Collective Decisions in Multi-Agent Systems

Frank Schweitzer

Chair of Systems Design, ETH Zurich, Kreuzplatz 5, 8032 Zurich, Switzerland
fschweitzer@ethz.ch

1 Decision making and herding effects

Decision making, in a simple sense, means a selection among alternatives. It is one of the fundamental processes in economics but also in social systems. If these systems consist of many interacting elements – which we call *agents* here – the system dynamics may be described on two different levels: the *microscopic* level, where the decisions of the individual agents occur and the *macroscopic* level where a certain collective behavior can be observed.

From the utilitaristic perspective of an individual agent, a decision should be made in a way that the result increases her private utility. This approach is often related to the *rational agent* model, one of the standard paradigms of neoclassical economic theory. It assumes that the agent is able to calculate her utility function based on (i) the *complete knowledge* of all possible actions and their outcomes, and (ii) the *common knowledge assumption*, i.e. that the agent knows that all other agents know exactly what he/she knows and are equally rational.

This implicitely requires an infinitely fast, loss-free and error-free dissemination of information in the whole system. A more realistic assumption would be based on the *bounded rationality* of agents, where decisions are not taken upon complete a priori information, but on incomplete, limited knowledge distributed with finite velocity. This however would require to model the information flow between the agents explicitely. A possibile approach to this problem is given by the *spatio-temporal communication field* [5]. It models the exchange of information in a spatially extended system with finite velocity, considering also the heterogeneous creation of information and memory effects due to the finite lifetime of information.

Based on incomplete information, how does an agent make her decision on a particular subject? A "simple" utility maximization strategy may fail because in many social situations, for example in public votes, the private utility cannot be easily quantified, i.e., agents do not exactly know about it.

Moreover, in multidimensional problems decisions often lead to ambiguous solutions which do not satisfy all needs. So, agents have to involve supplemented strategies to make their decisions.

In order to reduce the risk of making the wrong decision, it seems to be appropriate just to copy the decisions of others. Such an *imitation* strategy is widely found in biology, but also in cultural evolution. Different species including humans imitate the behavior of others of their species to become successful or just to adapt to an existing community.When agents only observe the decision of other agents and tend to imitate them, *without* complete information about the possible consequences of their decisions, this is commonly denoted as *herding behavior*. It plays a considerable role in economic systems, in particular in financial markets, but also in human and biological systems where *panic* can be observed.

Herding behavior is based on a non-linear feedback between the decisions of agents, where sometimes different kind of information is involved. In this short contribution, we will entirely focus on the role of such *non-linear feedbacks* on the decision of agents, while leaving out the possible influence of some private utility maximization. This restriction allows us to pass by most of the problems in defining social utilities; it further makes more clear to what extent the outcomes of decisions is already determined by these feedback processes.

When focussing on collective decisions, we are interested in the *aggregated* outcome of many individual decisions. As is known from a large body of research in the field of complex systems, the interaction of agents on the *microscopic* level – mentioned above – may lead to the emergence of new systems qualities on the *macroscopic* scale. While these emergent properties cannot be reduced to the dynamics of the agents, it is also important to notice that many of the individual agent features do not play a crucial role in establishing the macrofeatures because they are simply averaged out. So, it seems possible to derive a collective dynamics that sufficiently describes the aggregated outcome without depending on all details of the microscopic agent configurations.[1]

Suitable examples of a collective decision processes are public polls [2]. In many cases, these are based on *binary decisions*, i.e. in favor or against a given proposal, either for candidate A or B, etc. So, there are only two alternatives (or opinions), $\{0, 1\}$. Real examples from the year of 2005 include the French vote for/against the European constitution on May 29 – the result was 45% in favor and 55% against this proposal, or the Swiss vote for/against the Schengen treaty on June 5 – the result was 54.6% in favor and 45.6% against this proposal. Other well known examples are the two most recent US presidential elections, where voters had to decide between two candidates. Common to most of these examples, there is no simple utility maximization

[1] I am aware that I am selling here the methodological approach of statistical physics to social scientists. After all, this was a very *fruitful* research program with a lot of impact on the theory of complex systems.

involved. Further, the winning opinion in these collective decisions was always hard to predict, as it was in many cases about 50/50.

If the collective decision process results in the exclusive dominance of only *one* opinion, the system has reached *consensus*. If more than one opinion remain, the system state is characterized by *coexistence* - in the binary case of two (opposite) opinions. Both opinions most likely have a different share x (in percent), thus, if an opinon – in the binary case – has a share of $x > 0.5$, it is the opinion of the *majority*, and with $x < 0.5$ it is the opinion of the *minority*.

Our aim is now to find a *minimalistic agent model* that may describe the generic dynamics of collective decision processes as mentioned above. This model shall focus on the non-linear interaction of the agents rather than making assumptions about their individual utility maximization. In particular, we concentrate on the role of local and neighborhood effects on the aggregated outcome. Our aim is to derive a dynamics for a macroscopic parameter, such as the share x of a particular opinon, not to predict individual decisions.

2 Nonlinear Voter Models

Let us assume a population of agents ($i = 1, ..., N$) where each agent i is characterized by two individual variables: (i) her spatial position i (for simplicity just consecutively numbered) and (ii) her "opinion" $\theta_i(t)$ which is either 0 or 1. In this setting, "decision" simply means to keep or change opinion $\theta_i(t)$ in the next time step, i.e.

$$\theta_i(t+1) = \begin{cases} \theta_i(t) & \text{keep} \\ 1 - \theta_i(t) & \text{change} \end{cases} \tag{1}$$

The rate (number of events per time unit) to change the opinion shall be denoted by $w(1 - \theta_i|\theta_i)$. It remains to specify what the decision of agent i depends on. In social systems, this may depend on the many (internal or external) interdepencencies of an agent community that push or pull the individual decision into a certain direction, such as peer pressure or external influences. The *social impact theory* [3] that intends to describe the transition from "private attitude to public opinion" has covered some these collective effects in a way that can be also formalized within a physical approach.

Here, we follow a much simpler modeling approach by just assuming that the rate to change the opinion depends on other agents in the neighborhood in a nonlinear manner:

$$w(1 - \theta_i|\theta_i) = \kappa(f)\; f_i^{1-\theta_i} \tag{2}$$

where $0 \leq f_i^{1-\theta_i} \leq 1$ denotes the frequency of agents with *opposite* opinions in "neighborhood" of agent i and $\kappa(f)$ means a nonlinear response to the frequency of other opinions. This model class of frequency dependent processes is known as *voter models* [4, 6]. In the most simple case, $\kappa(f) = 1$, the transition rate towards the opposite opinion is simply proportional to the frequency of agents with that particular opinion, this is known as the *linear voter model.*

In order to determine f_i, we have to specify the meaning of neigborhood. For the simulations described below, we use a regular grid, where each agent has four nearest neighbors and eight next-nearest neighbors. In a more general setting, the neighbors are defined by the the social network of agent i, i.e., two agents are direct neighbors if there is a link between them. The structure of the social network can then be described by an adjacency matrix $\mathcal{C}$ which contains as entries all the existing links between any two agents. So, there is no principle limitation to set up the dynamics for any kind of networks. Only for visualization purposes, we restrict ourselves to the regular grid, which means a specific form of the matrix $\mathcal{C}$.

We further have to specify the nonlinear response function $\kappa(f)$, which gives a weight to the influence of agents with opposite opinions on the decision process of agent i. Figure 1 shows some possible cases.

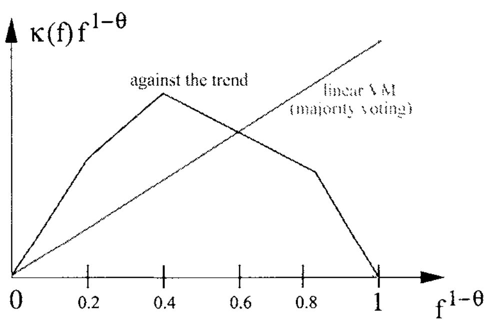

Fig. 1. Different (non)linear depencencies for the transition towards the opposite opinion.

The linear voter model, $\kappa(f) = 1$, where the decision rate of changing towards the opposite opinion directly increases with the frequency of the other opinion, is an example of *majority voting*, i.e., agent i tends to follow the majority of agents in the neighborhood. However, it can be also possible that agent i tends to follow the *minority* in his neighborhood. This means the more agents have the opposite opinions, the less agent i is convinced to follow them. Eventually, there is also the possibility to decide against the trend, i.e. agent i switches to the opposite opinion only as long as it is not the opinion of the majority. Note that various other nonlinear responses to the frequency of opposite opinion are possible.

In the following, we investigate some special cases of the nonlinear response function regarding their impact on the collective decision outcome. The linear voter model, $\kappa(f) = 1$, is used here as a reference case. Our computer simulations always start from a random initial distribution of opinions, i.e. agents get randomly assigned either a black or a white label. As the simulations in Figure 2 show, the individual decisions of agents result in the formation of spatial domains of like-minded agents. This is based on the herding effect mentioned above. Any configuration with more than two agents in the neighborhood having the same opinion results in a positive feedback towards the decisions of the other agents to adopt that "majority" opinion. On medium time scales, we observe the *local coordination* of decisions visible in the emergence of domains, and the *coexistence* of the two different opinions. On large time scales, however, one of these opinions takes over and the collective decision process converges to the *consensus* state.

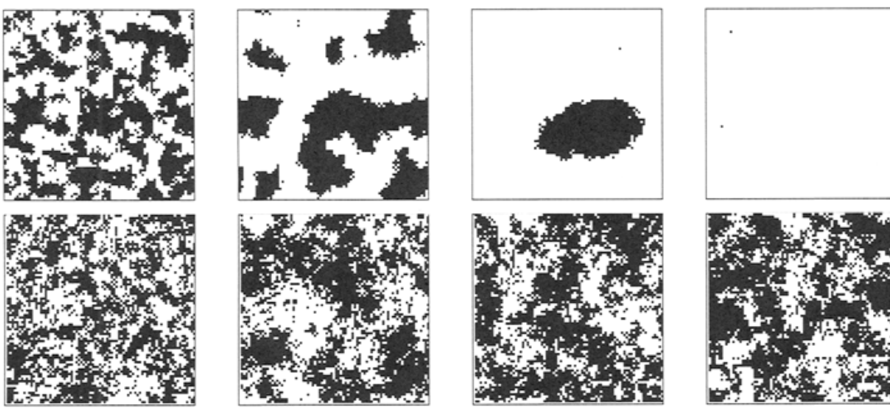

Fig. 2. Spatial distribution of opinions in the case of the linear voter model (top) and a nonlinear voter model (bottom) for different time steps: t = 10^1, 10^2, 10^3, 10^4. [4]

Such a collective dynamics is nice, but also boring because the time until the system reaches consensus, r, is the only interesting feature. Many investigations on the voter model have concentrated on *r(N),* i.e., how the consensus time depends on the system size [6]. The results show that *r* for the regular two-dimensional lattice scales as r ~ iVTniV, whereas for regular lattices with dimension *d* > 2 *r* ~ TV holds. This scaling is also observed for small-world networks.

More interesting are scenarios which would lead to the *coexistence* of the two opinions even on large time scales. This can be obtained by choosing nonlinear response functions *K,(f)* similar to the case of *minority voting,* shown in Figure 1. It means that every local majority trend is immediately teared down, and consensus is never reached. However, this case is boring again, because such a nonlinear feedback alone just reenforces the random equal distribution of opinions. While there is an ongoing dynamics, it does not allow for spatial coordination of decisions. What we are really intersted in, is the *coordination of decisions* together with *non-stationary coexistence* of opinions

– even on large time scales. This can be indeed observed for certain choices of the nonlinear response function $\kappa(f)$, as Figure 2 demonstrates. It should be noted that the spatial domains of opposite opinions continue to coexist while slightly changing in size and shape over time.

A closer inspection of the problem allowed us to derive a phase diagram in the parameter space of $\kappa(f)$ that distinguishes settings leading to random (trivial) coexistence from those leading to nonstationary coexistence with coordination of decisions (domains formation) and from those leading to coordination of decisions on medium time scales, but only to consensus on large time scales.

3 Conclusions

The nonlinear voter model used here as a framework for modeling collective decision processes follows the KISS principle, as it is simple and stupid enough to allow also for analytical investigations. This is, however, not the end of the story, because we extended this model gradually towards more realistic scenarios. A major step, not discussed in this short paper involves the *heterogeneity* of the agents, i.e. agents may have a different nonlinear response functions, $\kappa_i(f)$ dependent on individual attitudes. A variant of this heterogeneity includes memory effects, i.e. the past experiences of agents in their local neighborhood are taken into account. Further, we have considered ageing effects which affect the rate at which agents make a decision, and have also included dependency on the second-nearest neighbors [1].

References

[1] Behera, L.; Schweitzer, F. (2003). On spatial consensus formation: Is the Sznajd model different from a voter model? *International Journal of Modern Physics C* **14(10)**, 1331–1354.

[2] Galam, S.; Zucker, J.-D. (2000). From Individual Choice to Group Decision Making. *Physica A* **287(3-4)**, 644–659.

[3] Nowak, A.; Szamrej, J.; Latané, B. (1990). From Private Attitude to Public Opinion: A Dynamic Theory of Social Impact. *Psychological Review* **97**, 362–376.

[4] Schweitzer, F.; Behera, L.; Mühlenbein, H. (2003). Frequency dependent invasion in a spatial environment. `arXiv:cond-mat/0307742`.

[5] Schweitzer, F.; Zimmermann, J.; Mühlenbein, H. (2002). Coordination of decisions in a spatial agent model. *Physica A* **303(1-2)**, 189–216.

[6] Suchecki, K.; Eguiluz, V. M.; San Miguel, M. (2005). Voter model dynamics in complex networks: Role of dimensionality, disorder and degree distribution. *Physical Review E* **72**, 036132.