SIAM J. APPL. MATH.
Vol. 82, No. 4, pp. 1246–1266

DOUBLY STOCHASTIC PAIRWISE INTERACTIONS FOR AGREEMENT AND ALIGNMENT*

THOMAS DAGÈS† AND ALFRED M. BRUCKSTEIN†

Abstract. Random pairwise encounters often occur in large populations or groups of mobile agents, and various types of local interactions that happen at encounters account for emergent global phenomena. In particular, in the fields of swarm robotics, sociobiology, and social dynamics, several types of local pairwise interactions were proposed and analyzed leading to spatial gathering, clustering, agreement, or coordinated motion in teams of robotic agents, in animal herds, or in human societies. We here propose a very simple stochastic interaction at encounters that leads to agreement or geometric alignment in swarms of simple agents and analyze the process of converging to consensus. Consider a group of agents whose "states" evolve in time by pairwise interactions: the state of an agent is either a real value (a randomly initialized position within an interval) or a vector that is either unconstrained (e.g., the location of the agent in the plane) or constrained to have unit length (e.g., the direction of the agent's motion). The interactions are doubly stochastic in the sense that, at discrete time steps, pairs of agents are randomly selected and their new states are independently and uniformly set at random in (local) domains or intervals defined by the states of the interacting pair. We show that such processes lead, in finite expected time (measured by the number of interactions that occurred) to agreement in case of unconstrained states and alignment when the states are unit vectors.

Key words. control, gathering, pairwise interactions, decentralized, multiagent

MSC codes. 93A14, 93C05, 93C10, 93E1

DOI. 10.1137/21M1394680

1. Introduction.

We consider a group of N agents with states described by the set $\{x_1, \ldots, x_N\}$. The states x_i can be either real scalar values in some interval $I \subset \mathbb{R}$, vectors in a D-dimensional box $C \subset \mathbb{R}^D$, or unit vectors on the circle S. Typical agents could be people in a social group or a large population, ants in a colony, man-made robots designed to act in a swarm, fish in a school, gas molecules moving around in a container, or even software bots migrating from computer to computer on a network like the internet.

The "state" of an agent may therefore be the opinion of a person on some issue, which can be measured by a real value on $\mathbb{R}$, such as how much you like a product on a (continuous) scale from 0 to 10, where on the political spectrum you are from the far left ($-\infty$) to the far right ($+\infty$), the location of an ant or a robot in a planar domain $C \subset \mathbb{R}^2$, or the direction of motion of a mobile robot. Of course, "state" might also stand for the classical memory content of a (finite-state) machine or a bot-program, but we shall not consider such discrete states here.

We assume that agents are identical in their capabilities and behavior and their states change in time only due to interactions with other agents. The interaction rules must be given and depend only on the current states of the interacting agents, not on their identities. We say that the agents are identical, anonymous, and oblivious. Given some rules of interaction and their timing schedule, we are interested in the evolution of the states of the agents, the evolution reflecting some "emergent behaviors" of the

*Received by the editors January 29, 2021; accepted for publication (in revised form) April 11, 2022; published electronically July 21, 2022.
https://doi.org/10.1137/21M1394680
Funding: This work was partially supported by the Israel Science Foundation grant 2306/18.
†Department of Computer Science, Technion Israel Institute of Technology, Haifa, Israel (thomas.dages@cs.technion.ac.il).

swarm of agents, like agreement in a community of people, or gathering, grouping, or clustering of robotic agents, or coordinated motion due to the alignment of directions of movement in some herd of animals such as bison or insects like locusts.

In this paper we shall analyze a novel, doubly stochastic rule of interaction: we assume that each agent can interact with any other agent at all times, i.e., the complete interaction graph, and that interactions are pairwise only between randomly selected agents and occur at distinct and discrete times, denoted sequentially as t_k for $k = 0, 1, 2 \ldots$. At the interaction moments, two randomly selected agents exchange information about their states and decide to randomly and independently update their own state. This leads to an evolution of the set of states in time and hopefully to convergence to some interesting globally emergent swarming pattern.

2. An overview of previous results.

To set the stage for our proposed rules of interaction and the consequent emerging phenomena, we shall first describe several previously proposed pairwise interaction rules and the results obtained on the consequent global behaviors.

A fundamental problem in distributed computation, as well as in opinion dynamics, is to achieve agreement or consensus via a sequence of local interactions. Suppose that N agents have as initial states randomly selected real values and we would like the agents to eventually agree on a common real value. If agents could see all others' states, they could agree say on the average value of all the states. Suppose, however, that agents can only see neighboring agents or agents connected to them, such as in a given fixed neighborhood graph, determined by the designer of the network or by geometrical proximity. Then agents can average sequentially but locally, only within their neighborhood. The question is, Will this process eventually converge? This problem is not too difficult, and we can in many cases prove that, indeed, in time, the agents will agree on a value that is the average of the initial states thanks to average-preserving rules of motion.

However, consider a stochastic setting in that at distinct time instances t_k ordered increasingly for $k = 0, 1, 2 \ldots$ random pairs of agents are selected, and they replace their states by the averages of the corresponding values. How does this random pairwise averaging process behave? The result is a stochastic state equalizing process, and all states will converge to the average of the initial states of the agents as the motion rule is average preserving. This was first analyzed in [20, 17]. There is extensive research work on such processes under the name of "distributed gossiping." These "gossiping" works analyze the evolution in time of the gathering process to the average value of the initial states. They consider the moment when the normalized distance vector of individual states to the average reaches a value less than a small preset ε. The conclusion of these results is an upper- and lower-bound on the as so defined time to convergence proportional to $\log(\frac{1}{\varepsilon})$, where the constant factors depend on the number of agents N and on the size of the initial spread of the states; see [3, 8].

Random pairwise interactions were also proposed as suitable models for achieving consensus in social studies on opinion dynamics in populations. Several studies proposed to consider societies of individuals as holding initial opinions, or states, quantified by some real values in an interval I of the real line $\mathbb{R}$ and the following rule of evolution: at discrete time instants, if two random members of the society meet, they change their opinion so as to "approach" each other by a deterministic fixed proportion of the size of the difference between their opinions [12, 1, 4]. In a more complex and realistic setting, this is done only if the difference between their opinions is smaller than a certain threshold; otherwise the meeting results in no changes of opinion whatsoever. This later idea is the so-called pairwise interactions based "bounded confidence model" considered by the "French school" of opinion dynamics led by Deffuant; see [7, 11]. These works also lead to clustering and convergence of opinions either to the average of the initial states (if no bounded confidence threshold is assumed) or to several clusters in bounded confidence models [7, 13]. Recently, [10] studied a generalization of the unbounded confidence model in nonconvex opinion spaces, e.g., on the unit circle, but simplified the agent interaction graph to a ring (for a finite number of agents) and more generally to $\mathbb{Z}$ (for an infinite number of agents) coined as the "compass model." As opposed to [10], we focus on the complete graph for interactions.

The idea of using pairwise interaction models in analyzing the emergence of various collective dynamics phenomena is also prevalent in swarm robotics. It was, for example, proposed to model cooperative localization processes in swarms of robots, to improve their self-location estimates by averaging those at random pairwise encounters, when the agents know that they are co-located; hence their estimates should ideally coincide [9]. The idea of encounter averaging of self-location $2D$-vector estimates was there shown to significantly improve the cooperative odometric location evaluations, even under the assumption that the pairwise agent encounters are totally random, which obviously is not the case. This idea is, of course, prevalent in physics. In thermodynamics, one considers gas particles (molecules) moving and colliding, their self-propelling motion manifested as thermal energy and their collisions modeled with several types of deterministic or randomized interactions. The emergent collective behavior in this case is quantified by globally measured properties of the system of molecules like variations in density, temperature, and pressure as functions of container geometry and external, perhaps even temporally changing, factors [2].

3. Overview of results. We consider here three types of problems concerning systems with multiagent pairwise interactions. The interactions that we define are stochastic, and we prove that a desired behavior eventually emerges. We also provide evaluations about the expected time (measured by the number of interactions) to the convergence to a state that is very close to the desired global behavior. The problems are the following: systems of N agents, with states defined by either real numbers in an interval $I \subset \mathbb{R}$ or by real vectors in a D-dimensional box in $\mathbb{R}^D$ or by unit vectors on the circle, are considered to evolve due to random pairwise interactions that result in changes of the states of the interacting agents. The rules of the evolution are the following.

1. The interaction moments are discrete times $t_1, t_2, \ldots$ starting from $t_0 = 0$ when a random initialization is done.
2. At each moment t_k a random pair of agents is selected uniformly from the $\frac{N(N-1)}{2}$ possible pairs of agents.
3. The selected agents i and j uniformly choose new states in the "interval" defined by their states $\{x_i^k, x_j^k\}$ as follows:
   - if the states are real numbers, the "interval" is just chosen to be as $[\min\{x_i^k, x_j^k\}, \max\{x_i^k, x_j^k\}]$;
   - if the states are D-dimensional vectors, the "interval" is then the 1-dimensional line segment $\{\lambda x_i^k + (1-\lambda)x_j^k \mid \lambda \in [0,1]\}$ embedded in $\mathbb{R}^D$;
   - if the states are two unit vectors, the "interval" is the geodesic circle arc between the two points defined by x_i^k and x_j^k on the unit circle.

The main question we address is, How does such a stochastic system evolve in time, measured by the indices of the interaction times $t_0, t_1, t_2 \ldots$ (i.e., $0, 1, 2 \ldots$)? We prove that, in all the cases above, the system gathers the agents' states, with probability one, to a common random point on the real line in the first case, to a single random point in $\mathbb{R}^D$ in the second case, and to a random unit vector in the third case. We also show that the expected time to ε-convergence is finite and provide bounds on it, where ε-convergence is defined as the expected time (or number of iterations) for which the spread of the states of all the agents is smaller than ε.

We list the main results below:

- **Evolution of real values in an interval.** We prove that ε-convergence is achieved almost surely and in finite expected time by deriving an upper-bound of $O(N\log(\frac{N}{\varepsilon^2}))$ on the expected ε-gathering. We illustrate our theory with extensive numerical simulations, and they reveal the quality and tightness of our bound.
- **Evolution of real values in a D-dimensional box.** Similarly to the $1D$ case, we again prove almost sure ε-convergence in finite expected time by deriving an $O(N\log(\frac{DN}{\varepsilon^2}))$ upper-bound. Extensive numerical simulations are performed to show the quality of the bound.
- **Evolution of unit vectors on the unit circle.** We prove almost sure ε-convergence in finite expected time. This problem is significantly more challenging: we here provide a simplistic approach yielding a crude upper-bound, as revealed in extensive experiments. We also provide and detail several promising approaches for deriving a more reasonable upper-bound but leave the refinement issue as an open challenge.

Specifically, we proved the following bounds for the expectation of the time to reach ε-convergence T_ε.

THEOREM 3.1. *In the* 1*-dimensional case, the opinions almost surely converge. Furthermore, if the original opinions lie in a bounded interval $I = [a, b]$, we have*

$$\mathbb{E}(T_\varepsilon) \leq \frac{3}{2}N\ln\left(\frac{N}{\varepsilon^2}\right) + \frac{3}{2}N\left(\ln\left(\frac{(b-a)^2}{2}\right)+1\right).$$

THEOREM 3.2. *In the D-dimensional case, the opinions almost surely converge. Furthermore, if the original opinions lie in a bounded D-dimensional box $C = [a, b]^D$, we have*

$$\mathbb{E}(T_\varepsilon) \leq \frac{3}{2}N\ln\left(\frac{DN}{\varepsilon^2}\right) + \frac{3}{2}N\left(\ln\left(\frac{(b-a)^2}{2}\right)+1\right).$$

LEMMA 3.3. *In the circle case, the unit vector opinions reach a configuration where they are all contained within a half-disk in finite expected time. In other words, denote T_{HD} the time to reach such a configuration; then there exists a finite upper-bound B_N^{HD} depending only on the number of agents N such that*

$$\mathbb{E}(T_{HD}) \leq B_N^{HD}.$$

We prove that $B_N^{HD} = (\frac{81}{2}N(N-1)^2)^{\lfloor \frac{N}{2} \rfloor} + 2\lfloor \frac{N}{2} \rfloor$ is such an upper-bound.

THEOREM 3.4. *In the circle case, the unit vector opinions almost surely converge. Furthermore, we have*

$$\mathbb{E}(T_\varepsilon) \leq B_N^{HD} + \frac{3}{2}N\ln\left(\frac{N}{\varepsilon^2}\right) + \frac{3}{2}N\left(\ln\frac{\pi^2}{2}+1\right).$$

To test the above results, we carry out extensive numerical simulations. They show that the derived bounds in Theorems 3.1 and 3.2 provide a good description of the asymptotic behavior of the expected convergence time when either $\varepsilon \to 0$ or $N \to \infty$. However, the simulations suggest that the bound obtained in Lemma 3.3 is too crude, as we seem to have $\mathbb{E}(T_{HD}) = \Omega(N \ln N)$. The reason for the discrepancy in the bounds is mostly likely due to our naive proof strategy. However, we have not managed to successfully find a more elaborate approach that provides a better bound B_N^{HD}. We leave it as a challenge to reduce the gap in B_N^{HD} to get closer to the empirical results.

4. 1-dimensional case. A preliminary simple model to study social gathering is to assume that people's opinions solely depend on a unique parameter $x_{opinion}$ that lives on the real line $\mathbb{R}$. This model well suits systems where opinions exist along a simple spectrum, with notions of "left-wing" and "right-wing" opinions. The larger $x_{opinion}$ is, the more the opinion is "right-wing," and the smaller $x_{opinion}$ becomes, the more the opinion is "left-wing." It is important to note that, in this model, if $x_{opinion}$ increases, then the opinion becomes more and more "right-wing." We may also assume that the space of opinions is either bounded, which models systems well with limited "left-wing" and "right-wing extremism," or unbounded, which models better systems with unlimited "extremism" in one or both directions.

Mathematically, we define the space of opinions to be an interval $I \subset \mathbb{R}$ of the real line. Our system comprises $N \geq 2$ indistinguishable agents, each with their own opinion: $x_i \in I$ for $i \in \{1, \ldots, N\}$. The initial opinions are random, but we will work conditionally on them, and so they form a given set $\{x_i^0\}_{1 \leq i \leq N}$, where $x_i^0 \in I$ for all i. Opinion dynamics are modeled in discrete time, conditionally on the state of opinions at the previous time step $X_k = (x_1^k, \ldots, x_N^k)^T$. The evolution law at time step t_{k+1} for all $i \in \{1, \ldots, N\}$ is

$$x_i^{k+1} = \mathbb{1}_{i \notin \{A_{k+1}, B_{k+1}\}} x_i^k + \mathbb{1}_{i = A_{k+1}} U_1^{k+1} + \mathbb{1}_{i = B_{k+1}} U_2^{k+1}, \tag{4.1}$$

where (A_{k+1}, B_{k+1}) is a random uniform sampling of two indices of $\{1, \ldots, N\}$ without replacement independent of the past, $\mathbb{1}_{\mathcal{A}}$ is the indicator function of the event $\mathcal{A}$, and, conditionally on X_k, A_{k+1}, and B_{k+1}, U_1^{k+1} and U_2^{k+1} are independent random uniform variables in the interval $[\min\{x_{A_{k+1}}^k, x_{B_{k+1}}^k\}, \max\{x_{A_{k+1}}^k, x_{B_{k+1}}^k\}]$. Concretely, at each time step, two random agents A_{k+1} and B_{k+1} are selected, and they then independently and uniformly resample their opinion in the interval between their previous opinions. See Figure 4.1 for an example.

Note that the updated agents' opinions do not necessarily preserve "ordering." The problem and its analysis would be identical if we forced order preserving updates by renaming agents A_{k+1} and B_{k+1} to be closer to A_k and B_k, respectively.

Before diving into a detailed analysis of our model, we would first like to point out its novelty with respect to the existing literature.

Our model is not a Deffuant-like approach. In the Deffuant model [7], randomly selected pairs of agents are both updated by the same constant "fraction" μ of the interval defined by their opinions. Formally, for the pair of agents i and j selected at time step k, the opinion updates are $x_i^{k+1} = (1-\mu)x_i^k + \mu x_j^k$ and $x_j^{k+1} = \mu x_i^k + (1-\mu)x_j^k$.

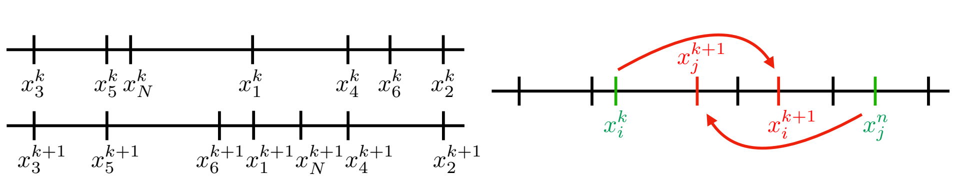

FIG. 4.1. *One-step opinion evolution in the* 1*-dimensional model.*

In a nonstandard randomized Deffuant model, the fraction μ is random, typically uniformly in $[0, \frac{1}{2}]$. In our model, the agents also update their opinions by some "fraction" of the interval of opinions; however, they do not share the same random fraction. Formally, the updates in our model can be written as $x_i^{k+1} = (1-\mu)x_i^k + \mu x_j^k$ and $x_j^{k+1} = \lambda x_i^k + (1-\lambda)x_j^k$, where λ and μ are independently chosen. In particular, we almost surely no longer have $\lambda = \mu$ as in the Deffuant model. Furthermore, our "fractions" λ and μ are sampled in $[0, 1]$ and not just in $[0, \frac{1}{2}]$. This choice implies that the agents cannot only flip their ordering after the update but also that the opinions may lie in the same half of the opinion interval. Both of these behaviors are not possible in the Deffuant-like approaches.

Our model also does not correspond to standard averaging models using doubly stochastic matrices for updates such as in gossiping. Before presenting doubly stochastic matrices, let us rewrite our update model in matrix form. Like the Deffuant approaches, our model is linear, albeit random, in the sense that opinions may be organized into a vector X_k and then updated by a product with a random matrix P: $X_{k+1} = PX_k$. In the Deffuant model, given the pair of agents for update i and j, P is equal to the identity matrix except on the ith and jth rows and columns. These rows and columns are filled with 0 except for two terms per row and per column: $P_{i,i} = \mu$, $P_{i,j} = 1 - \mu$, $P_{j,i} = 1 - \mu$, and $P_{j,j} = \mu$, where μ is either fixed in the standard Deffuant case or random in its randomized variant. In our model, given the pair for update, the difference is that, due to the independent choice of "fractions" for update for each selected agent, we have $P_{i,i} = \mu$, $P_{i,j} = 1 - \mu$, but $P_{j,i} = 1 - \lambda$ and $P_{j,j} = \lambda$, where μ and λ are independent and identically distributed (iid) random variables taking values in $[0, 1]$, hence almost surely $\lambda \neq \mu$. Interestingly, the Deffuant matrix P satisfies the following property: each column and each row sum to 1. The fact that each row sums to 1 is natural and means that updates belong to convex combinations of the two selected opinions. However, the columns summing to 1 implies that the selected agents contribute equally to the sum of all opinions, and thus the update is average preserving. In our model, the update matrix P does not satisfy the property that each column sums to 1. For that to occur, we would necessarily need to have $\lambda = \mu$, which is almost surely never the case.

Generally, matrices with positive entries and with rows and columns summing to 1 are called doubly stochastic matrices. If they are used in an update rule, then the model satisfies that updates belong to convex combinations of other opinions (row assumption) and that the random coefficients in the convex combinations cannot be independent between agents as each agent must also equally contribute to the update of the whole state vector (column assumption). Models using doubly stochastic matrices for updates have been extensively studied [17, 20, 15], especially in the gossiping literature [3, 8]. In our model, we no longer have such doubly stochastic matrices. Our concept of "double stochasticity" is thus different. In some sense, we have even stronger stochasticity due to both the selection of the pair of agents and the two random and independent updates.

Our model is also different from projected consensus models considered, for example, in [15, 18, 19]. In these approaches, a potential is defined such that consensus is reached at its local minima. Consensus is then achieved by performing gradient descent or ascent via a projected gradient step in order to maintain the opinions on the allowable set or manifold. For instance, [18] uses implicitly the sum of squared pairwise distances potential and updates the agents' opinions by subtracting the summed difference between the agents' current opinions and those of their neighbors in the fixed and possibly noncomplete interaction graph. However, for all projected consensus models, due to the gradient step philosophy, given the interaction graph (which may also randomly change), the update step is a given deterministic process. As such, there is no double randomness in these approaches, and they fundamentally differ from our own.

Hence, while our model shares a similar flavor to other approaches in the literature, it fundamentally differs from them as discussed above. We now focus on our model and analyze it.

PROPOSITION 4.1. *The quantity* $\max_{i\neq j}|x_i^k - x_j^k|$ *is a positive nonincreasing function and converges when* $k \to \infty$.

Proof. The result is immediate due to the rules of motion: no point can become more extreme than the already most extreme points. Mathematically, in the inclusion sense, the smallest closed interval containing all points is nonincreasing and will thus converge to a nonempty limit interval. ∎

PROPOSITION 4.2. *For all* $k \in \mathbb{N}$, *we have*

$$\mathbb{E}\left(\sum_{i\neq j}(x_i^{k+1} - x_j^{k+1})^2 \Big| X_k\right) = \left(1 - \frac{2N+1}{3N(N-1)}\right)\sum_{i\neq j}(x_i^k - x_j^k)^2.$$

Proof. We here give an overview of the proof, for which a detailed one can be found in the supplementary material subsection SM1.2.

Denote $\mathcal{L}_{i,j}^k$ and $\mathcal{L}^k$ the studied quantities

$$\mathcal{L}_{i,j}^k = (x_i^k - x_j^k)^2, \tag{4.2}$$

$$\mathcal{L}^k = \sum_{i\neq j}(x_i^k - x_j^k)^2 = \sum_{i\neq j}\mathcal{L}_{i,j}^k. \tag{4.3}$$

We can calculate the expected $\mathcal{L}^{k+1}$ by conditioning on the chosen pair (m,l) for (A^{k+1}, B^{k+1}) and use the linearity of the expectation to focus on the conditional expectation of the (i,j) term $\mathcal{L}_{i,j}^{k+1}$. By expanding the square in this term and once again using the linearity of the expectation, we thus only need to know the first two moments of the updates of the opinions, conditioned on the choice of pair for update (m,l). These are simply given by the well-known first order moments of 1-dimensional uniform random variables (see the supplementary material Proposition SM1.2). Some straightforward calculations give the final result. ∎

Unlike the range $\max_{i\neq j}|x_i^k - x_j^k|$, $\mathcal{L}^k$ is not monotonic. Nevertheless, its usage of the expectation rather than the maximum operator eases its analysis. Both quantities relate to each other through the following proposition.

PROPOSITION 4.3. *For all* $k \in \mathbb{N}$, *we have*

$$N\max_{i\neq j}\left|x_i^k - x_j^k\right|^2 \leq \mathcal{L}^k \leq \frac{N^2}{2}\max_{i\neq j}\left|x_i^k - x_j^k\right|^2.$$

Proof. This is a well-known result from [16] for the upper-bound and [14] for the lower-bound after noticing that, up to normalization and a constant factor, the Lyapunov sum of square differences of N points $x_1, \dots, x_N$ is the (biased) empirical variance of the points. We rederive the proof as supplementary material in subsection SM1.3. ∎

DEFINITION 4.4. *For any* $\varepsilon > 0$, *we denote* T_ε *the stopping time with respect to the natural filtration induced by the* (X_k) *sequence defined as*

$$T_\varepsilon = \inf\left\{k \in \mathbb{N} \mid \max_{i\neq j}\left|x_i^k - x_j^k\right| \leq \varepsilon\right\}.$$

DEFINITION 4.5. *For any $\varepsilon > 0$, we denote T'_ε the stopping time with respect to the natural filtration induced by the (X_k) sequence defined as*

$$T'_\varepsilon = \inf\left\{k \in \mathbb{N} \mid \sum_{i \neq j}(x_i^k - x_j^k)^2 \leq N\varepsilon^2\right\}.$$

PROPOSITION 4.6. *For all $\varepsilon > 0$, $T_\varepsilon \leq T'_\varepsilon$.*

Proof. The result follows directly from Proposition 4.3. Indeed, if $\mathcal{L}^k \leq N\varepsilon^2$, then $\max_{i\neq j}|x_i^k - x_j^k| \leq \frac{\mathcal{L}^k}{N} \leq \varepsilon^2$. Thus, the range of opinions at time k is smaller than or equal to ε, and we have already reached convergence. Therefore, the first occurrence of the event $\{\mathcal{L}^k \leq N\varepsilon^2\}$ happens at the same time or later than the first occurrence of $\{\max_{i\neq j}|x_i^k - x_j^k| \leq \varepsilon\}$, i.e., of convergence. ∎

THEOREM 4.7. *For a system evolving according to* (4.1)*, for any $\varepsilon > 0$, we have*

$$\mathbb{E}(T_\varepsilon \mid X_0) \leq \frac{3N(N-1)}{2N+1}\ln\left(\frac{\mathcal{L}^0}{N\varepsilon^2}\right) + \frac{3N(N-1)}{2N+1} \leq \frac{3}{2}N\ln\left(\frac{\mathcal{L}^0}{N\varepsilon^2}\right) + \frac{3}{2}N.$$

Proof. We here give an overview of the proof, for which a detailed one can be found in the supplementary material subsection SM1.4.

Due to Proposition 4.6, it suffices to find an upper-bound on the expectation of T'_ε. The idea of the proof is to write this expectation as the sum over k of tail distributions: $\mathbb{P}(T'_\varepsilon > k \mid X_0)$. If $\mathcal{L}^k$ is lower than or equal to $N\varepsilon^2$, then by definition T'_ε is lower than or equal to k. By contraposition, it thus suffices to find an upper-bound on the sum over k of tails of a new distribution: $\mathbb{P}(\mathcal{L}^k > N\varepsilon^2 \mid X_0)$. Luckily, we know the expectation of these variables, using Proposition 4.2, by induction on expectations:

$$\mathbb{E}(\mathcal{L}^k \mid X_0) = \left(1 - \frac{2N+1}{3N(N-1)}\right)^k \mathcal{L}^0. \tag{4.4}$$

We can then apply the Markov inequality on each term of the sum to get an upper-bound. However, the Markov inequality tends to be of poor quality in the first terms of the summation, as it there yields huge unrealistic bounds. We can compensate for this by simply upper-bounding the first terms by 1. We find that, for $k \geq \frac{3N(N-1)}{2N+1}\ln(\frac{\mathcal{L}^0}{N\varepsilon^2})$, the Markov inequality provides bounds lower than 1. Thus we split the summation into two parts, the first $\frac{3N(N-1)}{2N+1}\ln(\frac{\mathcal{L}^0}{N\varepsilon^2})$ terms of the sum, which together contribute at most to that amount, and the rest which contributes to at most an infinite geometric series with first term that is quite small. We can show that the second part of the sum can be upper-bounded by $\frac{3N(N-1)}{2N+1}$, which concludes the proof. ∎

We can now simply prove Theorem 3.1 by seeing it as a corollary of Theorem 4.7.

Proof. It follows immediately from Theorem 4.7 since using Proposition 4.3 we have

$$\mathcal{L}^0 \leq \frac{N^2}{2}(b-a)^2. \tag{4.5}$$ ∎

PROPOSITION 4.8. *If I is bounded, say $I = [a, b]$, and if the opinions in X_0 have uniform identical independent distributions in I, then*

$$\mathbb{E}(\mathcal{L}^0) = \frac{N(N-1)}{6}(b-a)^2.$$

Proof. The result is straightforward since it uses well-known first moments of iid uniform random variables (see Proposition SM1.2). ∎

THEOREM 4.9. *If I is bounded, say $I = [a, b]$, and if the initial opinions in X_0 have uniform identical independent distributions in I, then*

$$\mathbb{E}(T_\varepsilon) \leq \frac{3}{2} N \ln\left(\frac{N}{\varepsilon^2}\right) + \frac{3}{2} N \left(\ln\left(\frac{(b-a)^2}{6}\right) + 1\right).$$

Proof. The proof is similar to that of Theorem 4.7. To remove the conditioning on X_0, we use $\mathbb{E}(T'_\varepsilon) = \mathbb{E}\big(\mathbb{E}(T'_\varepsilon \mid X_0)\big)$. We have

$$\begin{aligned}
&\mathbb{E}(T'_\varepsilon) = \mathbb{E}\big(\mathbb{E}(T'_\varepsilon \mid X_0)\big) \\
&\leq \mathbb{E}\left(\sum_{k=0}^{\infty} \mathbb{P}(\mathcal{L}^k > N\varepsilon^2 \mid X_0)\right) = \sum_{k=0}^{\infty} \mathbb{E}\big(\mathbb{P}(\mathcal{L}^k > N\varepsilon^2 \mid X_0)\big) && (4.6)\\
&\leq \sum_{k=0}^{\infty} \mathbb{E}\left(\min\left\{\frac{\mathbb{E}(\mathcal{L}^k \mid X_0)}{N\varepsilon^2}, 1\right\}\right) && (4.7)\\
&\leq \sum_{k=0}^{\infty} \min\left\{\frac{\mathbb{E}(\mathbb{E}(\mathcal{L}^k \mid X_0))}{N\varepsilon^2}, 1\right\} = \sum_{k=0}^{\infty} \min\left\{\left(1 - \frac{2N+1}{3N(N-1)}\right)^k \frac{\mathbb{E}(\mathcal{L}^0)}{N\varepsilon^2}, 1\right\}, && (4.8)
\end{aligned}$$

where the inversion in (4.6) is achieved by positivity of the terms. We can thus replace $\mathcal{L}^0$ by its expectation when removing the conditioning. We then use Proposition 4.8 and continue the proof as in Theorem 4.7 to get the desired result. ∎

THEOREM 4.10. *For a system evolving according to* (4.1), *we have*

$$\left(1 - \frac{2N+1}{3N(N-1)}\right)^k \frac{2\mathcal{L}^0}{N^2} \leq \mathbb{E}\left(\max_{i \neq j} \left|x_i^k - x_j^k\right|^2 \mid X_0\right) \leq \left(1 - \frac{2N+1}{3N(N-1)}\right)^k \frac{\mathcal{L}^0}{N}.$$

Proof. Using Proposition 4.3, if we denote $r_k = \max_{i \neq j} |x_i^k - x_j^k|$ the range of opinions at step k, then

$$\frac{2}{N^2}\mathcal{L}^k \leq r_k^2 \leq \frac{\mathcal{L}^k}{N}. \tag{4.9}$$

We get the final result by taking the expectation and applying Proposition 4.2. ∎

THEOREM 4.11. *If I is bounded, say $I = [a, b]$, and if the opinions in X_0 have uniform identical independent distributions in I, then if r_k is the range at time step k,*

$$\frac{(N-1)\left(1 - \frac{2N+1}{3N(N-1)}\right)^k}{3N}(b-a)^2 \leq \mathbb{E}(r_k^2) \leq \frac{(N-1)\left(1 - \frac{2N+1}{3N(N-1)}\right)^k}{6}(b-a)^2.$$

Proof. The result immediately follows from Theorem 4.10 and Proposition 4.8. ∎

THEOREM 4.12. *A system evolving according to* (4.1) *converges to a single point $x_\infty \in I$ almost surely.*

Proof. The result immediately follows from Theorem 4.7 and Proposition 4.1. Note that the limit point x_∞ is random in I. ∎

Denote for conciseness $\bar{X}_k = \frac{1}{N} \sum_{i=1}^{N} x_i^k \in \mathbb{R}$ the average opinion at step k.

PROPOSITION 4.13. *For all $k \in \mathbb{N}$, we have $\mathbb{E}(\bar{X}_k \mid X_0) = \bar{X}_0$.*

Proof. The proof is straightforward and uses the well-known first moment of uniform random variables (see supplementary material Proposition SM1.2). See supplementary material subsection SM1.5 for a detailed proof. Note that this result is intuitively clear as the update rule is average preserving in mean. ∎

Denote $1_N = (1, \ldots, 1)^T$ the vector of ones and $(\beta_2, \ldots, \beta_N) \in \mathbb{R}^{N-1}$ an arbitrary orthonormal basis of the space orthogonal to the 1-dimensional space generated by 1_N. Define columnwise the following unitary matrix $U = \begin{pmatrix} \frac{1}{\sqrt{N}} 1_N & \beta_2 & \cdots & \beta_N \end{pmatrix}$. Let $\mathrm{Diag}(\lambda_1, \ldots, \lambda_N)$ be the diagonal matrix with entries $\lambda_1, \ldots, \lambda_N$.

PROPOSITION 4.14. *For all $k \in \mathbb{N}$, we have*

$$\mathbb{E}(X_k \mid X_0) = U \,\mathrm{Diag}\left(1, \left(1 - \frac{1}{N-1}\right)^k, \ldots, \left(1 - \frac{1}{N-1}\right)^k\right) U^T X_0.$$

Proof. The proof is also straightforward by working conditionally on the choice of pair (i, j) for update. We find that

$$\mathbb{E}(X_{k+1} \mid X_k) = U \,\mathrm{Diag}\left(1, \left(1 - \frac{1}{N-1}\right), \ldots, \left(1 - \frac{1}{N-1}\right)\right) U^T X_k, \tag{4.10}$$

which gives the final result by induction. A detailed proof is given in the supplementary material in subsection SM1.6. ∎

THEOREM 4.15. *The limit point for a system evolving according to* (4.1) *has the expectation:*

$$\mathbb{E}(x_\infty \mid X_0) = \bar{X}_0 = \frac{1}{N} \sum_{i=1}^{N} x_i^0.$$

Proof. Using Theorem 4.12, all opinions converge almost surely to the same finite but random value, and since the opinions are all bounded by the initial interval $I_0 = [\min X_0, \max X_0]$, we have by bounded convergence that

$$\begin{cases} \lim\limits_{k \to \infty} \mathbb{E}(x_i^k \mid X_0) = \mathbb{E}(x_\infty \mid X_0), \\ \lim\limits_{k \to \infty} \mathbb{E}(\bar{X}_k \mid X_0) = \mathbb{E}(x_\infty \mid X_0). \end{cases} \tag{4.11}$$

We then conclude using (4.11) and either Propositions 4.13 and 4.14. ∎

Note the following important remark: while Proposition 4.14 gives that the convergence of the expectations of each opinion is exponentially fast, it does not provide any guarantee on the speed of the convergence of sequences $(X_0, X_1, X_2, \ldots)$. This is given by Theorem 4.7.

We can compare our result with those from the gossip literature. Note that, in gossiping, x_∞ is deterministic and equals $\frac{1}{N} \sum_{i=1}^{N} x_i^0$ almost surely and that the communication graph may be assumed different from the complete graph, leading to different convergence times.

DEFINITION 4.16. *For any $\varepsilon > 0$, we denote $T_{gossip}(\varepsilon)$ the "ε-averaging time," which is the deterministic quantity studied in the gossip algorithms' literature and used there as the convergence time, defined as*

$$T_{gossip}(\varepsilon) = \sup_{X_0 \in I^N} \inf_{k \in \mathbb{N}} \left\{ k \mid \mathbb{P}\left(\frac{\|X_k - x_\infty (1, \ldots, 1)^T\|_2}{\|X_0\|_2} \geq \varepsilon \,\middle|\, X_0 \right) \leq \varepsilon \right\}.$$

THEOREM 4.17. *Assume I is bounded, say $I = [a, b]$, then*

$$T_{gossip}(\varepsilon) \leq \frac{-3\ln\varepsilon + \ln\left(2(N-1)(1-q_{a,b})\right)}{-\ln\left(1 - \frac{2N+1}{3N(N-1)}\right)} \leq \frac{3}{2}N\ln\left(\frac{N}{\varepsilon^3}\right) + \frac{3}{2}N\ln\left(2\left(1-q_{a,b}\right)\right),$$

where $q_{a,b} = \frac{a^2}{b^2}\mathbb{1}_{a\geq 0} + \frac{b^2}{a^2}\mathbb{1}_{b\leq 0}$.

Proof. We here give an overview of the proof, for which a detailed one can be found in the supplementary material subsection SM1.7.

Using Proposition 4.3, we can link the squared range to the Lyapunov quantity. Thus the tail distribution in the definition of $T_{gossip}(\varepsilon)$ can be upper-bounded by a tail distribution of $\mathcal{L}^k$. We can then apply the Markov inequality to get an upper-bound of this tail distribution. Unlike us, the gossip literature is solely interested in one tail distribution: the one that passes the ε-threshold. It thus suffices to find for which k the bound given by the Markov inequality is smaller than ε. ∎

Our result is similar to those in the gossip literature. However, the quantity we study is slightly different, and we need to add the result of the Markov inequality at all levels and not just at a specific level depending on ε. However, if unconventionally we change the definition of $T_{gossip}(\varepsilon)$ by using the squared 2-norms instead of the plain 2-norms, then the same calculations would give an upper-bound with dominant term $\frac{3}{2}N\ln(\frac{N}{\varepsilon^2})$. This suggests that the unconventional gossip convergence time is similar to our expected convergence time. This is a general result that is due to the fact that the quantity of interest, for us and for gossiping, is $\mathcal{L}^n$ and that it is an exponentially decreasing positive supermartingale. We can prove that if (Y_k) is a positive exponentially decreasing supermartingale, i.e., there is $\alpha > 0$ such that $\mathbb{E}(Y_{k+1} \mid Y_k) = (1-\alpha)Y_k$, and if Y_k converges to 0 almost surely, then if we slightly change the convergence definition of T_ε to $T_{cv}(\varepsilon^2)$ by normalizing by Y_0, i.e., looking at the first occurrence of $\{\frac{Y_k}{Y_0} \leq \varepsilon^2\}$, then $\mathbb{E}(T_{cv}(\varepsilon^2) \mid Y_0) \leq \frac{-2\ln\varepsilon}{-\ln(1-\alpha)} + \frac{1}{\alpha}$ and $T_{gossip}(\varepsilon) \leq \frac{-2\ln\varepsilon}{-\ln(1-\alpha)}$. Thus both quantities of interested have similar upper-bounds. In our case, Y_k would be $\mathcal{L}^k$, up to multiplicative factors, different in each case, in order to get a comparison with the same threshold ε. A proof of this similarity between expected convergence time and gossiping is given as supplementary material in subsection SM1.8.

5. Unconstrained D-dimensional case. We can generalize the previous 1-dimensional model to a similar D-dimensional one. Now, we assume that people's opinions depend on several parameters. However, we will say that the space of opinions is "unconstrained" in the sense that the domain for opinions is a convex set of $\mathbb{R}^D$. This model well suits cases when extreme opinions correspond to at least one "extreme" parameter, i.e., very big or very small, and any nonextreme point lying inside the convex hull of the extreme opinions is a feasible opinion. Once again, the space of opinions can either be bounded or unbounded.

Mathematically, we define the space of opinions to be a convex set $C \subset \mathbb{R}^D$. Our system comprises $N \geq 2$ indistinguishable agents, each with their own opinion: $x_i = (x_{i,1}, \ldots, x_{i,D})^T \in C$ for $i \in \{1, \ldots, N\}$. The initial opinions are random, but conditionally on them they form a given set $x_i^0 \in C$. Opinion dynamics are modeled in discrete time, conditionally on the state of opinions at the previous time step $X_k = (x_1^k, \ldots, x_N^k)^T \in \mathbb{R}^{N\times D}$. The evolution law at time step $k+1$ for all $i \in \{1, \ldots, N\}$ is

$$x_i^{k+1} = \mathbb{1}_{i\notin\{A_{k+1},B_{k+1}\}}x_i^k + \mathbb{1}_{i=A_{k+1}}U_1^{k+1} + \mathbb{1}_{i=B_{k+1}}U_2^{k+1}, \tag{5.1}$$

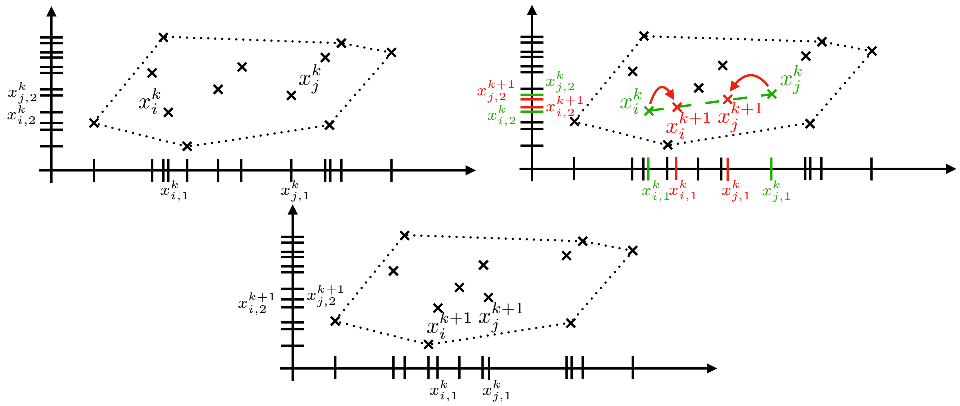

FIG. 5.1. *One-step opinion evolution in the unconstrained* 2*-dimensional model.*

where (A_{k+1}, B_{k+1}) is a random uniform sampling of two indices of $\{1, \ldots, N\}$ without replacement independent of the past and where, conditionally on X_k, A_{k+1}, and B_{k+1}, U_1^{k+1} and U_2^{k+1} are independent random uniform variables in the line segment $\{(1-\lambda)x_{A_{k+1}}^k + \lambda x_{B_{k+1}}^k \mid \lambda \in [0,1]\}$. Concretely, at each time step, two random agents A_{k+1} and B_{k+1} are selected, and they then independently and uniformly resample their opinion in the segment between both previous opinions. See Figure 5.1 for an example when $D = 2$.

We define the ε-convergence time T_ε as the first time when all agents are at most ε distant from each other. Its definition is the same as for the 1-dimensional case in Definition 4.4 except that $|\cdot|$ has been replaced by $\|\cdot\|_2$.

It is fairly simple to see that, along each dimension, the projected coordinates of the opinions rigorously obey (albeit simultaneously) the 1-dimensional rules of motion (4.1). As such, we can easily generalize the proofs done in the 1-dimensional case to the general D-dimensional unconstrained convex case by working per coordinates. In particular, the maximum L^2 distance between two opinions is nonincreasing, the opinions will almost surely converge to a random limit point x_∞ that preserves the initial average on mean, i.e., $\mathbb{E}(x_\infty \mid X_0) = \bar{X}_0$, and the expectation of the time to reach ε-convergence is also finite, and we give an upper-bound for $\mathbb{E}(T_\varepsilon)$ in Theorem 3.2. We provide detailed proofs on how to generalize the 1-dimensional results to the D-dimensional case as supplementary material in section SM2. The key message is that, thanks to the convexity of the opinion space, increasing the dimensionality does not fundamentally change the behavior of the system nor the tools for analyzing it.

6. Constrained 2-dimensional case. The limitation of the previous model is the convex opinion space assumption, which is well adapted to situations where "extreme" opinions correspond to at least one "big" opinion parameter. However, in some cases, it is more accurate to also consider some "extreme" cases with neither parameter being "big." This happens when the opinion space is no longer convex.

For nonconvex opinion spaces, it is then necessary to redefine how agents interact. In the convex case, we modeled an interaction along the line segment linking the two states. In a nonconvex, yet arc-connected space, a reasonable possibility to model interactions between opinions is to consider a geodesic between the opinions in the opinion space. In this paper we will study a simple case that naturally generalizes the previous models: the unit circle, which is interesting from two aspects. First it is nonconvex in $\mathbb{R}^2$. Secondly, a reparametrization of S using the oriented angle from the x-axis in $[0, 2\pi)$ leads to a new parameter space $[0, 2\pi)$ for S which is convex. However, it fundamentally differs from the previous convex models for the two following reasons: first, as opinions communicate along geodesics of S, if $(\theta_1, \theta_2) \in [0, 2\pi)^2$, then, depending on the size of $|\theta_2 - \theta_1|$, communication can happen in the $[\min\{\theta_1, \theta_2\}, \max\{\theta_1, \theta_2\}]$ interval or in its closed complement in $[0, 2\pi)$ which is equal to $[\max\{\theta_1, \theta_2\}, 2\pi) \cup [0, \min\{\theta_1, \theta_2\}]$. Thus convexity in the parameter space is not enough to use the previous models, as we require convexity in the embedding space. Secondly, if $\theta \in [0, 2\pi)$ increases, then as θ reaches the right boundary of the interval, then it simultaneously reaches the left boundary as well. This violates the principle that we assumed in the previous cases where a more and more "right-wing" opinion could not simultaneously become more and more "left-wing." Furthermore, the circle naturally appears as the opinion space when working on alignment of 2-dimensional agents moving with unit velocities in the plane. Indeed, circular opinions correspond to headings, and alignment only occurs when all headings are the same, i.e., in consensus on the circle.

Mathematically, we define the space of opinions to be $S \in \mathbb{R}^2$ the unit circle embedded in 2-dimensional space. Our system comprises $N \geq 2$ indistinguishable agents, each with their own opinion: $x_i = (x_{i,1}, x_{i,2})^T$ following the circle constraint $x_{i,1}^2 + x_{i,2}^2 = 1$. The opinion state is $X_k = (x_1^k, \dots, x_N^k)^T \in \mathbb{R}^{N\times 2}$. It will be useful to consider the equivalent reparametrization by angles $\theta_i \in [0, 2\pi)$ with $(x_{i,1}, x_{i,2}) = (\cos(\theta_i), \sin(\theta_i))$. The initial opinions are random, but conditionally on them they form a given set $\theta_i^0 \in [0, 2\pi)$. Here the state-evolution dynamics are modeled in discrete time, conditionally on the state of opinions at the previous time step $\Theta_k = (\theta_1^k, \dots, \theta_N^k)^T \in \mathbb{R}^N$. The evolution law at time step $k+1$ for all $i \in \{1, \dots, N\}$ is now done along geodesics and is

$$\theta_i^{k+1} = \mathbb{1}_{i \notin \{A_{k+1}, B_{k+1}\}} \theta_i^k + \mathbb{1}_{i = A_{k+1}} U_1^{k+1} + \mathbb{1}_{i = B_{k+1}} U_2^{k+1}, \tag{6.1}$$

where (A_{k+1}, B_{k+1}) is a random uniform sampling of two indices of $\{1, \dots, N\}$ without replacement independent of the past and where, conditionally on Θ_k, A_{k+1}, and B_{k+1}, U_1^{k+1} and U_2^{k+1} are independent random uniform variables in $G(\theta_{A_{k+1}}, \theta_{B_{k+1}}) \subset [0, 2\pi)$, which is the geodesic circle arc between opinions $x_{A_{k+1}}$ and $x_{B_{k+1}}$:

$$G(\theta, \tilde{\theta}) = \begin{cases} \left[\min\{\theta, \tilde{\theta}\}, \max\{\theta, \tilde{\theta}\}\right] & \text{if } \left|\tilde{\theta} - \theta\right| \leq \pi, \\ \left[\max\{\theta, \tilde{\theta}\}, 2\pi\right) \cup \left[0, \min\{\theta, \tilde{\theta}\}\right] & \text{if } \left|\tilde{\theta} - \theta\right| > \pi. \end{cases} \tag{6.2}$$

Concretely, at each time step, two random agents A_{k+1} and B_{k+1} are selected, and they then independently and uniformly resample their opinion on the shortest circle arc between both previous opinions. See Figure 6.1 for an example.

Note that, for the pathological case of two agents at an angular distance of exactly π, then we chose a deterministic geodesic. This work would be similar if we chose a random geodesic in that case and even if we chose for the two agents to not necessarily choose the same one. This is because this π distance configuration almost surely never happens, as the updates are continuous random variables, except for the eventual cases in Θ_0 were angles are initially set to be at such a distance.

Note that, as in the convex case, we allow the updated opinions to change their order on the circle. By simple renaming of the unit vectors,[1] we could preserve

[1]Note that we do assume agents to be identical and indistinguishable anyway.

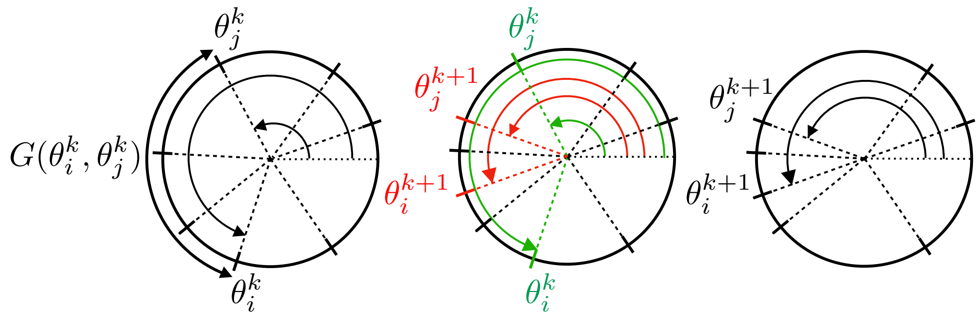

FIG. 6.1. *One-step opinion evolution in the constrained 2-dimensional model.*

"ordering," but nothing would change in our analysis of the problem. However, when viewing the opinions as headings for agents that move with unit velocities in the plane, "agents selected for interaction" could correspond to physical encounters of two agents in the plane. In such cases, depending on physical assumptions on the type of their interactions, one might wish to either keep or exchange their ordering on the circle. As mentioned above, we account for both possibilities by not constraining the ordering of the updated agents.

We are not the first to propose to use the circle as nonconvex opinion space and its natural angular parametrization. Indeed, the circle is the most natural and simplest nonconvex connected-by-arc opinion space. Notably, it was also used in [10]. Nevertheless, the assumed rules of interaction in [10] radically differ from ours. First, the interaction graph corresponds to a predetermined ring in the finite case (and $\mathbb{Z}$ in the infinite one), rather than to the complete graph like in our case. Secondly, the dynamics are not doubly stochastic since the randomness lies only in the selected pair of agents: once selected, the agents will update their opinions in a deterministic way as is standard in the Deffuant model. We further incorporate randomness in the opinion update along the geodesic between opinions. While we proposed to work in the same domain, our work is different in both the assumed interaction graph and the update rules as well as the tools for the analysis of convergence.

Other less relevant approaches exist for opinion dymanics on nonconvex manifolds. Motivated by alignment of agents in the plane by considering the heading of planar agents as their opinions, [19] uses a projected gradient approach, where agents simply perform a given step along the gradient of a potential function involving the other opinions (and the weighted interaction graph). However, there is no randomness in the update whatsoever. The approach is purely deterministic, and such papers only prove convergence, but they do not study the expected time to reach it.

We now focus on our model. Similarly to the convex case, we define the following stopping time on the angle parametrization.

DEFINITION 6.1. *For any $\varepsilon > 0$, we denote T_ε the stopping time, with respect to the natural filtration induced by the (Θ_k) sequence, defined as the first time step when all unit vector opinions are within a circle arc of length ε:*

$$T_\varepsilon = \inf\left\{k \in \mathbb{N} \mid \exists i, j \text{ s.t. } \left|G(\theta_i^k, \theta_j^k)\right| \leq \varepsilon \text{ and for all } l, \theta_l^k \in G(\theta_i^k, \theta_j^k)\right\},$$

where $|G(\alpha, \beta)| = \min\{|\alpha - \beta|, 2\pi - |\alpha - \beta|\}$ is the geodesic angular distance between angles α and β.

DEFINITION 6.2. *We denote T_{HD} the stopping time, with respect to the natural filtration induced by the (Θ_k) sequence, of the event that all unit vector opinions are within a half-disk:*

$$T_{HD} = \min\{k \in \mathbb{N} \mid \exists \theta_{HD}^k \in [0, 2\pi) \text{ for all } i \in \{1, \ldots, N\}, \cos(\theta_i^k - \theta_{HD}^k) > 0\}.$$

PROPOSITION 6.3. *For any system evolving according to* (6.1), *if at time step $k \in \mathbb{N}$ all unit vector opinions are within a half-disk, then for all $k' \geq k$, all unit vector opinions are within a half-disk.*

Proof. This result is obvious after realizing that the rules of motion in (6.1) are equivalent to those in the 1-dimensional case (4.1) once all opinions are within a half-disk. To better see this, we present a detailed proof.

Let θ_{HD}^k be an angle such that, for all i, $\cos(\theta_i^k - \theta_{HD}^k) > 0$. Perform the change of angular parametrization around θ_{HD}^k for all future time steps $k' \geq k$:

$$\tilde{\theta}_i^{k'} = \theta_i^{k'} - \theta_{HD}^k. \tag{6.3}$$

In this new parametrization in $[-\pi, \pi)$, all geodesics are contained within the one between the two most extreme opinions, i.e., $G(\tilde{\theta}_i^k, \tilde{\theta}_j^k) \subset G\big(\min\{\tilde{\theta}_l^k\}, \max\{\tilde{\theta}_l^k\}\big)$ for all $(i, j) \in \{1, \ldots, N\}^2$. Due to the rules of motion (6.1), this in turn implies that all angles at the next step are within that same geodesic. By induction, we can claim that, for all future time steps $k' \geq k$, $\tilde{\theta}_i^{k'} \in G(\min\{\tilde{\theta}_l^k\}, \max\{\tilde{\theta}_l^k\})$, which implies that all angles are contained within a half-disk forever. ∎

We first prove Lemma 3.3.

Proof. The proof consists in showing the existence of a sequence of finitely many updates with probability lower-bounded by a strictly positive constant that drives the system from any configuration that is not contained in a half-disk to a configuration contained in one. To do this, it suffices to find a finite sequence of events leading to the half-disk configuration from any other configuration. Due to the finite number of agents, each of these events will have lower-bounded probabilities, and since the sequence is finite, we have a nonzero lower-bound for the probability to have such a sequence occur from any given configuration. In turn this gives that almost surely all agents will be located within a half-disk and that the expected time for this to occur is finite. A detailed proof is given as supplementary material in subsection SM3.1. ∎

As we later show in numerical experiments, the bound in Lemma 3.3 is rather loose due to the proof strategy that worked. We hope that better bounds can be derived. To also take them into account, we define the set of finite upper-bounds of the expected time of convergence, and the bound we proved is just one of its elements.

DEFINITION 6.4. *Let $\mathcal{B}_N^{HD}$ be the set of finite upper-bounds of $\mathbb{E}(T_{HD} \mid \Theta_0)$:*

$$\mathcal{B}_N^{HD} = \left\{B_N^{HD} \in \mathbb{R} \mid \mathbb{E}(T_{HD}) \leq B_N^{HD}\right\}.$$

We can now prove Theorem 3.4.

Proof. The result immediately follows from Lemma 3.3 and Theorem 4.7. Indeed, it suffices to notice that once all agents are within a half-disk, then the dynamics of the system using (6.1) are equivalent to the 1-dimensional dynamics (4.1) using the angles for the opinions. Note that the "initial" 1-dimensional Lyapunov once we have reached the state where all unit vector opinions are within a half-disk, that is, the Lyapunov at time step T_{HD}, is a random value: however, since the pairwise angular distance is then less than π for any pair of opinions, we have bounded it using Proposition 4.3 by

$$\mathcal{L}^{T_{HD}} \leq \frac{N^2}{2}\pi^2. \tag{6.4}$$ □

7. Open problems on the constrained 2-dimensional case. Many issues remain unsolved for the constrained 2-dimensional case. We propose them in this paper as open questions. The main problem was to obtain a better bound than the crude $O((\frac{9\sqrt{2}}{2})^N N^{\frac{3}{2}N})$ provided in Lemma 3.3 for the expectation of the time for all agents to get within a half-disk. We provide three interesting approaches based on different quantities for which we do not have a final solution. Details can be found in our technical report [6].

The first approach consists in studying the vector sum of all the unit vector opinions $\mathcal{S}^k = \sum_{i=1}^N x_i^k$. The purpose of studying this vector is that convergence of opinions in S is equivalent to convergence of $\mathcal{S}^k$ in $\mathbb{R}^2$ and to convergence of its 2-norm to its upper-bound, N, by finiteness of the problem. Intuitively and experimentally, if $\|\mathcal{S}^k\|_2^2$ is "large," then there is a "large" number of opinions positively oriented with $\mathcal{S}^k$, and furthermore opinions positively oriented with $\mathcal{S}^k$ tend to be updated in a way that further increases the norm of $\mathcal{S}^k$. However, $\|\mathcal{S}^k\|_2$ is upper-bounded by N which can only happen for opinions arbitrarily close to each other. Therefore, we can simply study the evolution of $\|\mathcal{S}^k\|_2^2$, which is an upper-bounded random real quantity, show that it converges to its upper-bound, and study its speed of convergence. Another possibility would be to analyze $\langle \mathcal{S}_{k+1}, \mathcal{S}_k \rangle$ in order to take into account reinforcement drift in the direction of S_k when its norm is sufficiently large. We propose to introduce the geodesic bisectors $\beta_{i,j}^k$ between each pair of agents $\{i, j\}$ and the half angle $\alpha_{i,j}^k$ of the geodesic circle arc between them. Many interesting properties and formulae can be derived; unfortunately we are faced with summations of quantities that are difficult to bound.

The second approach consists in analyzing the evolution of the maximal empty angle $\gamma_{\max}^k$, which is the angle of the longest circle arc between two consecutive opinions on the circle. Note that this arc is not necessarily geodesic. Obviously, there is equivalence between $\gamma_{\max}^k > \pi$, and all unit vector opinions are within a half-disk. Thus we could study $\gamma_{\max}^k$ as a random walk on $[0, 2\pi)$ starting in $[0, \pi)$ and look for the first time it passes the π threshold. Ideally, $\gamma_{\max}^k$ would be a submartingale which would then give us almost sure convergence and convergence time bounds. We can show that, while the opinions are not yet contained within a half-disk, $\gamma_{\max}^k$ is biased to increase and in particular that $\mathbb{P}(\gamma_{\max}^{k+1} < \gamma_{\max}^k \mid \gamma_{\max}^k) \leq \frac{1}{2}(1 - \frac{1}{N})$. Unfortunately, simply having that the probability of decrease is upper-bounded by a value strictly smaller than $\frac{1}{2}$ is not enough; we need to study with more detail the probability distribution of $\gamma_{\max}^k$ for its expectation. The proof provides a reasonable approach to bound the entire distribution of the decrease event of the maximal empty angle. However, analyzing the increase is significantly harder and remains an open challenge.

The third approach consists in designing and analyzing Markov chains. We studied a Markov chain with $n + 1 = \lfloor \frac{N}{2} \rfloor + 1$ states, which is an extension of the naive proof in Lemma 3.3. It is essentially a doubly chained graph with probability c of increase and $1 - c$ of decrease, and the last state is absorbing. On average, reaching the absorbing state takes longer than reaching a half-disk configuration. In Lemma 3.3, we analyze n successive increases. In reality, we tolerate some decreases in the process. Explicit calculation of the expected time to reach the absorbing state is possible by inverting an almost tridiagonal Toeplitz matrix using the Sherman–Morrison formula and the well-known invert of a tridiagonal Toeplitz matrix [5]. As $c = \frac{4}{27N^2(N-1)^2} < \frac{1}{2}$, the expected time is approximately $(\frac{1-c}{c})^n$. This yields a bound similar to the one given in Lemma 3.3. The problem is that c was derived using a pessimistic worst case geometry per state. In practice, closer to half-disk configurations, thus with higher state number, the geometry is biased far away from the worst case scenario giving on average significantly higher state increase probabilities. We believe that it should be possible to find an alternative simple Markov chain with higher probabilities for getting to the absorbing state that provides a reasonable upper-bound .

8. Numerical results.

While the theory provides a guarantee of finite expected time convergence in all previous cases, it also provides explicit bounds, which we can compare to empirical results in numerical simulations.

8.1. 1-dimensional case.

The chosen domain is the unit interval $I = [0, 1]$. The initial opinions in X_0 follow an iid uniform distribution on I. We tested the grid of configurations defined by the number of agents $N \in \{5, 10, 100, 250, 500, 750, 1000\}$ and convergence threshold $\varepsilon \in \{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1\}$. For each configuration, $n_{trials} = 1000$ independent trials were performed. Each trial ran until we reached ε-convergence. We denote $\hat{T}_\varepsilon$ the natural estimator of $\mathbb{E}(T_\varepsilon)$ by simply taking its empirical average.

A summary of the empirical dependency of the average convergence time on the convergence threshold ε is done in Figure 8.1, where we plot $\hat{T}_\varepsilon$ against ε and against $-\ln \varepsilon$. We find that $\hat{T}_\varepsilon$ has a minus logarithmic dependency on ε as expected from Theorem 4.9. Furthermore, the slopes of the curves and their respective bounds from Theorem 4.9 in Figure 8.1(b) seem to be approximately the same for high N, suggesting that in fact the convergence time is not only upper-bounded but also lower-bounded by a similar term with approximately the same dominant coefficient: $T_\varepsilon \approx c_N \frac{3}{2} N \ln(\frac{N}{\varepsilon^2}) + O(1)$, where $O(1)$ represents a function bounded with respect to ε (but not with respect to N) and $c_N \in [0, 1]$ represents a constant depending on N such that

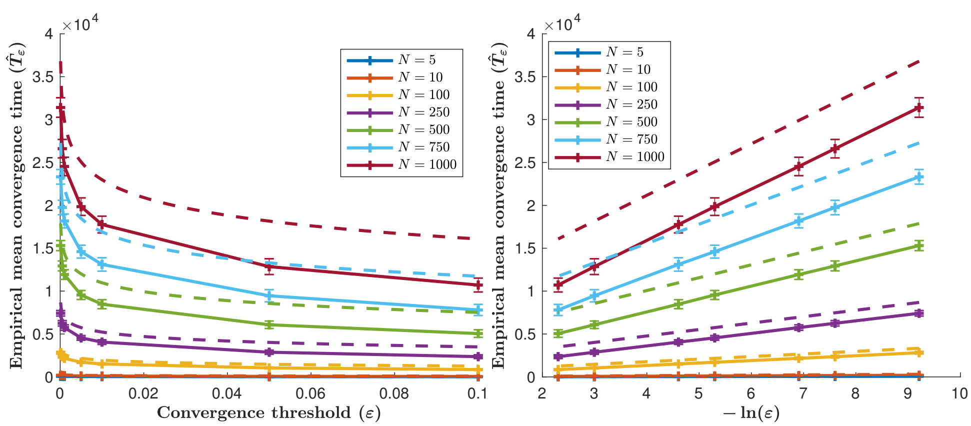

FIG. 8.1. 1-*dimensional evolution: Dependency of the empirical mean convergence time on the convergence threshold* ε. *Left:* ε *abscissa. Right:* $-\ln \varepsilon$ *abscissa. The plain curves correspond to the empirical results, whereas the dashed ones correspond to the theoretical bounds. We superimpose on the empirical curves the classic unbiased estimator of the standard deviation of each data point.*

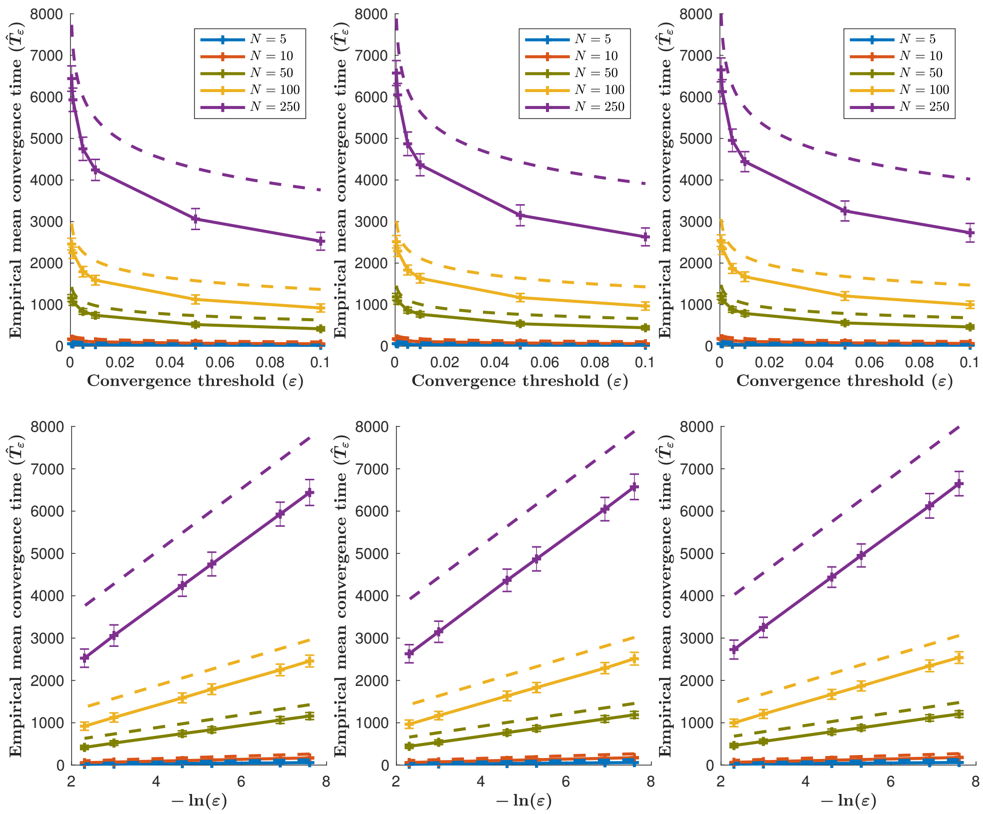

FIG. 8.2. *D-dimensional evolution: Dependency of the empirical mean convergence time on the convergence threshold* ε *in the* 2-, 3-, *and* 4-*dimensional cases from left to right. Top:* ε *abscissa. Bottom:* $-\ln\varepsilon$ *abscissa. The plain curves correspond to the empirical results, whereas the dashed ones correspond to the theoretical bounds. We superimpose on the empirical curves the traditional unbiased estimator of the standard deviation of each data point.*

$c_N \xrightarrow[N\to\infty]{} 1$ and $c_N \approx 1$ when $N \geq 100$. A further analysis of the dependency on N of $\hat{T}_\varepsilon$ gives that empirically $\hat{T}_\varepsilon \approx -3c_N N \ln\varepsilon + 0.89N\ln N - 2.3N + 5.8$. This is done in the supplementary material subsection SM4.1, as we are primarily interested in the ε-dependency in this paper.

8.2. Unconstrained *D*-dimensional case. The chosen domain is the unit cube $C = [0,1]^D$, where the dimension D ranges in $\{2,3,4\}$. The initial opinions in X_0 follow an iid uniform distribution on C. We tested the grid of configurations defined by the number of agents $N \in \{5, 10, 50, 100, 250\}$ and convergence threshold $\varepsilon \in \{0.0005, 0.001, 0.005, 0.01, 0.05, 0.1\}$. For each configuration, $n_{trials} = 1000$ independent trials were performed. Each trial ran until we reached ε convergence. A similar estimator was used as in the 1-dimensional case.

A summary of the empirical dependency of the average convergence time on the convergence threshold ε is done in Figure 8.2. Once again, we find that $\hat{T}_\varepsilon$ has a minus logarithmic dependency on ε and the slopes approximately correspond to those derived in the upper-bound. On the other hand, for the tested values of D, the displacement between the true convergence time and the bounds seems to be the same. Furthermore, we see a slight increase in convergence time with respect to D. However, it would require extensive trials with high D to be able to claim that the dependency is indeed logarithmic, which would be computationally too expensive for our purposes. These three observations lead us to generalize naturally the conjecture made in the 1-dimensional case: $T_\varepsilon \approx c_N \frac{3}{2} N \ln(\frac{DN}{\varepsilon^2}) + O(1)$, where $O(1)$ represents a function bounded with respect to ε (and perhaps also with respect to D but not with respect to N). See the supplementary material subsection SM4.2 for an analysis of the dependency on N and in particular for the confirmation of the presence of c_N.

8.3. Constrained 2-dimensional case. The initial opinions in Θ_0 follow an iid uniform distribution on $[0, 2\pi)$. We tested the grid of configurations defined by the number of agents $N \in \{5, 10, 100, 250, 500, 750, 1000\}$ and convergence threshold $\varepsilon \in \{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1\}$. For each configuration, $n_{trials} = 1000$ independent trials were performed. Each trial was stopped as soon as ε-convergence was reached, i.e., $\gamma^k_{\max} \geq 2\pi - \varepsilon$. A similar estimator was used as in the convex case.

A summary of the empirical dependency of the average convergence time on the convergence threshold ε is done in Figures 8.3(a) and 8.3(b). As in the convex case, we find that $\hat{T}_\varepsilon$ has a minus logarithmic dependency on ε as predicted in Theorem 3.4 with similar slope. However, the bound is many orders of magnitude larger than our estimator even for large ε. This is due to our poor bound B_N^{HD} deriving from a Borel–Cantelli-like idea when studying T_{HD}. Since T_{HD} does not depend on ε as soon as $\varepsilon \leq \pi$, the dependency on ε is naturally inherited from the 1-dimensional case, as the angle of the opinions follows the $1D$ case evolution when all unit vector opinions are within a half-disk. We can therefore extend the conjecture to the circle case: $T_\varepsilon \approx c_N \frac{3}{2} N \ln(\frac{N}{\varepsilon^2}) + O(1)$, where $O(1)$ represents a function bounded with respect to ε (but not with respect to N).

The key part in the circle evolution, and the hardest one to analyze, is the initial regime when not all agents are within a half-disk, i.e., $k < T_{HD}$. To better understand the behavior of the systems in this regime, a summary of the empirical dependency of the average stopping time to a half-disk configuration on the number of agents N is done in Figures 8.3(c) and 8.3(d). We find that $\hat{T}_{HD}$ depends quasi-linearly on N (in

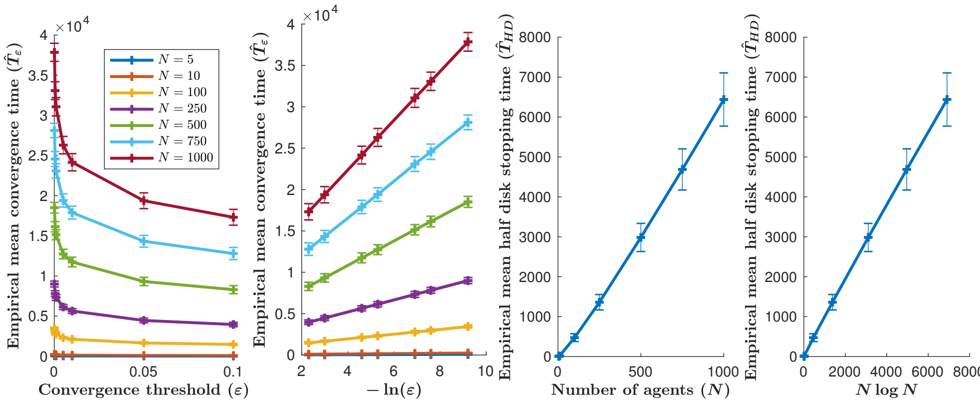

(a) (b) (c) (d)

FIG. 8.3. *Circle evolution: The two plots on the left represent the dependency of the empirical mean convergence time on the convergence threshold ε, while the two on the right display the dependency of the empirical mean half-disk stopping time $\hat{T}_{HD}$ on the number of agents N.* (a) *Empirical results with abscissa ε.* (b) *Empirical results with abscissa $-\ln \varepsilon$.* (c) *Empirical results with abscissa N.* (d) *Empirical results with abscissa $N \ln N$. We superimpose the traditional unbiased estimator of the standard deviation of each data point.*

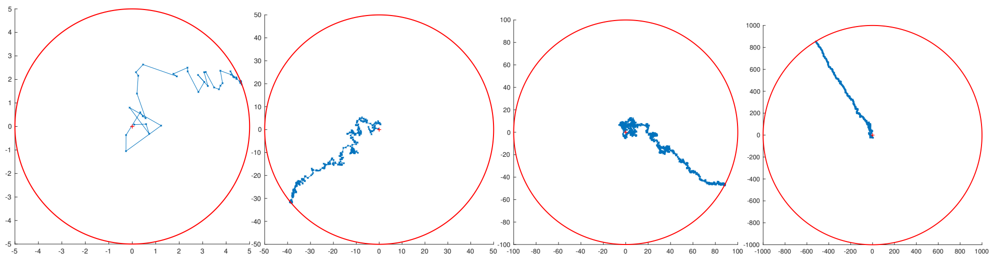

(a) $N = 5$ (b) $N = 50$ (c) $N = 100$ (d) $N = 1000$

FIG. 8.4. *Circle evolution: Evolution of the sum of unit vector opinions $\mathcal{S}^k$ in a single trial for various number of agents with random uniform initial distribution on the circle. The red circle corresponds to $\|\mathcal{S}^k\|_2 = N$. We stopped the evolution after* 1000 *steps for $N \in \{5, 50\}$, after* 2000 *steps for $N = 100$, and after* 30000 *steps for $N = 1000$.*

fact a linear regression gives that $\hat{T}_{HD} \approx 0.92N \ln N + 100$) to be compared with the $O((\frac{9\sqrt{2}}{2})^N N^{\frac{3}{2}N})$ bound from Lemma 3.3, which is many orders of magnitude larger than our estimator even for the smallest number of agents. Further work is needed to find a better theoretical $B_N^{HD} \in \mathcal{B}_N^{HD}$ that should be a $O(N \log N)$.

A further analysis of the dependency on N of $\hat{T}_\varepsilon$ gives that empirically we have $\hat{T}_\varepsilon \approx -3c_N N \ln \varepsilon + 0.93N \ln N + 4.1N - 18$. This is done in the supplementary material subsection SM4.3, as we are primarily interested in the ε-dependency in this paper.

We also plot examples of evolutions of $\mathcal{S}^k$, the vector sum of all unit vectors, in single trials for various number of agents in Figure 8.4. It seems that $\mathcal{S}^k$ is initially random around 0 and then after a small threshold distance drifts in its current direction, suggesting that $\|\mathcal{S}^k\|_2^2$ or $\langle \mathcal{S}^{k+1}, \mathcal{S}^k \rangle$ would be interesting quantities to analyze.

9. Conclusion. We analyzed in detail models of doubly stochastic pairwise interactions for N agents with states described by a single real value, by a D-dimensional real vector, or by a constrained unit vector on the circle. The evolution in time of the states of the N-agent system was found to exhibit convergence to ε-agreement in finite expected time, and we provide upper-bounds on the expected time that seem to be tight in the case of unconstrained states. However, for unit vector states the dependence on N in the upper-bound is quite far from the empirical results. This is due to the difficulty in proving a fast gathering of unit vectors into a half-circle as a result of the assumed doubly stochastic pariwise interactions. This problem is challenging, but we hope to address it in the near future, along the lines outlined in section 7.

Acknowledgments. We thank our referees for their thorough reading of our manuscript and for their useful comments that contributed to the improved presentation of our results.

REFERENCES

[1] D. Aldous and D. Lanoue, *A lecture on the averaging process*, Probab. Surv., 9 (2012), pp. 90–102.

[2] E. Bertin, M. Droz, and G. Grégoire, *Boltzmann and hydrodynamic description for self-propelled particles*, Phys. Rev. E, 74 (2006), 022101.

[3] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, *Randomized gossip algorithms*, IEEE Trans. Inform. Theory, 52 (2006), pp. 2508–2530.

[4] S. Chatterjee and P. Diaconis, *Note on Repeated Random Averages*, preprint, arXiv:1911.02756, 2019.

[5] C. da Fonseca and J. Petronilho, *Explicit inverses of some tridiagonal matrices*, Linear Algebra Appl., 325 (2001), pp. 7–21.

[6] T. Dagès and A. M. Bruckstein, *Doubly Stochastic Pairwise Interactions for Agreement and Alignment*, preprint, arXiv:2012.13727, 2020.

[7] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, *Mixing beliefs among interacting agents*, Adv. Complex Syst., 3 (2000), pp. 87–98.

[8] A. G. Dimakis, S. Kar, J. M. Moura, M. G. Rabbat, and A. Scaglione, *Gossip algorithms for distributed signal processing*, Proc. IEEE, 98 (2010), pp. 1847–1864.

[9] Y. Elor and A. M. Bruckstein, *A "thermodynamic" approach to multi-robot cooperative localization*, Theoret. Comput. Sci., 457 (2012), pp. 59–75.

[10] N. Gantert, M. Heydenreich, and T. Hirscher, *Strictly weak consensus in the uniform compass model on* $\mathbb{Z}$, Bernoulli, 26 (2020), pp. 1269–1293.

[11] O. Häggström, *A pairwise averaging procedure with application to consensus formation in the Deffuant model*, Acta Appl. Math., 119 (2012), pp. 185–201.

[12] R. Hegselmann and U. Krause, *Opinion dynamics and bounded confidence models, analysis, and simulation*, J. Artif. Soc. Soc. Simul., 5 (2002).

[13] J. Lorenz, *Continuous opinion dynamics under bounded confidence: A survey*, Internat. J. Modern Phys. C, 18 (2007), pp. 1819–1838.

[14] J. Nagy, *Über algebraische Gleichungen mit lauter reellen Wurzeln.*, Jahresber. Dtsh. Math.-Ver., 27 (1918), pp. 37–43.

[15] A. Nedić, A. Ozdaglar, and P. A. Parrilo, *Constrained consensus and optimization in multi-agent networks*, IEEE Trans. Automat. Control, 55 (2010), pp. 922–938.

[16] T. Popoviciu, *Sur les équations algébriques ayant toutes leurs racines réelles*, Mathematica, 9 (1935), pp. 129–145.

[17] F. Proschan and M. Shaked, *Random averaging of vector elements*, SIAM J. Appl. Math., 44 (1984), pp. 587–590.

[18] A. K. Sahu, S. Kar, J. M. Moura, and H. V. Poor, *Distributed constrained recursive nonlinear least-squares estimation: Algorithms and asymptotics*, IEEE Trans. Signal Inform. Process. Netw., 2 (2016), pp. 426–441.

[19] A. Sarlette and R. Sepulchre, *Consensus optimization on manifolds*, SIAM J. Control Optim., 48 (2009), pp. 56–76.

[20] D. Tanny and J. Wellner, *Random Pairwise Averaging and Products of Random Matrices*, technical report, Department of Statistics, University of Rochester, Rochester, NY, 1981.