The expectation  $\mathbb{E}[f(X)]$  remains a dot product between series coefficients of  $f$  and corresponding expectations of  $X$ , so time step transitions can be represented as matrix multiplication. Topics for future work include extending the model to handle more complex interaction schemes, such as simultaneous interaction of distinct pairs, migration between populations, and interaction on network topologies.

## 9. CONCLUSIONS

Analytical models with hard error bounds are a viable alternative to Monte Carlo simulation for study of some stochastic multiagent systems. Benefits of analytical approaches include the possibility of obtaining exact answers, independence of computational complexity and population size, and ability to transform the complexity of prediction. Symbolic expressions can be incorporated to efficiently predict effects of model parameters, infer information about past distributions, tune model parameters to fit empirical data, and find parameter values to optimize a performance metric. In the Bounded Confidence model with  $\varepsilon = 1$ , modeling distributions with their raw moments and functions with Taylor polynomials closely predicts future moments of the system, outperforming Monte Carlo simulations and previous analytical approaches in terms of speed and accuracy. By incorporating additional types of mathematical tools to approximate functions and distributions, this analytical technique can be extended to make effective predictions about the behavior of a wide range of systems.

## 10. ACKNOWLEDGMENTS

This material is based upon work partially supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 0946825, a Woobank Fellowship from the University of Tulsa, and the Tulsa Undergraduate Research Challenge.

## 11. REFERENCES

- [1] R. S. Bucy. Linear and nonlinear filtering. *Proceedings of the IEEE*, 58(6):854–864, 1970.
- [2] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(01–04):87–98, 2000.
- [3] J. Gómez-Serrano, C. Graham, and J.-Y. Le Boudec. The bounded confidence model of opinion dynamics. *Mathematical Models and Methods in Applied Sciences*, 22(02), 2012.
- [4] M. P. Hassell. *The dynamics of arthropod predator-prey systems*. Princeton University Press, 1978.
- [5] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.
- [6] S. Huet, G. Deffuant, and S. Huet. Bounded Confidence with rejection: The infinite population limit with perfect uniform initial density. *Modelling and leArning in Social and Human Sciences (MASHS) 2008*, 2008.
- [7] L. P. Kadanoff. More is the same; phase transitions and mean field theories. *Journal of Statistical Physics*, 137(5-6):777–797, 2009.
- [8] R. E. Kalman and R. S. Bucy. New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83(3):95–108, 1961.
- [9] V. N. Kolokoltsov. Nonlinear Markov games on a finite state space (mean-field and binary interactions). *International Journal of Statistics and Probability*, 1(1):p77, 2012.
- [10] Q. Pan, F. Yang, L. Ye, Y. Liang, and Y.-m. Cheng. Survey of a kind of nonlinear filters — UKF. *Control and Decision*, 20(5):481, 2005.
- [11] S. Solomon. Generalized Lotka-Volterra (glv) models of stock markets. *Advances in Complex Systems*, 3(01–04):301–322, 2000.
- [12] V. Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13(4):354–356, 1969.
- [13] G. Szabó, J. Vukov, and A. Szolnoki. Phase diagrams for an evolutionary prisoner's dilemma game on two-dimensional lattices. *Physical Review E*, 72(4):047107, 2005.
- [14] G. Welch and G. Bishop. An introduction to the Kalman filter. Technical report, University of North Carolina at Chapel Hill, 1995.
- [15] E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione. Discrete opinion dynamics with stubborn agents. *Available at SSRN 1744113*, 2011.
- [16] S. Zhang, T. H. Hansson, and S. Kivelson. Effective-field-theory model for the fractional quantum Hall effect. *Physical Review Letters*, 62(1):82, 1989.