On the Impact of Social Media Recommendations on Consensus of Discrete Opinions (Short Version)

Vincenzo Auletta[1], Antonio Coppola[1] and Diodato Ferraioli[1,*]

[1]Università degli Studi di Salerno, via Giovanni Paolo II, 132, Fisciano, SA, 84084, Italy

Abstract

We consider a discrete opinion formation problem in a setting where agents are influenced by both information diffused by their social relations and from recommendations received directly from the social media manager. We study how the "strength" of the influence of the social media and the homophily ratio affect the probability of the agents of reaching a consensus and how these factors can determine the type of consensus reached.

In a simple 2-symmetric block model we prove that agents converge either to a consensus or to a persistent disagreement. In particular, we show that when the homophily ratio is large, the social media has a very low capacity of determining the outcome of the opinion dynamics. On the other hand, when the homophily ratio is low, the social media influence can have an important role on the dynamics, either by making harder to reach a consensus or inducing it on extreme opinions.

Finally, in order to extend our analysis to more general and realistic settings we give some experimental evidences that our results still hold on general networks.

Keywords
Opinion Dynamics, Consensus, Social Learning

1. Introduction

Over the last years, we witnessed a rapid rise of the role of online social networking platforms, such as Facebook or Twitter, in our life. As a consequence, individuals increasingly rely on these social platforms to get news and form their opinions. E.g., according to Pew Research Center survey in 2018 [2] 68% of American adults get news on social media, a significant rise from 49% of 2012. Moreover, it has been observed that social media may have a relevant effect in many real-world critical settings, such as in electoral campaigns [3, 4]. For example, some studies showed that the social media may lead to extremism [5] and polarization in individuals' opinions [6].

Hence, it urges to understand how the social media may affect the process of opinion formation of their users. To this aim, several models have been introduced to describe how the opinions of agents evolve under the effect of the social influence. The first such model, due to [7], states that each agent adopts an opinion that averages among the ones of individuals which she interacts

with. One of the most relevant extensions of this model is, undoubtedly, the dynamics described by [8] (see also, the work of [9]), that limits the effects of social influence by holding agents close to their original ideology. These models assume that opinions may take values in a continuous space, and agents may express any value in this space. However, in several real settings, i.e., electoral contexts, the number of alternatives around which opinions should converge are limited. Moreover, even if opinions can take values that do not match any alternative, these cannot be expressed due to the limitedness of the options according to which opinions are expressed (e.g., polls, finite-precision ranks, etc.). For these reasons, continuous models turn out to be scarcely representative in some settings, and discrete versions of these models have been proposed in which agents' opinions must belong to a discrete set [10, 11].

However, in several settings it is not sufficient to take into account only the social influence among agents', but we have also to understand how the social media may influence the opinion formation process, and whether and how it is necessary to mitigate in some way the effects it provokes.

There has been recently an increasing interest on these questions. In particular, most of the recent literature in the social choice area focuses on the opportunity for the social media to manipulate the opinion formation process in order to support a target opinion. Different forms of manipulations have been studied, such as seeding, edge addition/deletion, and alteration of the order of changes (see Related Works section for more details).

In this work, we deviate from this approach, and we do not consider the social media as a manipulator. That is, the social media does not have a target that should be promoted, but it only acts as a platform for sharing information. However, social media's goal is to maximize the activity of the agents on the platform and it implements policies about which, when, and to whom information are shared, in order to maximize engagement of users to their service. While the actual implementation of these policies is private, it is evident that users are more likely to be exposed to information closer to their own opinion [12, 13]. [14] have proved that agents have larger probability of interacting (by viewing, liking, or re-sharing) with this kind of information, witnessing in this way their major engagement with the social media.

In this paper we want to answer the following question: how much a social media implementing these policies can influence the opinion formation process? This problem has been recently addressed by [15] in the context of continuous opinion formation processes. Their answer depends on the strength of the influence of the social media platform on individuals: if this is high, then agents' opinions tend to extremes; if low, agents' opinion tend to converge; in the middle, instead, some non-extreme disagreement can occur.

However, the continuous approach adopted by [15] does not fit with many real world critical contexts, such as in voting, in which we usually have a discrete and limited number of candidates around which opinions should converge. For this reason, in this work, we will depart from the work of [15], by focusing on the discrete opinion formation process, as defined by [11].

In our work (a preliminary version appeared as [16]), we evaluate the impact of social media recommendations with respect to their influence on the ability of users to reach a consensus. Indeed, the likelihood that a consensus is reached has been widely adopted for comparing different opinion models, and for evaluating the impact that variations on the model may have on opinion formation [7, 17, 18]. Note also that consensus is a required goal in many practical settings: from the analysis of collective behaviour of flocks and swarms [19, 20], to sensor fusion

[21], to formation control for multi-robot systems [22, 23, 24].

Our Model We consider $n \geq 2$ agents whose relationships are embedded into a social network modelled as an undirected weighted graph $G = (V, E, w)$, where each vertex of the graph represents an agent. Each agent i keeps an opinion $x_i^0 \in \Theta = \{-1, -1 + \delta, \ldots, -\delta, 0, \delta, \ldots, 1 - \delta, 1\}$ for some $0 < \delta \leq \frac{1}{2}$. We will sometimes denote δ as the *discretization factor* of Θ. One may think about Θ as the set of alternatives (e.g., candidates to an election) on which agents' opinions need to converge: note that we are assuming that there is no way for an agent to express an opinion that does not corresponds to an alternative. Observe that $|\Theta| = 2 \lceil \frac{1}{\delta} \rceil + 1$. Let $\mathbf{x} = (x_1, \ldots, x_n)$ be a profile of opinions held by players, where x_i is the opinion kept of player i.

The opinions of agents are influenced by their social relationships. Specifically, we assume that, for each edge $(i, j) \in E$, opinions of agents i and j are mutually influenced and the weight $w_{ij} > 0$ of the edge models the strength of this influence.

Moreover, we assume that the opinion of an agent can be also influenced by recommendations received directly from the social media and not diffused through their own neighbours. We assume that the social media can present to the agents different recommendations, tailoring them on their interests. In particular, we assume that the social media has a discrete subset Ω of $[-1, 1]$, representing the available information, and it decides to present to an agent with opinion x the information $s(x) \in \Omega$, where the function $s \colon \Theta \to \Omega$ models the recommendation procedure adopted by the media. Clearly, since the social media is interested in increasing the engagement of their users to the platform, it is interested in advertising to users information that best matches their profile. Thus, e.g., in an electoral setting, the social media will recommend right parties to right-oriented agents, left parties to to left-oriented agents, and moderate party to remaining agents.

Thus, at each time step t agents update their opinions depending on the opinions held by their social relations and the recommendations received by the social media. We denote by $\mathbf{x}^t$ the profile of opinions held by agents at time t.

In this work, following the model introduced by [15], we will consider a specific choice for Ω and s: in particular, we assume $\Omega = \{-1, 0, 1\}$ (we will sometimes refer to the elements of Ω as "extreme left", "extreme moderate", and "extreme right" information or opinions), and assume s being a symmetric threshold function such that $s(x) = -1$ if $x < -\lambda$, $s(x) = 1$ if $x > \lambda$, and $s(x) = 0$ otherwise, for some $0 < \lambda < 1$. While this choice is clearly simplifying the model, it still leads to interesting results about how these social media recommendations may affect the chance that agents may reach a consensus.

The combined influence of neighbours and social media recommendations may lead an agent to update her opinion. In this work, we follow the principles of the model presented by [7] to represent how the opinion is updated. Specifically, since our focus is on a setting with discrete opinions, we will adapt to our model the discrete generalization of the DeGroot model defined by [11]: at each step $t \geq 1$, agent i will choose the opinion x that minimizes $c_i(x, \mathbf{x}^{t-1}) = b(x - s(x_i^{t-1}))^2 + \sum_{j \colon (i,j) \in E} w_{ij}(x - x_j^{t-1})^2$, where $b > 0$ is the weight of the influence of the social media on agents, and $\mathbf{x}^{t-1} = (x_1^{t-1}, \ldots, x_n^{t-1})$ is the opinion profile at the previous time step. We notice that this setting can be equivalently described as a game:

agents are the players, opinions are their strategies, and the function c_i is the cost function of player i. According to this game-theoretic viewpoint, the opinion update consists essentially of selecting the *best-response* strategy, i.e. the one that minimizes the cost of the player given the strategies currently selected by other players and the social media.

We say that an opinion profile $\mathbf{x}^t = (x_1^t, \ldots, x_n^t)$ is a *consensus* (on opinion $\overline{x}$) if $x_i^t = \overline{x}$ for every i. Moreover, we say that an opinion profile $\mathbf{x}^t = (x_1^t, \ldots, x_n^t)$ is *stable* if it is a Nash equilibrium of the corresponding game, i.e. x_i^t minimizes $c_i(x, x^t)$ for every agent i. It is easy to see that a consensus on an extreme opinion, say, e.g., 1, is always a stable profile. Hence, in this opinion game, a Nash Equilibrium always exits.

Although a stable profile always exists, for given G and b, there may be multiple stable opinion profiles, and which one is reached depends on the way in which agents update their opinions. In the literature, the DeGroot model has been associated to different update rules: The most popular rules are: i) *synchronous rule*, where at each time step t all the agents update their options; ii) *asynchronous rule*, where at each time step t, a single agent, arbitrarily chosen, is allowed to update his/her opinion.

Our Contribution   In our work we first focus on a very simple class of networks, namely *symmetric two-block model*, already analyzed by [15], in which agents are separated in two components, and agents from the same component have the same initial opinion and receive the same influence from individuals inside and outside their component. Despite of the simplicity of this network, it highlights a very important difference with respect to the results given by [15]: namely, the impact of the social media not only depends on the strength of the social media influence, but also on the *homophily ratio*, that is how much individuals weight their similar compared to others. This measure has been often showed to be a key attribute in opinion formation dynamics (see, e.g., [25]). Hence, our results show a better alignment with respect to the previous literature than the one given by [15].

Specifically, we showed that whenever the strength of the social media influence is large, consensus is essentially impossible to achieve whenever the initial opinions of the two groups are far from each other. Interestingly, for these initial opinions, consensus is also impossible to achieve when the homophily ratio is large, but the strength of the social media is very small. We also showed how the chance of reaching a consensus changes with respect to how extreme are the initial opinions in the two groups. Finally, when initial opinions are instead close to each other, we show that consensus is always possible, but the likelihood of reaching a consensus increases when the homophily ratio is large or the strength of the social media is low.

Future Directions   We conjecture that these findings hold not only for the simple symmetric two-block model, but also for more complex networks whenever initial opinions can be partitioned in two macro-blocks. This conjecture is supported by a massive set of experiments [26] both on synthetic and on real networks: all our experiments show that the dynamics essentially follows the behaviour prescribed by results on the symmetric two-block model as the strength of the social media, the homophily ratio, and the value of initial opinions change. It would be of primary interest to settle this conjecture.

Moreover, while our work focuses mainly on how and how much the social media may

influence the likelihood that agents reach a consensus, it would be interesting also to deepen our analysis by evaluating how the social media can influence, not only the probability of consensus, but also the kind of equilibria that can be reached by the opinion formation process.

In our work we focused on a classical opinion formation model. However, we believe that it would be undoubtedly interesting to analyze whether our results extend to more complex (but more realistic) opinion formation models.

In our analysis, we restricted the opinion space of the social media Ω to the three extreme values. To consider different choices of these values or an higher cardinality of Ω would be clearly of interest, though we will expect that such an analysis will give results very similar in spirit to the ones proved in our work (but with an explosion of possible cases). Similar considerations can be done about extending our mono-dimensional representation of opinions to higher dimensional representations.

Even if, our experimental results highlight a large adherence to the theoretical findings obtained for the symmetric two-block model, some small differences exist among the results for different network structures. It would be then interesting to understand whether and how these differences may be motivated through a detailed study of the relationship among the impact of the social influence and the structural and topological properties of the social network.

Related Works Several extensions have been recently proposed to the seminal models by DeGroot and by Friedkin and Johnsen (and their discrete counterparts), by considering only limited interaction by agents [27, 28], or an evolving environment [17, 29, 30, 31, 18], or both repulsive and attractive interaction [32, 33]. Despite their larger adherence with many real world aspects, however none of these variants has received the same level of interest as the models by DeGroot and by Friedkin and Johnsen. Moreover, the simplicity of the latter models allows a more clear analysis of the influence of social media, by untying it from the complexities of the former models.

Consensus in opinion formation has been object of intense research since the seminal work of [7]. Indeed, most works aim to evaluate opinion formation models based on their ability to reach a consensus [17, 18]. Many other works try to characterize the parameters that enable a given dynamics to reach consensus [34, 35, 36]. In this work we pursue both approaches: on one side, we investigate on how the social media recommendations may vary the probability that a consensus is reached; on the other side, we identify the settings, in terms of homophily ratio, strength of the social media influence, and initial agents' opinions, where the probability of consensus is larger.

The study of the influence of a (non-manipulating) social media on the opinion formation process has been initiated by [15], where, as described above, the focus is on continuous opinions, while we here consider discrete opinions.

Many works instead focus on manipulation of the opinion formation process in social networks, in particular in the framework of election manipulation. The first and most studied manipulation technique is *seeding*, that consists in selecting a set of sources of news from which to start a successful viral campaign in favour of a designed candidate or against her competitors [37, 38, 39, 40, 41, 42]. Another kind of manipulation that received large interest consists in adding or deleting links [43, 44, 42, 41]: these may be implemented by social media by hiding

the content of a "friend" or "neighbour" in the social network, or promoting the content of non-friends (e.g., as advertised content or through the mechanism of friend suggestion). A last kind of manipulation that recently received a lot of interest consists in guiding the dynamics by influencing the order in which agents are prompted to update their opinion (e.g., by delaying the visualization of a news) so that they will update only when there are enough friends to push them towards the desired candidate [45, 46, 47, 36, 48]. We note that, as described above, our work differs from all these works, since we are not considering a social media operating with the goal of promoting a specific candidate.

Acknowledgments

References

[1] R. De Benedictis, M. Castiglioni, D. Ferraioli, V. Malvone, M. Maratea, E. Scala, L. Serafini, I. Serina, E. Tosello, A. Umbrico, M. Vallati, Preface to the Italian Workshop on Planning and Scheduling, RCRA Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion, and SPIRIT Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy (IPS-RCRA-SPIRIT 2023), in: Proceedings of the Italian Workshop on Planning and Scheduling, RCRA Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion, and SPIRIT Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy (IPS-RCRA-SPIRIT 2023) co-located with 22th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2023), 2023.

[2] E. Shearer, K. E. Matsa, News use across social media platforms 2018, Available at https://www.pewresearch.org, 2018.

[3] A. I. Androniciuc, Using Social Media In Political Campaigns. Evidence From Romania, SEA - Practical Application of Science (2016) 51–57.

[4] T. Fujiwara, K. Müller, C. Schwarz, The effect of social media on elections: Evidence from the United States, Technical Report, National Bureau of Economic Research, 2021.

[5] M. C. Benigni, K. Joseph, K. M. Carley, Online extremism and the communities that sustain it: Detecting the isis supporting community on twitter, PLOS ONE 12 (2017) 1–23.

[6] H. Allcott, L. Braghieri, S. Eichmeyer, M. Gentzkow, The welfare effects of social media, American Economic Review 110 (2020) 629–76.

[7] M. DeGroot, Reaching a consensus, Journal of the American Statistical Association 69 (1974) 118–121.

[8] N. Friedkin, E. Johnsen, Social influence and opinions, The Journal of Mathematical Sociology 15 (1990) 193–206.

[9] D. Bindel, J. M. Kleinberg, S. Oren, How bad is forming your own opinion?, Games and Economic Behavior 92 (2015) 248–265.

[10] F. Chierichetti, J. Kleinberg, S. Oren, On discrete preferences and coordination, Journal of Computer and System Sciences 93 (2018) 11 – 29.

[11] D. Ferraioli, P. Goldberg, C. Ventre, Decentralized dynamics for finite opinion games, Theoretical Computer Science 648 (2016) 96–115.

[12] E. Bakshy, S. Messing, L. A. Adamic, Exposure to ideologically diverse news and opinion on facebook, Science 348 (2015) 1130–1132.

[13] R. Levy, Social media, news consumption, and polarization: Evidence from a field experiment, American economic review 111 (2021) 831–70.

[14] Y. Halberstam, B. Knight, Homophily, group size, and the diffusion of political information in social networks: Evidence from twitter, Journal of public economics 143 (2016) 73–88.

[15] J. Anunrojwong, O. Candogan, N. Immorlica, Social learning under platform influence: Extreme consensus and persistent disagreement, SSRN (2020).

[16] V. Auletta, A. Coppola, D. Ferraioli, On the impact of social media recommendations on opinion consensus, in: AIxIA 2021, 2022.

[17] R. Hegselmann, U. Krause, Opinion dynamics and bounded confidence: Models, analysis and simulation, Journal of Artificial Societies and Social Simulation 5 (2002) 1–24.

[18] V. Auletta, A. Fanelli, D. Ferraioli, Consensus in opinion formation processes in fully evolving environments, in: AAAI, 2019, pp. 6022–6029.

[19] R. Olfati-Saber, Flocking for multi-agent dynamic systems: Algorithms and theory, IEEE Transactions on automatic control 51 (2006) 401–420.

[20] A. V. Savkin, Coordinated collective motion of groups of autonomous mobile robots: Analysis of vicsek's model, IEEE Transactions on Automatic Control 49 (2004) 981–982.

[21] R. Olfati-Saber, J. S. Shamma, Consensus filters for sensor networks and distributed sensor fusion, in: CDC, 2005, pp. 6698–6703.

[22] M. Egerstedt, X. Hu, Formation control with virtual leaders and reduced communications, IEEE Transactions on Robotics and Automation 17 (2001) 947–951.

[23] H. G. Tanner, G. J. Pappas, V. Kumar, Leader-to-formation stability, IEEE Transactions on Robotics and Automation 20 (2004) 443–455.

[24] Z. Lin, B. Francis, M. Maggiore, Necessary and sufficient graphical conditions for formation control of unicycles, IEEE Transactions on Automatic Control 50 (2005) 121–127.

[25] P. Dandekar, A. Goel, D. T. Lee, Biased assimilation, homophily, and the dynamics of polarization, Proceedings of the National Academy of Sciences 110 (2013) 5791–5796.

[26] A. Coppola, On the impact of social media recommendations on opinion-consensus, https://github.com/CoppolaAntonio/On-the-Impact-of-Social-Media-Recommendations-on-Opinion-Consensus, 2022.

[27] D. Fotakis, D. Palyvos-Giannas, S. Skoulakis, Opinion dynamics with local interactions, in: IJCAI, 2016, pp. 279–285.

[28] D. Fotakis, V. Kandiros, V. Kontonis, S. Skoulakis, Opinion dynamics with limited information, in: WINE, 2018, pp. 282–296.

[29] K. Bhawalkar, S. Gollapudi, K. Munagala, Coevolutionary opinion formation games, in: STOC, 2013, pp. 41–50.

[30] V. Bilò, A. Fanelli, L. Moscardelli, Opinion formation games with dynamic social influences, Theoretical Computer Science 746 (2018) 444–458.

[31] D. Ferraioli, C. Ventre, Social pressure in opinion games, in: IJCAI, 2017, pp. 3661–3667.

[32] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Generalized discrete preference games, in: IJCAI, 2016, pp. 53–59.

[33] E. Acar, G. Greco, M. Manna, Group reasoning in social environments, in: AAMAS, 2017, pp. 1296–1304.

[34] M. Feldman, N. Immorlica, B. Lucier, S. M. Weinberg, Reaching consensus via non-bayesian asynchronous learning in social networks, in: APPROX/RANDOM, 2014, pp. 192–208.

[35] E. Mossel, J. Neeman, O. Tamuz, Majority dynamics and aggregation of information in social networks, Autonomous Agents and Multi-Agent Systems 28 (2014) 408–429.

[36] V. Auletta, D. Ferraioli, G. Greco, On the complexity of reasoning about opinion diffusion under majority dynamics, Artificial Intelligence 284 (2020) 103288.

[37] D. Kempe, J. Kleinberg, E. Tardos, Maximizing the spread of influence through a social network, in: KDD, 2003, pp. 137–146.

[38] B. Wilder, Y. Vorobeychik, Controlling elections through social influence, in: AAMAS, 2018, pp. 265–273.

[39] F. Corò, E. Cruciani, G. D'Angelo, S. Ponziani, Exploiting social influence to control elections based on scoring rules, in: IJCAI, 2019, pp. 201–207.

[40] M. Abouei Mehrizi, F. Corò, E. Cruciani, G. D'Angelo, Election control through social influence with unknown preferences, in: COCOON, 2020, pp. 397–410.

[41] M. Castiglioni, D. Ferraioli, N. Gatti, G. Landriani, Election manipulation on social networks: Seeding, edge removal, edge addition, Journal of Artificial Intelligence Research 71 (2021) 1049–1090.

[42] R. Bredereck, E. Elkind, Manipulating opinion diffusion in social networks, in: IJCAI, 2017, pp. 894–900.

[43] S. Sina, N. Hazon, A. Hassidim, S. Kraus, Adapting the social network to affect elections, in: AAMAS, 2015, pp. 705–713.

[44] V. Auletta, D. Ferraioli, V. Savarese, Manipulating an election in social networks through edge addition, in: AI*IA, 2019, pp. 495–510.

[45] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Minority becomes majority in social networks, in: WINE, 2015, pp. 74–88.

[46] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Information retention in heterogeneous majority dynamics, in: WINE, 2017, pp. 30–43.

[47] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Robustness in discrete preference games, in: AAMAS, 2017, pp. 1314–1322.

[48] V. Auletta, D. Ferraioli, G. Greco, Optimal majority dynamics for the diffusion of an opinion when multiple alternatives are available, Theoretical Computer Science 869 (2021) 156–180.