Consensus formation in the Deffuant model of opinion dynamics

Author: Joris Veefkind Castelló
Facultat de Física, Universitat de Barcelona, Diagonal 645, 08028 Barcelona, Spain.

Advisor: Dr. Luce Prignano

Abstract: This final project aims at exploring the emergence of consensus states in the already classic Deffuant model of opinion dynamics. In order to do so, four definitions are proposed and serve as the ground of the subsequent treatment. The results start off with a brief study of the general behavior of the model, as well as of the influence that both initial conditions and the convergence parameter μ exert upon it. After this, a comparison of the original model with two proposed variations of it is carried out, providing further insight into the mechanisms behind the phenomena of interest. Finally, a probabilistic study based on a simplified model of these mechanisms is introduced.

I. INTRODUCTION

A distinctly modern understanding of society presents us with a picture of a bunch of individuals interacting in different ways with a reduced number of their peers, usually much smaller than the total amount of people. What needs explanation within this picture is the "stunning global regularities" that are nonetheless displayed by social systems[1]. Assuming a secularized notion of universality, statistical explanations seem to emerge as the natural way to account for these kinds of regularities. It is owing to this that over the past few decades the mathematical, physical and computational tools developed by statistical physics, which in the course of last century achieved great success in explaining diverse physical phenomena, have started to appear useful in understanding specifically social phenomena and processes.

Opinion Dynamics, one of the manifold paths taken by these interdisciplinary explorations, focuses on "the dynamical processes of the diffusion and evolution of public opinions and social norms in human population"[1]. This encapsulates a wide array of phenomena such as collective decision making, fashion, minority opinion survival or the formation of consensus. The field as such can be traced back to the second half of the last century with the proposal of a number of models such as the well-known voter model put foward by Clifford and Sudbury in 1973[2]. These first models mainly considered a discrete number of possible opinions; more recently, models based on a finite continuum of opinions have been gaining prominence[3].

II. THE DEFFUANT MODEL

Consider the following situation: two people, each of whom holds their own judgement on a certain subject matter, meet for a coffee and have a chat concerning this particular topic. It might be plausible to assume that as long as their stances are not too different, such an interaction will bring the opinions of both sides closer than they were before. If, on the contrary, they hold very different - even opposed - opinions, these will not be affected by the interaction.

This assumption - that an interaction between two individuals with strongly differing opinions will not modify the opinions in question - is often called "bounded confidence", and lies at the bottom of the computational model proposed in [4] and henceforth referred to as Deffuant model (DM). The model at issue purports to describe the evolution of opinion states of an array of social agents, defined by an opinion on continuum from 0 to 1, by turning the situation described above into a simple norm of interaction.

To do so, it builds upon an array of N agents i with continuous opinions x_i. At each time step t, two agents x, x' are randomly picked out. If the difference between their opinions is greater than a certain range of confidence d, they will not interact. If, on the contrary, $|x - x'| < d$, both agents will re-adjust their opinions according to a convergence parameter μ such that:

$$x_{t+1} = x_t + \mu \cdot (x'_t - x_t)\,, \tag{1}$$

$$x'_{t+1} = x'_t + \mu \cdot (x_t - x'_t)\,. \tag{2}$$

As is the case with many opinion dynamics models, the behavior of agents in DM exhibits a clear tendency towards clustering, with final states consisting of groups of agents that are no longer able to interact with agents outside of their group. This project is centered around the emergence of consensus states - i.e., final states consisting in a single cluster of agents.

III. CONCEPTUAL CLARIFICATIONS

In their 2000 paper, Deffuant et al. swiftly mention the appearance of "populations of wings (corresponding to a few percent of the population) in the vicinity of the extreme opinions 1 and 0"[4]. They do not, however, discuss the formal features of their treatment of this phenomenon, leading in turn to a not all too clear notion of what they call a cluster.

Here, in the interest of providing an accurate characterisation of the different behaviors exhibited by the model, a cluster is defined as an ensemble of agents such that they are all within each other's range of confidence, outside of any other agent's range of confidence and constitute more than 5% of the population of any such ensemble. A wing, then, is any set of agents that meets the first two but not the last of these conditions. Neither wings nor clusters can split or merge with other wings or clusters. This, in turn, serves to define the final state of a system: a state in which the totality of the agents are distributed amongst clusters and wings. As these can neither split nor merge, the final state is a stationary one with regard to the distribution of wings and clusters.

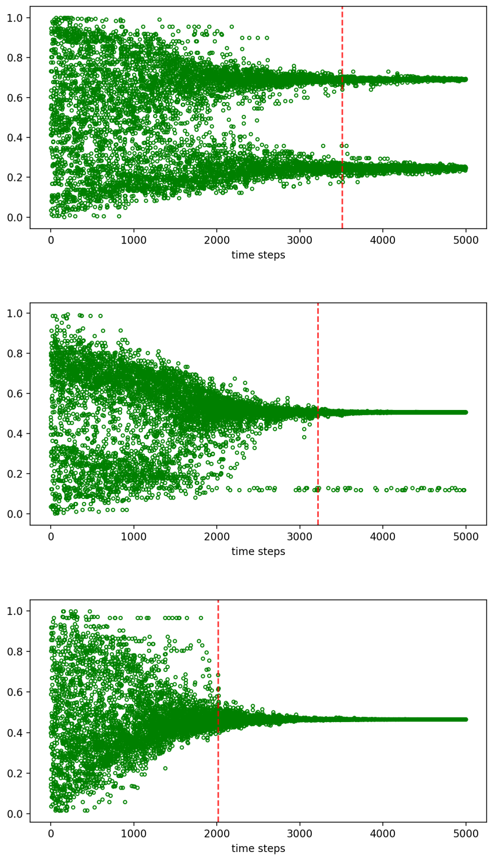

FIG. 1: Time chart of opinions exhibiting no consensus (upper), weak consensus (middle) and strong consensus (lower). The vertical red line is set at the time step at which the final stationary state is reached, in which all opinions are distributed into clusters and wings.

Given these definitions, it is possible to proceed to a fourth one, that of consensus. A consensus state is a stationary state of the system in which all agents are distributed into a single cluster and an arbitrary number of wings. On the grounds of the distinction between cluster and wing, it will prove useful to make a further distinction between consensus in general (with an arbitrary number of wings) and consensus in which no wings are formed at all. The former will henceforth be referred to as "weak consensus"; the latter, as "strong consensus".

The time charts in Figure 1 - depicting single simulations of DM on an array of 200 agents, with $\mu = 0.5$ - illustrate the definitions given in this section. The upper chart, corresponding to $d = 0.3$ depicts a state in which two clusters emerge in the final state, meaning that no consensus is reached at all; the one at the center, with $d = 0.35$, exhibits weak consensus; finally, the bottom one, with $d = 0.35$, ends up in a strong consensus state.

IV. CONSENSUS STATES IN THE DEFFUANT MODEL

Figure 2 presents the fraction of simulations reaching both weak and strong consensus states as a function of parameter d for DM. The data has been obtained by carrying out a total of 1000 simulations on an array of $N = 200$ agents, with parameter μ fixed at $\mu = 0.5$, and with different initial conditions (initial distribution of opinions) for each simulation. Unless stated otherwise, these will serve as the standard parameters used for the different results obtained along the project.

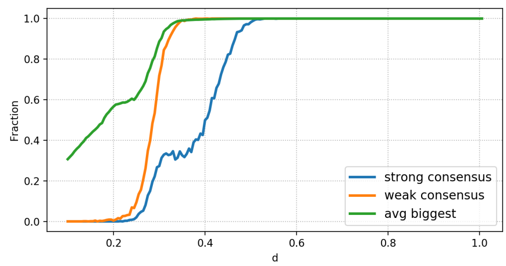

FIG. 2: Fraction of weak consensus, fraction of strong consensus and average fraction of agents that conform the biggest cluster as a function of d in final states of DM.

With regards to the consensus curves, the model exhibits different behaviors depending on the value of the confidence parameter d. In a first region, for values of d under 0.2, no consensus is reached whatsoever. At $d \approx 0.22$, the weak consensus curve undergoes a steep growth, with the weak consensus curve growing at a smaller rate. When the frequency of weak consensus states reaches its maximum value, the strong consensus curve plateaus at a frequency of around 0.35. At $d \approx 0.4$, the frequency of strong consensus starts rising again until it reaches its maximum at $d \approx 0.5$. From there onward, both weak and strong consensus are always attained by the system.

The disparity, in the $d \in (0.2, 0.4)$ range, between the weak and strong consensus curves, can be caused by two distinct factors: the purely stochastic process whereby a sequence of interaction is established in each simulation and the initial distribution of opinions, i.e. the initial conditions of the simulation.

A. Initial conditions

A better view of the way in which initial conditions determine the final states reached emerges when running a number of simulations on the same initial opinion distribution. Along these lines, Figure 3 depicts the fraction of strong consensus states obtained after running 1000 simulations on the same initial condition for 4 different initial conditions.

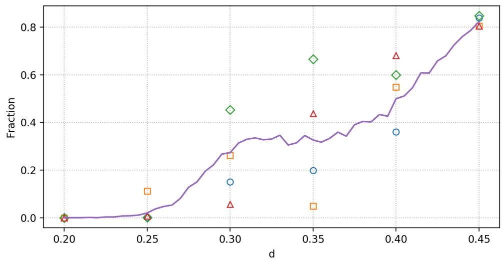

FIG. 3: Fraction of strong consensus for different values of d using 4 different initial conditions. Each marker depicts one initial condition. The solid line is the average fraction of strong consensus, and serves as a comparison.

The main conclusion to be extracted from these results is that the fraction of strong consensus obtained by simulating DM for a single initial conditions changes sharply when switching initial conditions in the $d \in [0.25, 0.45]$ range. In other words, some initial conditions are significantly more prone to the emergence of wings than others.

B. The impact of μ

Up to this point, the emergence of consensus states has been studied relative to the range of confidence parameter d and the initial conditions of the simulations. However, the convergence parameter μ might also play its role. To assess this possibility, Figure 4 compares the weak and strong consensus curves reached for different values of this parameter.

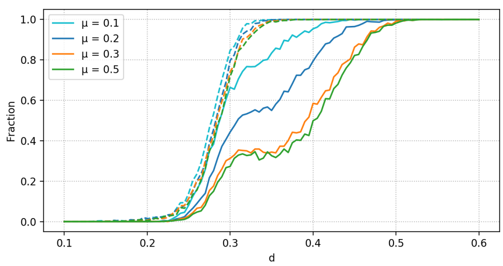

FIG. 4: Weak consensus (dashed lines) and strong consensus (solid lines) curves as a function of d for different values of μ.

As the figure shows, μ plays a decisive role upon the final states reached by the simulations: higher values of the parameter bring with them an increase in the formation of wings. This seems to contradict Deffuant et. al.'s paper, in which the following is claimed: "μ and N only influence the convergence time and the width of the distribution of final opinions (when a large number of different random samples are made)" [4]. However, the non-dependence on μ of the weak consensus curve suggests a possible explanation of their assertion on the conjecture that they employed a definition of cluster similar to the one used here, thus discarding wings from all their results.

C. The composition of wings

A further question in connection with the formation of weak and strong consensus states concerns the features of the wings, and splits up into two matters: the number of agents that constitute the wings and the total number of wings that appear.

Figure 5 depicts the number of agents per wing for different values of d and μ, chosen at the bounds of the plateau of the strong consensus curve, in a sample of 10000 weak consensus states with wings. The coloring of the bins informs about the number of wings sampled in states with a single wing (blue) and with two wings (orange). At the chosen values of μ and d, final states never display more than two wings.

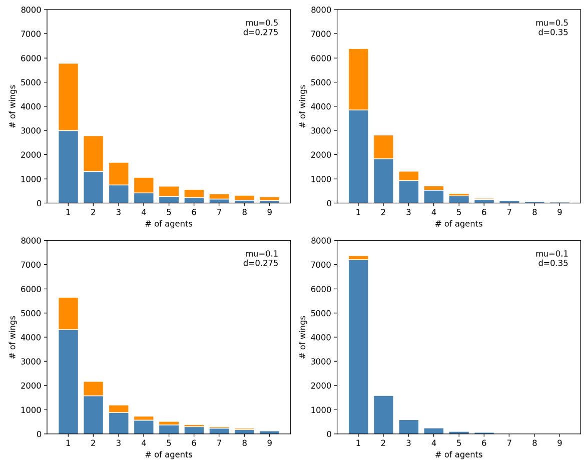

FIG. 5: Number of agents per wing for different values of d, μ in a sample of 10000 weak consensus states with wings, DM. In blue, wings pertaining to single wing final states. In orange, to final states with two wings.

The two parameters play a part in both the number of wings and the number of agents per wing. Nevertheless, without diving in too deep it is possible to say that as higher values of d tend to reduce the number of agents per wing, higher values of μ translate into a bigger chance of finding two wings instead of one in consensus.

V. VARIATIONS ON THE DEFFUANT MODEL

In order to further assess the mechanisms that explain the behavior of DM, two variations can be introduced. In the first one, which will be called "Asymmetrical Deffuant Model" (AM), only one of the two random interacting agents re-adjusts their opinion, making the interaction asymmetrical. This corresponds to a picture similar to that of the regular Deffuant model, but in which only one of the two agents of the interaction emits an opinion, while the other acts merely as a passive receiver disposed to adjust her opinion if the received opinion falls within her range of confidence. In the second variation, which will be called "Broadcasting model", only one agent x^* is randomly picked out. All the other agents x^i such that $|x^* - x^i| < d$ will re-adjust their opinions according to:

$$x^i_{t+1} = x^i_t + \mu \cdot \left(x^*_t - x^i_t\right). \quad (3)$$

The real world situation to which this model corresponds is that of a speaker that addresses a group of listeners ready to adjust their opinion if the received one turns out to be close enough to their own.

It is possible to put forward a classification of DM and its variants along two fundamental axes: symmetry and reach. In terms of their symmetry, the interactions that constitute both AM and BM are asymmetrical, while DM is based on a symmetrical interaction. Based on reach, the interaction of BM is a one-to-many interaction, while DM and its asymmetrical variation have one-on-one interactions.

A. Reach: the Broadcasting model

Figure 6 presents the fraction of both weak and strong consensus as well as the average fraction of agents in the biggest cluster as a function of parameter d for BM.

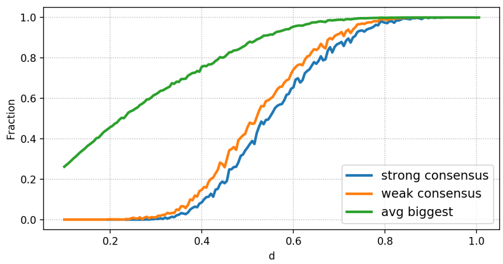

FIG. 6: Fraction of weak consensus, Fraction of strong consensus and average fraction of agents that conform the biggest cluster as a function of d for BM.

The comparison between Figure 6 and Figure 2 shows two very prominent differences. First, there is a disparity in the slopes of the curves of both models: DM reaches weak consensus states with a frequency of 1 at $d \approx 0.35$, with the same happening for strong consensus states at $d \approx 0.51$. That is significantly quicker than BM, which doesn't behave like that until past $d = 0.8$. Besides this, DM presents a behavior that has no counterpart in BM for values of d between 0.3 and 0.4. In this range of the parameter, almost all - if not all - the simulations reach weak consensus states, while the frequency of strong consensus remains more or less steady. In other words, for a certain range of the d value, DM reaches consensus states almost always, while displaying wings (and hence, weak consensus) in about 60% of the simulations carried out (see Figure 2).

The differences in the behavior exhibited suggest deeper dissimilarities in the mechanisms that underlie consensus formation in both models. With symmetry and reach being the two main differences between the protocols of interaction of the models discussed here, a question about the influence of each of these factors on the behaviors exhibited comes up.

B. Symmetry: the Asymmetrical Deffuant Model

Figure 7 presents the fraction of both weak and strong consensus as well as the average fraction of agents in the biggest cluster as a function of parameter d for AM.

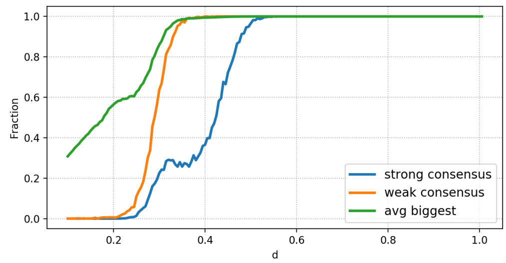

FIG. 7: Fraction of weak consensus, fraction of strong consensus and average fraction of agents that conform the biggest cluster as a function of d for AM.

A comparison between Figure 7 and Figure 2 shows that the behaviors of the curves obtained from simulating AM do not depart from those exhibited by the regular DM in any significant way.

Owing to the classification made above, the assessment of the comparison between DM and its asymmetrical variation leads to a further conclusion: the differences between the regular DM and BM are caused by the difference in their reach rather than by the symmetry of the models. A corollary is that one on one interactions favor the emergence of both weak and strong consensus.

VI. THE PROBABILITY OF WING FORMATION

In order to get a deeper insight into the mechanisms at work in the emergence of consensus states within DM, a simple probabilistic calculation has been carried out. Its main aim is to approach, in a very simplified way, the mechanisms by which an agent at one of the limits of the opinion array might be left outside of the range of influence of every other agent, thus constituting a wing.

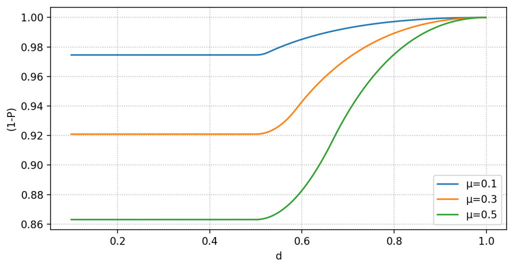

FIG. 8: $(1-P)$ as a function of d for different values of μ.

Consider, in the framework of the model, an agent $x_0 = 0$ and one of her neighbors x_1 such that $x_1 < d$. It is interesting to study the probability of x_1 falling outside of x_0's range of influence/neighborhood after interacting with a third agent x_2. For a given x_1, this can be calculated by dividing Δx_2, the interval in which x_2 can be chosen such that x_1 is pulled out of x_0's neighborhood, by the total length of x_1's neighborhood. Then, integrating over x_1 provides the total probability P that after an interaction between a neighbor of x_0 and a neighbor of a neighbor of x_0, the former is pulled out of x_0's neighborhood. The probability function $P(\mu, d)$ obtained, displayed in Figure 8, is the following:

$$P = \frac{\mu - 2}{\mu}\left[\ln 2 - \ln(2-\mu) + \frac{\mu}{\mu - 2}\right], \tag{4}$$

$$P = \frac{1 - 4d^2 - \mu + 4d\mu}{2d\mu} - \frac{(-2+\mu)\ln(2d - d\mu)}{\mu}, \tag{5}$$

$$P = -\frac{(d-1)^2\mu}{2d(\mu-1)}, \tag{6}$$

where (3) corresponds to $d \leq 0.5$, (4) to $0.5 < d \leq 1/(2-\mu)$ and (5) to $d > 1/(2-\mu)$.

Higher values of μ imply a higher probability of neighbor loss, diminishing the probability of strong consensus states. This similitude might, to some extent, justify the claim that the simplified model resembles actually occurring mechanisms behind wing formation.

VII. CONCLUSIONS

This study contributes to the growing understanding of simple opinion dynamics models through the analysis of consensus formation in the Deffuant model. A first significant result has been the thorough establishment of the influence of both the initial distribution of opinions and the convergence parameter μ in strong consensus formation. Moreover, the comparison between the original model and its proposed variations has led to the conclusion that the symmetry of the interaction is not relevant to the behavior of the model. By contrast, reach has shown decisive in this respect. Finally, a tentative approach at a probabilistic explanation of the behavior of the model through the singling out of its most basic mechanisms has been laid out.

At a more general level, the project put forward the importance of establishing clear definitions that enable detailed description of different behaviors and, most importantly, adequate comparison of results with other research. Further insight could be gained by conducting similar inquiries with other models or variations and mapping out the similarities and differences between these. Most importantly, the design and implementation of real-life experiments that admit a mapping in terms of these models appears indispensable to further delve into the probabilistic study of social processes

Acknowledgments

I am profoundly grateful to my advisor, Dr. Luce Prignano. Her guidance into the field of opinion dynamics, the insight she has provided me with and the questions she has raised have been the main drive of the project.

[1] Xia, Haoxiang & Wang, Huili & Xuan, Zhaoguo. (2011). *Opinion Dynamics: A Multidisciplinary Review and Perspective on Future Research.* 591–646. IJKSS. 2. 72-91.

[2] Clifford, P., & Sudbury, A. (1973). A Model for Spatial Conflict. *Biometrika, 60*(3), 581–588.

[3] Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. *Statistical physics of social dynamics. Reviews of Modern Physics, 81*(2), 591–646.

[4] Deffuant, G., Neau, D. B., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems, 03*(01n04), 87–98.

[5] Fortunato, Santo. (2004). On the Consensus Threshold for the Opinion Dynamics of Krause-Hegselmann. *International Journal of Modern Physics C, 16*(02), 259–270.