Calibration of an agent-based model for opinion formation through a retweet social network

Loretta Mastroeni[a], Maurizio Naldi[b] and Pierluigi Vellucci[a]

[a]Dept. of Economics, Roma Tre University, Via Silvio D'Amico 77 00145 Rome, Italy

[b]Dept. of Law, Economics, Politics and Modern languages, LUMSA University, Via Marcantonio Colonna 19 00192 Rome, Italy

Abstract

Calibration of agent-based models (ABM) for opinion formation is needed to set their parameters and allow their employment in the real world. In this paper, we propose to use the correspondence between the agent-based model and the social network where those agents express their opinions, namely Twitter. We propose a calibration method that uses the frequency of retweets as a measure of influence and allows to obtain the influence coefficients in the ABM by direct inspection of the weighted adjacency matrix of the social network graph. The method has a fairly general applicability to linear ABMs. We report a sample application to a Twitter dataset where opinions about wind power (where turbines convert the kinetic energy of wind into mechanical or electrical energy) are voiced. Most influence coefficients (76%) result to be zero, and very few agents (less than 5%) exert a strong influence on other agents.

Keywords

Opinion formation, Agent-based models, Twitter, Calibration

1. Introduction

Agent-based models (ABM) are increasingly used to analyse opinion formation (see the survey in [1]), as an alternative to econophysics models [2]. Some examples with general applicability are described in [3, 4, 5, 6]; they are also applied to study specific phenomena such as equality bias [7] or personal finance decisions [8, 9].

However, models with no application to real world data may be too abstract. In the mathematical model describing the interactions among agents, we need to set the parameters governing those interactions, i.e. to calibrate those models. Calibration allows us to obtain realistic expectations about the behaviour of a social group. Very few studies have attempted to calibrate agent-based models. We can group them into two classes, where data are obtained respectively through a laboratory experiment of from the observation of a real social network.

One of the first example of the former class is the celebrated Friedkin-Johnsen model [10], which has been evaluated through data collected on small groups in a laboratory [11], where the authors asked their subjects to estimate the extent to which each other group member influenced their final opinion by means a mechanism based on a reward paid in poker chips. Another example is provided in [12], where participants in the experiment expressed their

WOA 2020: Workshop "From Objects to Agents", September 14–16, 2020, Bologna, Italy

loretta.mastroeni@uniroma3.it (L. Mastroeni); m.naldi@lumsa.it (M. Naldi); pierluigi.vellucci@uniroma3.it (P. Vellucci)

opinion about the best location for a new leisure center (on the line across two towns), and the influence could be observed through the evolution of one's opinion after examining the opinions of others.

Influence in real contexts has instead been studied through social networks. An advice network inside a manufacturing company was considered in [13] on the basis of the data provided in [14]. The influence of any individual was assumed to be proportional to the number of people who seek advice from him/her. A political context, namely the American Senate, was instead investigated in [15], where co-sponsorship of bills was taken as a measure of influence. The opinion value for each senator was assumed to be the fraction of votes when he/she was present and voted with the majority. Finally, a co-habitation context, namely a university student dormitory, was analysed in [16], where the Social Evolution dataset enclosed in [17] was employed. A similar context is also the stage for the study contained in [17], where surveys were conducted monthly on social relationships, health-related habits, on-campus activities and other issues.

Here we wish to propose a method to calibrate agent-based models for opinion formation, considering data extracted from an online social network, namely Twitter. This contrasts with the contributions appeared in the literature so far, where just physical social networks have been considered. The abundance of data appearing in online contexts makes them a natural choice to look at for our goal. We measure the influence by the frequency of retweeting, which makes our methods applicable to any social network where reposting of opinions is allowed.

Our major original contributions can be summarised as follows:

- we propose a calibration method based on the adjacency matrix in an online social network;
- we provide a systematic assessment of its applicability, considering a taxonomy of agent-based models;
- we demonstrate its application using Twitter opinions on wind power;
- for that specific context, we show that, though most actors exert an influence, very few exert a strong influence (i.e., being retweeted more than once);
- for that specific context, we show that the influence of any agent is limited to one other agent in most cases so that the matrix of influence coefficients is sparse.

2. The retweet network

We base our calibration on Twitter data. In this section, we report some basic information about Twitter and using retweets to calibrate an agent-based model.

Twitter is a popular messaging service, with 330 million monthly active users (see https://www.oberlo.com/blog/twitter-statistics), where messages (aka tweets) are no longer than 140 characters [18]. Message receiving uses an opt-in mechanism: you decide to *follow* somebody and receive all his/her updates. Messages may include *hashtags*, i.e., # symbol terms that associate

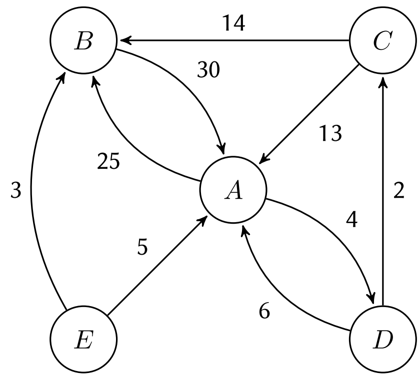

Figure 1: Toy retweet network

a topic to the message: all messages concerning a specific topic (as identified by the hashtags included in those messages) can be retrieved at once by searching for that hashtag.

Twitter allows retweeting a message, i.e., reposting somebody else's tweet. This is a sign of support for somebody else's opinion, stronger than just following him/her because it is an uncritical sharing of his/her opinions (retweeting through the Retweet button does not allow to add comments). It is also specific because it concerns a single message. Being retweeted is a sign of influence. We are using retweets to measure the influence of somebody's opinion on other participants in a social network. In particular, we build a retweet network. Retweet networks have been designated (see Chapter 4 of [19]) as a major tool to analyse the influence of twitterers. In our network each node represents a twitterer (an agent in the corresponding agent-based model) and an edge from node A to node B means that A has been retweeted by B. The weight associated to that edge is the number of retweets: the weighted out-degree of a node is the measure of the associated agent's influence. In the end, we obtain a weighted directed network. It is to be noted that this network allows us to quantify the influence of each agent on each other agent, which is what we need in our any-to-any agent-based model, rather than the overall influence of a twitterer (as in [20]).

In Figure 1, we show a toy retweet network made of five twitterers to illustrate this concept. We see that C has been retweeted 13 times by A, and 14 times by B, but has never retweeted either A or B. Retweeting is not symmetric in general.

3. Calibration methodology

Our aim is to set the parameters defining an agent-based model for opinion formation, i.e., to calibrate the model, exploiting the data retrieved from Twitter. In this section, we describe our calibration method and highlight its applicability.

For the time being, without loss of generality, we refer to an agent-based model such as described in [5], where the generic agent i features a quantity $x_i(t)$ representing its opinion at time t. Then we denote the opinions of all the agents at time t by the vector $\underline{x}(t)$; its evolution over time follows the state updating equation:

$$\underline{x}(t+1) = S\underline{x}(t), \tag{1}$$

with S being the matrix of influence coefficients. Namely, the element s_{ij} of S describes the effect of the opinion of the agent j on the opinion of the agent i.

The process leading from Twitter activity to the matrix of influence coefficients is made of the following phases:

1. scraping Twitter data;
2. building the retweet network and extracting the adjacency matrix;
3. mapping the edge weights on the influence coefficients.

The first phase consists in retrieving the list of tweets/retweets concerning a specific hashtag (or any combination of hashtags and keywords). This phase is often referred to as *Twitter Scraping* and can be easily accomplished using Twitter's API (Application Programming Interface), accessible upon opening a *Twitter developer account*. Tweets were retrieved using the R package twitteR [21]. The search index has a 7-day limit, which means that only tweets posted in the latest seven days will be retrieved. Our inspection interval falls in the week ending on December 9, 2019.

As to Phase 2, we build the retweet network as in Section 2, i.e. a social network where the nodes are the twitterers, who represent the agents in our ABM. That social network is described by its weighted adjacency matrix M, whose generic element m_{ij} is the number of times that the twitterer i has been retweeted by the twitterer j.

For the computation of influence coefficients (Phase 2), we assume that the set of twitterers is the set of agents. If the order of elements in the two sets is not the same, a permutation on either set is needed before applying the calibration method. Since the weights of the edges in the retweet network represent the influence exerted by the twitterers (i.e., the agents) on one another, we obtain the matrix of influence coefficients in the ABM by the simple equation $S = M$, if no normalization is needed. However, it is to be noted that the adjacency matrix elements $m_{ij} \in \mathbb{N}$. A further stage is therefore needed if the influence coefficients belong to a different domain. We can just examine the cases where the domain of the influence coefficients is either $\mathbb{N}$ or $\mathbb{R}^+$. We introduce the set V of values acceptable for the influence coefficients. For example, in [5], we have $V = [0, 1]$. Let's consider first the case where $V \subset \mathbb{R}^+$. Assuming that V always includes the value 0 to describe the case of no influence, defining v as the upper bound of V and $m = \max_{i,j} \quad m_{ij}$ we can map the values of the adjacency matrix M into influence coefficients by the following linear scaling

$$s_{ij} = \frac{v}{m} m_{ij} \qquad i, j = 1, 2, \ldots, n. \tag{2}$$

If $V \subset \mathbb{N}$, we have instead either a contraction mapping or a dilation mapping according to whether we have $m > v$ or $m < v$ respectively.

For the applicability of the method, we recall the following categories adopted to classify agent-based models in the survey [1]:

- opinion domain;
- interaction direction;
- interacting agents;
- updating equation;
- updating frequency;
- utility function.

Opinion domain. Though an opinion is intrinsically a qualitative and potentially multi-faceted feature, we have to describe it by a numeric variable and choose the domain where that variable can lie. The following three domains have been surveyed in [1]:

- discrete;
- continuous over a bounded interval;
- continuous over $\mathbb{R}$.

In the discrete case, the agent may choose its opinion within a limited set. Discretization lends itself well to represent a qualitative feature through a proper mapping. For example, if we consider the simplest discrete case where we have a binary opinion variable, the two values may represent respectively a positive versus a negative opinion.

If the domain is instead a continuous but bounded interval, a common choice is the $[0, 1]$ interval.

The method proposed here is generally applicable to any opinion domain, since the applicability condition concerns the values of the influence coefficients. However, the closure property may impose some conditions on the influence coefficients. By closure property we mean that the set V is closed with respect to the application of the state updating equation. For example, in [5], where $V = [0, 1]$ and the agents belong to one of c classes (n_i representing the number of agents in class i), and similarly in [4] for the pairwise case, the closure property requires that

$$s_{ii}(n_i - 1) + \sum_{\substack{k=1 \\ k \neq i}}^{c} s_{ik} n_k \leq 1 \qquad i = 1, 2, \ldots, \sum_{j=1}^{c} n_j \tag{3}$$

Interaction direction. That feature considers which way the agents influence each other. again, we refer to the classification established in [1]. In a bilateral interaction, any two agents always influence each other mutually. We have instead a unilateral influence when an agent may influence another agent without being influenced by it. Even in the bilateral case, the influence may not be perfectly symmetrical, since we can have different weights in the opinion updating equations, signalling that the impact of agent X on agent Y is different from that in the reverse direction.

The method is applicable when the interaction direction is bilateral, non symmetric since retweeting may take place in either direction.

Interacting agents. As to the number of agents that interact at each step, the classification adopted in [1] considered the following three:

- pairwise;
- any-to-any;
- closest neighbours.

In the pairwise case, just two agents interact at any single opinion updating round. Since the pairs may change at each round, any agent may interact (influencing or being influenced by) with any other agent in the long run. The any-to-any interaction case is obviously the case where all the agents change their opinion at each time step, since they are influenced by the opinions of the other agents at the previous time step. Finally, in the closest neighbour case, any agent interacts just the closest agents (where the notion of closest involves the use of some distance metric, which is natural in a social network).

In our method, since the interaction is associated with retweeting, and any twitterer can retweet any other twitterer, the calibration we propose applies to all three categories. For the case of pairwise interaction, the weights will be employed in pairs at each round, though they have been estimated considering the embedding social network as a whole. As to the closest neighbour case, if the distance employed for agents is that established on the embedding retweet network, there is actually no difference, since the closest neighbour are those retweeting and actually the only ones exhibiting a non-zero weight.

Updating equation. The updating equation is the function that relates the opinion of an agent to the opinions of the other agents. Here, the classification proposed in [1] is quit simple, considering a linear vs a nonlinear model.

Here we have described the method considering just linear updating equations so far, but it could also be applied to non-linear updating equations, though requiring a more complex mapping from M to V.

Updating frequency. This parameter can be considered as the speed of the opinion formation process. The survey in [1] considers periodic and aperiodic updating. In the periodic setting, each time step involves a change of opinions for all the agents. On the other hand, we fall in the aperiodic case when just a couple of agents changes their opinion at each time step (and it is not known in advance when their turn comes again), or opinions are updated just after a triggering event, or opinion change takes place for a random selection of agents at each time.

Our calibration method is agnostic to the choice of updating frequency, since it can be applied as many times as desired.

Utility function. Again, our calibration method is agnostic to the choice of utility function, as long as that does not impact on the influence coefficients.

4. The dataset

As recalled in the Introduction, for our calibration method we have chosen an application example concerning the influence of people' opinions about wind power. In this section, we describe our dataset and the procedure we have adopted to build it.

The procedure goes along the following four phases:

1. tweet retrieval;

2. duplicate removal;

3. selection of relevant tweets;

4. retweet network building.

In order to retrieve all relevant tweets, we exploit the Twitter API by searching for all the tweets containing either of the following word combinations:

- *wind* AND *power*;

- *wind* AND *energy*.

We deem those words to be fairly representative of the tweets associated to wind power for our calibration demo. Of course, if our aim were to go beyond a demo, we could devise an ampler set of words to obtain a dataset as exhaustive as possible. In this paper, we consider the tweets posted in the week ending on December 9, 2019. Again, the calibration could be made more accurate by extending the analysis horizon over several weeks or even longer periods.

Since many tweets contain both the above combinations, our basket after the retrieval phase may contain duplicates. In the second phase of our procedure, we remove all duplicates. The unique tweets after Phase 2 are 36539.

We must however recognize the possibility of including tweets that, despite binding the word *wind* to *power* or *energy* terms, are not relevant to our actual theme, i.e. the use of wind to get electrical power. Examples of such tweets are shown in Figure 2.

We need to eliminate as many as possible non-relevant tweets. In order to arrive at a set of relevant tweets, we employ a semi-automatic procedure, based on hashtags and the co-occurrence principle. Our procedure, which makes up Phase 3 of the overall procedure mentioned above, goes through the following steps:

1. select the k most frequent hashtags in the dataset of interest;

2. identify the hashtags that are surely relevant with out topic and form a group with them (say Group X);

3. form a group with all the other hashtags (say Group Y);

4. examine all tweets containing Group Y hashtags but not Group X ones, and move their hashtags to Group X if those tweets are relevant;

5. assign to Group X all the hashtags co-occurring with Group X hashtags (this is not done iteratively, but just once for each Group X hashtag).

The last step of the procedure is equivalent to:

1. building the network of hashtags, where an edge is drawn between two hashtags if those two hashtags co-occur in at least one tweet;

2. identifying Group X hashtags;

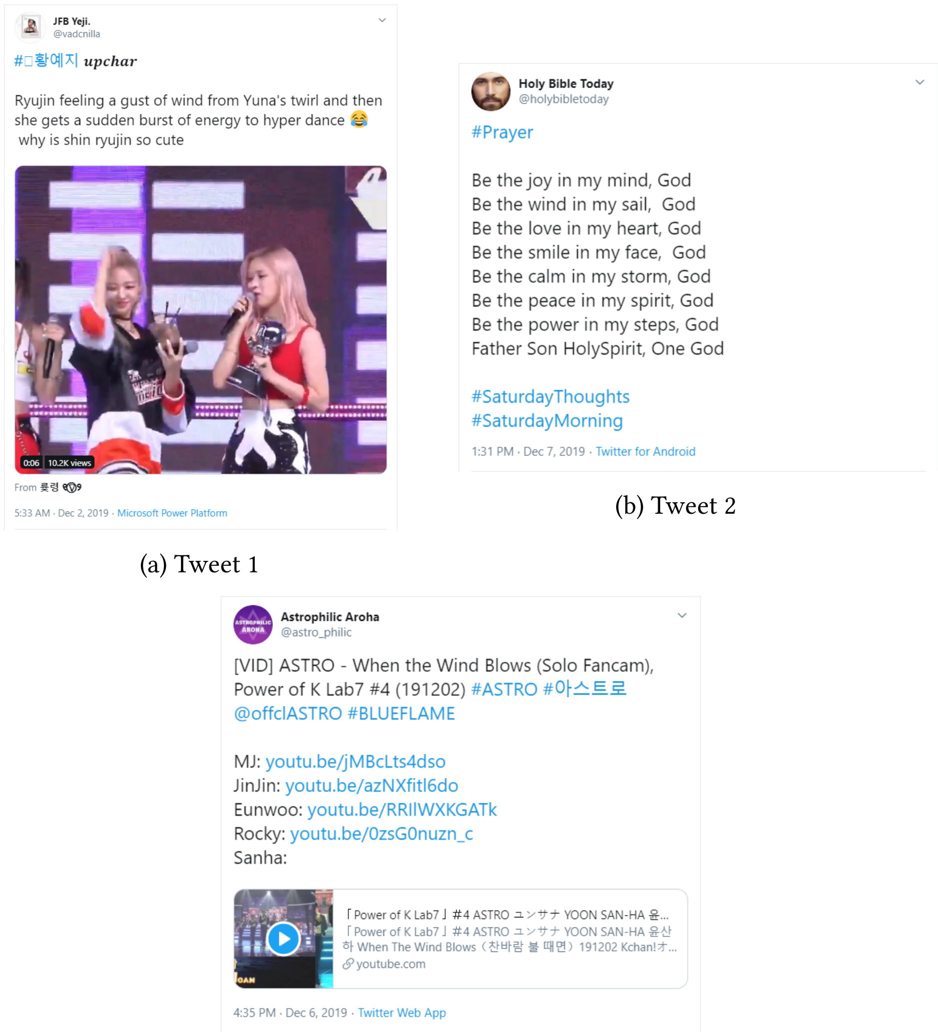

(a) Tweet 1

(b) Tweet 2

(c) Tweet 3

Figure 2: Example of non-relevant tweets.

3. adding neighbours of Group X hashtags to Group X.

A subset of the hashtag network is shown in Figure 3, where the nodes (hashtags) belonging to Group X are highlighted in dark colour.

At this point we have the Group X of relevant hashtags. We can consider a tweet as relevant if it contains any hashtag included in Group X and build the retweet network on the basis of those tweets as described in Section 3. The resulting network is built out of 4739 tweets and is made of 3528 nodes and 3617 edges.

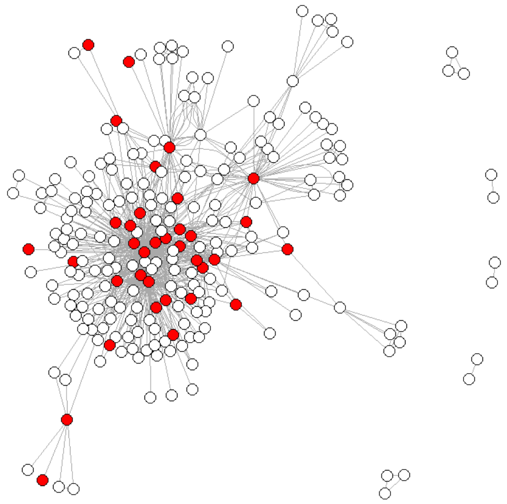

Figure 3: Network of hashtags

5. Experimental calibration results

After proposing the calibration method in Section 3, in this section we apply it to the dataset described in Section 4.

In Figure 4, we show the resulting retweet network. In order to avoid excessive garbling, we have drawn just the nodes with degree larger than one and have arranged it by degree, so that the most central nodes are located in the inner core of the graph. We see that there are just a few very much retweeted twitterers.

Aside from the sheer demonstration of the applicability of our calibration method, we exploit this application to investigate the following research questions (RQ):

1. how frequent is the retweet phenomenon?
2. how widespread is the influence of any single agent upon the community?
3. how heavy is the influence of any single agent on another specific agent?

As to RQ1, we already have a partial answer from Figure 4. There are wide imbalances in the connectivity of individual nodes. The sparsity of the adjacency matrix is 99%; this means that the network is extremely far from a fully-meshed network where every twitterer retweets all his/her fellows. However, this should be better investigated over time, since we considered a single week, and retweeting relationship accrue over time. In fact, though older tweets are often quickly forgotten, most users tend to "live in the present", forgetting or abandoning topics they followed just some hours or days after, the retweeting relationship lasts over time if the retweeting user is a convinced follower/supporter of the retweeted one.

In order to answer RQ2, we show the distribution of nodes by their degree in Figure 5(a). We see that 89.6% of the agents influences just another agent (the corresponding nodes have degree 1), though there is a very small number of agents exerting their influence on a large

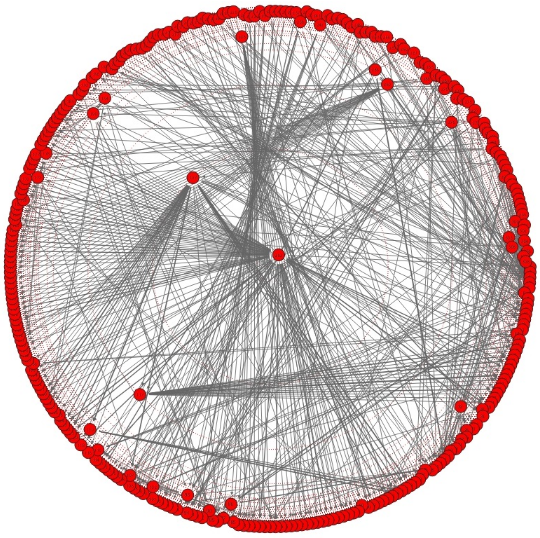

Figure 4: Retweet network arranged by degree centrality

number of other agents, even more than 10. An overall measure of the imbalance in influence is given by the graph centralization measure (see Chapter 5.3.1 of [22]), which takes values in the [0,1] range, with 0 corresponding to a network where all the agents have the same influence, and 1 corresponding to a star network (maximum possible imbalance). The centrality measure we consider is based on the degree (i.e., the number of links incident upon a twitterer in the embeddding retweet network). Mathematically, if g_v is the degree of node v, and v^* is the node exhibiting the highest degree the *graph centralization* is:

$$G := \frac{1}{(n-1)(n-2)} \sum_{i=1}^{n} \left[g_{v^*} - g_{v_i} \right]. \tag{4}$$

In our case, the graph centralization measure achieves an intermediate value, namely 0.4275. Though this could appear lower than expected, given the emergence of a few dominant twitterers, those large imbalances that inflate the sum in Equation (4) are probably countered by the large size of the network (hence, large value of n that enlarges the denominator in Equation (4)).

As to RQ3, in the correspondence between the retweeting network and the agent-based model, we have set the influence coefficients proportional to the edge weights. A measure of the level of influence exerted by twitterers on their fellows is therefore the weight: if some edge is associated to a large weight, that relationship bears a heavy influence. If we look at the distribution of edge weights in Figure 5(b), we see that most twitterers (94.8%) exert just a small influence on others (i.e., the weight of their edges is just 1), though there are a small minority that are retweeted more frequently (even 7 times over the week of observation) and therefore exert a heavier influence.

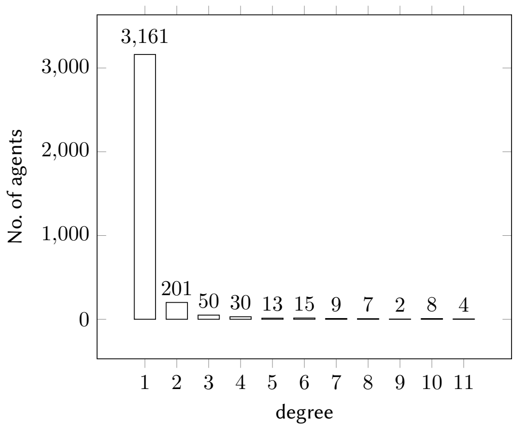

(a) Distribution of nodes by their degree

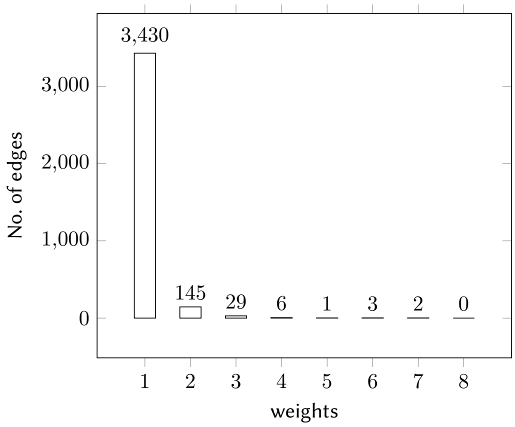

(b) Distribution of edges by their weight

Figure 5: Influence by agents

6. Conclusions

Our paper deals with a critical issue in the development of agent-based models, i.e., setting the parameters that govern the model and allow to apply such models in the real world (what we call the calibration of the model). Our method can be applied whenever the agents act in a social network. Though we showed an example using Twitter data, any social network allowing an opinion reposting mechanism can be used. Also, the class of agent-based models to which it can be applied is fairly large. The only significant limitation of the current approach is the linear form of the equations that govern the interaction among agents. The possible limitations on the correspondence between the ranges of influence coefficients in the agent-based model on one side and edge weights in the social network on the other side could be addressed easily by suitable operations, e.g., by translation and rescaling to realign the two ranges or possibly by nonlinear transformations.

We therefore envisage this calibration model to fill the gap between the theoretical analysis of an agent-based model and its applicability in a real world context. We wish to address its limitations as to the applicability to nonlinear models and different parameter ranges in our future work.

References

[1] L. Mastroeni, P. Vellucci, M. Naldi, Agent-based models for opinion formation: a bibliographic survey, IEEE Access (2019) 1–12.

[2] P. Vellucci, M. Zanella, Microscopic modeling and analysis of collective decision making: equality bias leads suboptimal solutions, Annali dell'Università di Ferrara (2017).

[3] Y. Shang, An agent based model for opinion dynamics with random confidence threshold, Communications in Nonlinear Science and Numerical Simulation 19 (2014) 3766–3777.

[4] S. Monica, F. Bergenti, An analytic study of opinion dynamics in multi-agent systems, Computers & Mathematics with Applications 73 (2017) 2272 – 2284.

[5] L. Mastroeni, M. Naldi, P. Vellucci, Opinion dynamics in multi-agent systems under proportional updating and any-to-any influence, in: M. Paolucci et al. (eds.), Advances in Optimization and Decision Science for Society, Services and Enterprises, AIRO Springer Series 3, Springer Nature Switzerland., 2019. URL: https://www.springer.com/gp/book/9783030349592#aboutBook.

[6] E. Dacrema, S. Benati, The mechanics of contentious politics: an agent-based modeling approach, The Journal of Mathematical Sociology 44 (2020) 163–198.

[7] L. Mastroeni, M. Naldi, P. Vellucci, Individual competence evolution under equality bias, in: European Modelling Symposium (EMS), Manchester, 2017.

[8] T. Lux, Estimation of an agent-based model of investor sentiment formation in financial markets, Journal of Economic Dynamics and Control 36 (2012) 1284–1302.

[9] L. Mastroeni, M. Naldi, P. Vellucci, An agent-based model on scale-free networks for personal finance decisions, in: Proceedings of the 20th Workshop "from Objects to Agents", 2019.

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

[11] N. E. Friedkin, E. C. Johnsen, Social influence networks and opinion change, Advances in Group Processes 16 (1999) 1–29.

[12] M. Mäs, A. Flache, Differentiation without distancing. Explaining bi-polarization of opinions without negative influence, PLOS ONE 8 (2013) 1–17.

[13] P. Jia, A. MirTabatabaei, N. E. Friedkin, F. Bullo, Opinion dynamics and the evolution of social power in influence networks, SIAM Review 57 (2015) 367–397.

[14] D. Krackhardt, Cognitive social structures, Social Networks 9 (1987) 109 – 134.

[15] K. R. Varshney, Bounded confidence opinion dynamics in a social network of Bayesian decision makers, IEEE Journal of Selected Topics in Signal Processing 8 (2014) 576–585. doi:10.1109/JSTSP.2014.2309945.

[16] L. Li, A. Scaglione, A. Swami, Q. Zhao, Consensus, polarization and clustering of opinions in social networks, IEEE Journal on Selected Areas in Communications 31 (2013) 1072–1083.

[17] A. Madan, M. Cebrian, S. Moturu, K. Farrahi, A. Pentland, Sensing the "health state" of a community, IEEE Pervasive Computing 11 (2012) 36–45.
[18] T. O'Reilly, S. Milstein, The twitter book, O'Reilly Media, Inc., 2011.
[19] S. Kumar, F. Morstatter, H. Liu, Twitter data analytics, Springer, 2014.
[20] E. Bakshy, J. M. Hofman, W. A. Mason, D. J. Watts, Everyone's an influencer: quantifying influence on twitter, in: Proceedings of the fourth ACM international conference on Web search and data mining, ACM, 2011, pp. 65–74.
[21] J. Gentry, Package 'twitteR', CRAN repository., 2016. URL: https://cran.r-project.org/web/packages/twitteR/twitteR.pdf.
[22] S. Wasserman, K. Faust, et al., Social network analysis: Methods and applications, volume 8, Cambridge university press, 1994.