Diemo Urbig and Robin Malitz (2007): Drifting to more extreme but balanced attitudes: Multidimensional attitudes and selective exposure. *to be presented at Fourth Conference of the European Social Simulation Association (ESSA)*, September 10-14, Toulouse, France.

Drifting to more extreme but balanced attitudes: Multidimensional attitudes and selective exposure

Diemo Urbig[1] and Robin Malitz[2]

[1] Humboldt-Universität zu Berlin, School of Management and Economics and Department of Computer Science, diemo@urbig.org
[2] Humboldt-Universität zu Berlin, Computer- und Medienservice, malitzro@cms.hu-berlin.de

Abstract. We present a model of opinion and attitude dynamics that incorporates multi-dimensionality of a single attitude, i.e. attitudes are composed of evaluations of several features. We include selective exposure that is modeled as a hierarchical version of the bounded confidence model. Individuals adapt their feature evaluations if their attitudes and the respective feature evaluation are sufficiently close. We show that this model can explain a balancing principle, such that features associated with the same object tend to be evaluated similarly. We also show that such individuals can get more extreme despite they might even interact only with agents that have less extreme attitudes.

1 Introduction

An 'attitude is a psychological tendency that is expressed by evaluating a particular entity with some degree of favor or disfavor' (Eagle & Chaiken, 1993). If the 'entity' is a behavior, then an attitude can be considered as an individual's tendency to engage in a specific behavior. Thus, attitudes have been at the heart of models that predict behavior of people. Behaviors of interest are, e.g., buying decisions by consumers, voting behavior, or the adoption of new technologies by firms. Hence, there are many models that describe and explore processes of attitude formation. While some models focus on how single individuals form their opinions in response to specific stimuli, other models focus on social processes of opinion formation. This article focuses on how attitudes evolve in groups of interacting individuals.

Simple models on interacting individuals and resulting changes in attitudes were studied by Festinger (1950), French (1956), Latané et al. (1994), Deffuant et al. (2000), Hegselmann and Krause (2002), Galam (2002), and others. A basic assumption is that interacting and especially communicating individuals mutually affect their attitudes. Regarding communicating individuals there has also been some work on explaining why individuals do not tell others their true attitudes (e.g., Kuran, 1987). Thus, there might be a gap between an individual's attitude, as the mental state, and his expressed opinion as element of communication. As Urbig (2003) argues such differences might also be caused by

limitations inherent in languages or mechanisms used for communication. Because we are not concerned with biases of the attitude–opinion relation we use the terms opinion and attitude as synonyms throughout this text. Furthermore, these communication processes are moderated by selective exposure, which describes a mechanism where people tend to maintain their beliefs by exposing themselves to information supporting their beliefs.

Models of attitude and opinion dynamics as introduced by Nowak et al. (1990), Hegselmann and Krause (2002), and Deffuant et al. (2000) represent some of the more simplistic approaches; but they allow for a clear analysis and understanding of the fundamental dynamics. A more complex model of consumer behavior involving attitudes has been introduced by Jager (2000). From my point of view, taking the simple models as starting point and increasing complexity step by step enables us to keep track of the sources of complexity and of specific dynamics and thus allow to include only those elements that have large effects (see Hegselmann, 2004). Most simple models of opinion dynamics in groups of interacting individuals consider an individual's opinion as a one-dimensional value, either binary as good or bad opinion, or with a finite number of stages between the extremes, or as a continuum between extreme values (see Hegselmann & Krause, 2002, Weisbuch et al., 2002). Sometimes such models also include uncertainty about the opinion (see for example Amblard et al., 2003). The idea of selective exposure is included in these models as bounded confidence or relative agreement, i.e. individuals only others if their attitudes towards the object under consideration are sufficiently close.

One-dimensional approaches to attitudes and opinions are beneficial for understanding basic dynamics, but it is questionable whether they can capture the complex dynamics displayed by opinion dynamics in real settings. For instance, when analyzing the diffusion of innovative products or the acceptance of complex programs of political parties it is somewhat heroic to assume that people either agree on all aspects or on no aspect, as it would be implied by one-dimensional approaches. According to consumer theory and psychological attitude theory a one-dimensional approach seems to be inappropriate. According to this research stream an attitude is composed of different *impressions* that themselves are composed of two elements: *beliefs* (also called *cognitions*) about the presence of some attributes and *evaluations* of these attributes (Ajzen, 1991). In fact, we end up with a qualitative distinction between knowledge and evaluations and a multidimensionality regarding a set of attributes related to an entity. This multidimensionality is at the core of most of the reliable approaches to attitude measurement (Ajzen, 2002). While the distinction between beliefs and evaluations might be questioned, even competing approaches to measurement of attitudes keep the multidimensional nature (see Trommsdorff, 1998). Measuring attitudes gets important when linking simulations to real measured data. Altogether, multidimensionality, but also – to a smaller extend – the distinction between evaluations and cognitions seem to be very general though relevant concepts. Although such a compositional approach to attitudes finally results in a one-dimensional attitude, we will see that an explicit modeling of the different

dimensions enables a modeling of a more differentiated and from our point of view more realistic communication process.

There are several attempts to include multi-dimensionality into simple models of opinion dynamics. But unfortunately the multi-dimensionality is then interpreted as a set of opinions instead of several elements of the same opinion or attitude. Deffuant et al. (2000), Weisbuch et al. (2002), as well as Fortunato et al. (2005) consider a vectors of binary values as a set of opinions toward different objects. Lorenz (2003) investigates a model of n-dimensional vectors of real values, where the real values represent the amount of money allocated to n projects. In Lorenz's model agents interact and adapt their opinions if the Euclidian distance of these opinions is smaller than a constant, which is a possible extension of the Bounded confidence concept by Hegselmann and Krause (2002) and similarly by Fortunato et al. (2005). These models assume an all-or-nothing strategy for adapting to other's opinions, which means that either all dimensions are adapted or no dimension. This is not plausible from my point of view because individuals might be willing to adapt in one dimension while at the same time they strongly resist any adaptation in another dimension.

Furthermore, for determining the condition for an effective interaction agents evaluate their whole vectors of opinions and not only their own but also the whole vector of their current communication partner. From my point of view this extensive comparison of all elements is a strong weakness of these models because it requires an extensive interaction even with those partners that differ extremely; one needs this extensive comparison just to figure out that they actually differ to such a large degree. It seems more plausible that there is some comparable aggregate or placeholder (aggregated evaluation or signal for group membership in cultural models) that finally determines if individuals get involved into a deeper discussion or reject any discussion and therefore get away being unaffected by each other without having had a deep interaction on the mutual differences.

In this article we want to incrementally develop a model of dynamics of structured attitudes that gets closer to psychological attitude theory, i.e. to the model of Fishbein and Ajzen (Ajzen, 1991). Still we are aware that this will only be an inspiration by psychological theory but not an implementation of the theory on attitudes, given that a consistent theory exists at all. Our general approach, which can be labeled as *hierarchical selective exposure* or *hierarchical bounded confidence* approach, can be described as follows:

> *Having first just a general interaction about the attitude regarding an object under consideration an interaction continues only in case of a sufficient closeness. If they continue they may talk about the different attributes related to the object. At this level they may again be bounded in their predisposition to consider their peer's opinion seriously.*

All together, there are different interpretations of vectors of opinions, but as far as we know there are only two interpretations that refer to structured attitudes or opinions, where the attitude is an aggregate of multiple dimensions that might be adapted independently, namely Urbig and Malitz (2005) and Deffuant and Huet (2007). The first paper is on an agent-based simulation of a preliminary

version of the model presented here. The latter article introduces a model that assumes that a 'filter' selects only important features with a higher threshold of importance when the attitude about the feature is incongruent with the global attitude about the object. Furthermore, individuals transmit only features that are congruent with their global attitude. It turns out that in such cases the sequence of perceiving information is highly significant for the overall evolution of attitudes. However, this model deviates in it basic structure significantly from the idea of continuous opinion dynamics.

2 Model

Let $a_{i,t} \in \mathbb{R}_{[0,1]}$ be the attitude of agent $i \in \{1..N\}$ regarding a focal object at time t. The attitude is composed of a set of impressions $e_{i,k,t} \in \mathbb{R}_{[0,1]}$, which are in fact evaluated features comprising an evaluative component and a belief component about the presence of some feature. The attitude is a summary evaluation of all features. We normalize the attitude by dividing it by the number of features.

$$a_{i,t} = \frac{1}{n} \sum_{k=1}^{K} e_{i,k,t} \tag{1}$$

In every step t an agent i interacts with a randomly chosen agent j. If they both adapt their attitudes and thus adapt on at least one feature evaluation, then they are considered as interacting effectively. The condition of interacting effectively on feature k is indicated by $ef_k(i,j) \in \{0,1\}$, which gives 1 for effective interaction and 0 otherwise. This function is symmetric, $ef_k(i,j) = ef_k(j,i)$, such that if one adapts also the other adapts. A convergence parameter $0 < \mu_k \leq 0.5$ describes how much agents adapt . For now let it be specific for the different features k. If $\mu_k = 0.5$ then agents' effective interaction on a feature leads to an agreement in the respective feature evaluations. The smaller μ_k is, the less quickly they converge in their feature evaluations. Given two randomly chosen agents i and j, equation 2 defines how agent i updates her feature evaluations and thus indirectly her attitude. The rule is symmetric for agent j. Note that an agent's attitude only changes if the two agents' attitudes are sufficiently close and if their feature evaluations are sufficiently close.

$$\begin{aligned} e_{i,k,t+1} &= e_{i,k,t} - ef_k(i,j) \cdot \mu_k \cdot (e_{i,k,t} - e_{j,k,t}) \\ &\text{with } ef_k(i,j) = \begin{cases} 1 \; if \; |a_{i,t} - a_{j,t}| < \varepsilon_a & \wedge \quad |e_{i,k,t} - e_{j,k,t}| < \varepsilon_f \\ 0 \; if \; |a_{i,t} - a_{j,t}| \geq \varepsilon_a & \vee \quad |e_{i,k,t} - e_{j,k,t}| \geq \varepsilon_f \end{cases} \end{aligned} \tag{2}$$

3 Get different by getting more alike

A first analysis of this model in an agent-based simulation has shown that there seems to be tendency that agents can get more extreme attitudes although they communicate with agents that have on average less extreme attitudes (Urbig &

Malitz, 2005). However, there was no systematic analysis of this effect. Therefore, we now show that the process described in our model converges, keeps average attitudes and average feature evaluations constant, but attitudes may get more distant through an interaction.

Consider two interacting agents, agent 1 with attitude $a_{1,t}$ and agent 2 with attitude $a_{2,t}$. Without loss of generality let agent 2 have the more optimistic attitude, i.e. $a_{1,t} < a_{2,t}$. Let C contain those features, on which the two agents adapt, i.e. $\forall 1 \leq k \leq K : |e_{1,k,t} - e_{2,k,t}| < \varepsilon_f$.

Because both agents' updating processes are characterized by the same convergence parameter μ_k, the mean of the corresponding impressions keep constant, i.e. $\frac{1}{2}\left(e_{1,k,t+1} + e_{2,k,t+1}\right)$ equals $\frac{1}{2}(e_{1,k,t+1} + \mu_k \cdot (e_{2,k,t} - e_{1,k,t}) + e_{1,k,t+1} + \mu_k \cdot (e_{1,k,t} - e_{2,k,t}))$, which finally translates into $\frac{1}{2}\left(e_{1,k,t} + e_{2,k,t}\right)$. Because the attitude is the average of impressions also the average of all agents' attitudes keeps constant. Although agents can get closer and can get more distant, these dynamics stabilize, which means that they do not move back and forwards for ever.[3]

The distance of two agents' corresponding impressions never increases due to our definition of updating. However, the distance in attitudes can increase. This is not surprising if the updating process was not symmetric for two interacting agents, but in our model it is. One can easily show that the changes in attitudes of two interacting agents have the opposite sign (proof available on request). An increasing distance thus implies that the change in one agents attitude needs to have the opposite sign of the difference between own and the other agent's attitude. The change in attitude of agent i interacting with agent j is given as $\frac{1}{n}\sum_{k \in C} \mu_k \left(e_{j,k,t} - e_{i,k,t}\right)$. Because we assume that $a_{1,t} < a_{2,t}$, we need $0 > \frac{1}{n}\sum_{k \in C} \mu_k \left(e_{2,k,t} - e_{1,k,t}\right)$ for having opposite signs. This implies that the two agents adapt on a feature where the order of their respective impressions is inversed compared to the order of their attitudes. Thus, the move of agent 1 has to have a different sign that the difference between the own and the other agent's attitude.

Our further analysis will focus on a special case. For this we assume that the convergence parameter is equal for all features, i.e. $\mu_k = \mu$. We also look at the case of two features since it is much easier to illustrate than case with more features. For these two instantiations we can derive more precise criteria when agents get more distant while they interact: If the convergence parameter μ is equal for all features, i.e. $\mu_k = \mu$, then agents can only get more extreme if they do not compromise on at least one feature. This is easily show by contradiction:

[3] If agents adapt their feature evaluation towards another agent's evaluation, then they assign the weight $1 - \mu_k$ to their own and μ_k to the other's feature evaluation. Since agents are defined symmetrically, either both interaction partners adapt their evaluations or none of them. Therefore, (1) every agent always assign his own evaluation a weight larger than zero, (2) if an agent assigns a positive weight to another agent's evaluations then this happen vice versa, and (3) all positive weights are always above or equal to $1 - \mu_k$. Hence, according to the stabilization theorem by Lorenz (2005) the process stabilizes in the feature evaluation space. Because the attitude space is just a mapping it stabilizes there as well.

Without loss of generality we assume that the second agent's attitude is larger the first agent's attitude, i.e. $a_{2,t} > a_{1,t}$. Let the convergence parameter be equal for all features, i.e. $\mu_k = \mu$. If all feature evaluations are affected, then we can rewrite the above stated condition as follows $0 < -\sum_{k \in C} (e_{2,k,t} - e_{1,k,t}) = -\sum_{k=1}^{K} (e_{2,k,t} - e_{1,k,t}) = -\sum_{k=1}^{K} e_{2,k,t} + \sum_{k=1}^{K} e_{1,k,t}$. This leads to $a_{1,t} > a_{2,t}$, which contradicts our assumption about the initial relation between these two agents' attitudes. From our analysis so far we can also conclude that if there are only two features, then two agents can only get more distant if they adapt on only one of the two features and this displays an inverse order compared to the attitudes.

4 Drifting to more extreme though balanced attitudes

Getting more distant by compromising is just a possibility and this effect might be wiped out in a population of interacting agents. Therefore we now look at the population level.

For analyzing our model in groups of interacting individuals there are several methods that could be applied. One is agent-based simulation that simulates populations of single agents, i.e. in every step one knows the attitudes of every single agent. We already did agent-based simulation on a similar model as introduced here and have recognized that exploring the full parameter space requires huge computational efforts (see Urbig & Malitz, 2005). Therefore we decided to numerically solve rate equations for analyzing the dynamics (Lorenz, 2006, calls this approach Interactive Markov Chains). Thus, we indirectly assume a population of infinite many agents, i.e. $N = \infty$, and in every step every agent gets a randomly chosen interaction partner and all could adapt in every step. The deviations have an impact, but it is rather small (see for instance for other models of the same class Deffuant & Huet, 2007 and compare Deffuant et al., 2000 with Ben-Naim et al., 2003).

We focus on a two-dimensional model. We define a discrete space of equidistant impressions or feature evaluations. They are labeled from 1 to R. The initial probability mass located in the two-dimensional feature evaluation space at point i, j is given by $F_0(i, j)$ with $1 \leq i, j \leq R$. The two-dimensional feature evaluation space is mapped into the one-dimensional attitude space. Because the feature evaluations are discrete, also the attitude space is discrete with $2 \cdot R - 1$ levels. The probability mass in this space is given by $A_0(i)$ with $1 \leq i \leq 2R - 1$. The mapping from the feature evaluation space onto the attitude space is given by (3). It represents an averaging of the feature evaluations.

$$A_t(i) = \begin{cases} \sum_{x=1}^{i} F_t(x, i - x + 1) & \text{if } i \leq R \\ \sum_{x=i}^{R} F_t(x, R - x + 1) & \text{if } i > R \end{cases} \tag{3}$$

The left part of Figure 1 illustrates the mapping. For a better understanding we plot the isoquants for the attitudes; along these lines agents have the same attitude. The attitude increases the more one gets to the upper right corner.

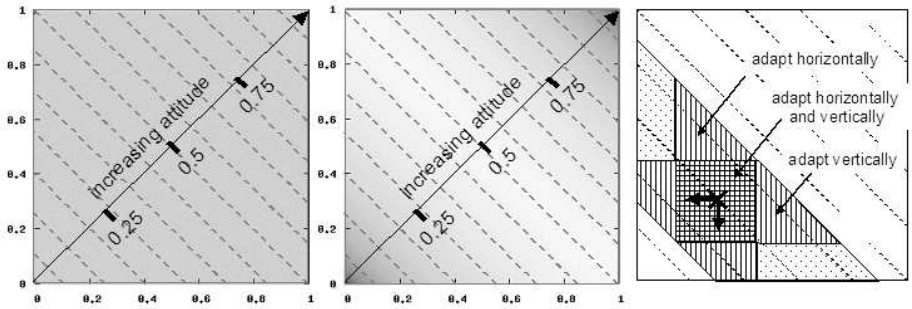

Fig. 1. Left: equally distributed feature evaluations, Center: equally distributed attitudes, Right: an agent's tendency to get more extreme

On can easily see that there are more agents, i.e. a larger probability mass, for intermediate attitudes. The equally distributed feature evaluations translate into a triangular distribution in attitudes.

If we want to compare our model with one-dimensional models that assume initially equally distributed attitudes, then we have to make an equivalent assumption on our initial distribution. We thus manipulate the initial distribution in the feature evaluation space such that at the attitude level agents are equally distributed. The required distribution is illustrated in the middle plot of Figure 1. We need a distribution that emphasizes those areas that imply more extreme attitudes. Our analysis first considers equally distributed feature evaluations to be compatible with traditional models of multidimensional opinions and later compares it with the second case of equally distributed attitudes.

We now define the rate equation for our model. Remember, we assume two dimensions and agreement on all features where agents interact effectively, i.e $K = 2$ and $\mu_k = \mu = 0.5$. We assumed that feature evaluations are discrete. Thus, if two agents compromise they might want to compromise on a value that is exactly in between of two or four of the available values. In such cases we assume that they assign randomly to the respective values. Due to the geometry of the space, only three cases can happen: (1) agents' new feature evaluations can be located at a single cell; (2) they are located on a border between two cells; (3) they are located at a corner, where four cells meet. In the first case we assume that the probability is 1.0 to move into this cell. If it is on a border, then the probability is 0.5 to get into either of two corresponding two cells. For cases when agents would be located at corners the probability is 0.25 to get into one of the four cells. Given this rule let us define two additional functions. Function $d^+(i, j, x_1, y_1, x_2, y_2) \in \{0.25, 0.5, 1.0\}$ gives the probability that an agent maintains feature evaluations according to the cell (x_1, y_1) and moves to cell (i, j) due to an interaction with an agent from cell (x_2, y_2). Furthermore, function $d^-(i, j, x, y) \in \{0.0, 1.0\}$ gives the probability that an agent from cell (i, j) moves somewhere due to an interaction with an agent from cell (x, y). Based on these two functions we now define a rate equation that describes how

the distribution of agents within the space of feature evaluations changes from one step to another.

$$f_{t+1}(i,j) = f_t(i,j) + D_t^+(i,j) - D_t^-(i,j) \qquad \text{with} \tag{4}$$

$$D_t^-(i,j) = \sum_{x=1}^{R}\sum_{y=1}^{R} F_t(x,y) \cdot F_t(i,j) \cdot d^-(i,j,x,y)$$

$$D_t^+(i,j) = \sum_{x_1=1}^{R}\sum_{y_1=1}^{R}\sum_{x_2=1}^{R}\sum_{y_2=1}^{R} F_t(x_1,y_1) \cdot F_t(x_2,y_2) \cdot d^+(i,j,x_1,y_1,x_2,y_2)$$

The sequential calculation of the rate equation is implemented in Java. It either stops after 500 steps or if the sum of all changes is smaller than 10^{-6}. We also calculated arrows that roughly indicate the drifts within the multidimensional space. Figures 2 and 3 illustrate the results of our analysis.

Figure 2 illustrates the distribution in the feature evaluation space for three steps and three different settings. We look at the second setting, where there is in fact no selective exposure at the attitude level. It gets closest to previously analyzed multidimensional dynamics. Because the condition at the attitude level is not restricting the dynamics at all, the dynamics are almost equal to those of two combined one-dimensional dynamics, see for instance the discussion by Fortunato et al. (2005). As Fortunato et al. (2005), we also observe the formation of four equally sized major clusters. Interesting to note, in this scenario agents might get more extreme, but on the population level this effect is wiped out.

If we look at the setting where the selective exposure is restricting the dynamics at both levels, we can indeed observe a drift to feature evaluations, where both features tend to be evaluated similarly and to more extreme attitudes. The difference to the previous case gets obvious in Figure 2. Obviously the selective exposure at the attitude level makes a difference. Right part of Figure 1 illustrates the reason. Agents from the dotted, squared, and striped areas succeed with respect to selective exposure at the attitude level. These restrictions together with the general limits at zero and one create the situation that there is a larger probability to be affected by a more extreme agent at a specific feature evaluation than to be affected by a more moderate agent.

In case that the selective exposure is not restricting the dynamics at the feature evaluation level, we end up with a dynamics that is close to a one-dimensional dynamics. Due to the selective exposure of 0.25 we might expect two clusters, but because the initial distribution is a triangular distribution there is a stronger tendency to form only one large central cluster. However, we observe two small minorities.

While Figure 2 only looks at the distribution in the feature evaluation space and only considers equally distributed initial feature evaluations, Figure 3 plots the evolution of distributions in the attitude space and additionally plots the results for settings, where the initial distribution of attribute evaluations are adjusted such that the initial attitudes are equally distributed.

For the case without a binding restriction at the feature evaluation level we observe a typical evolution of the attitude given that the initial attitudes are

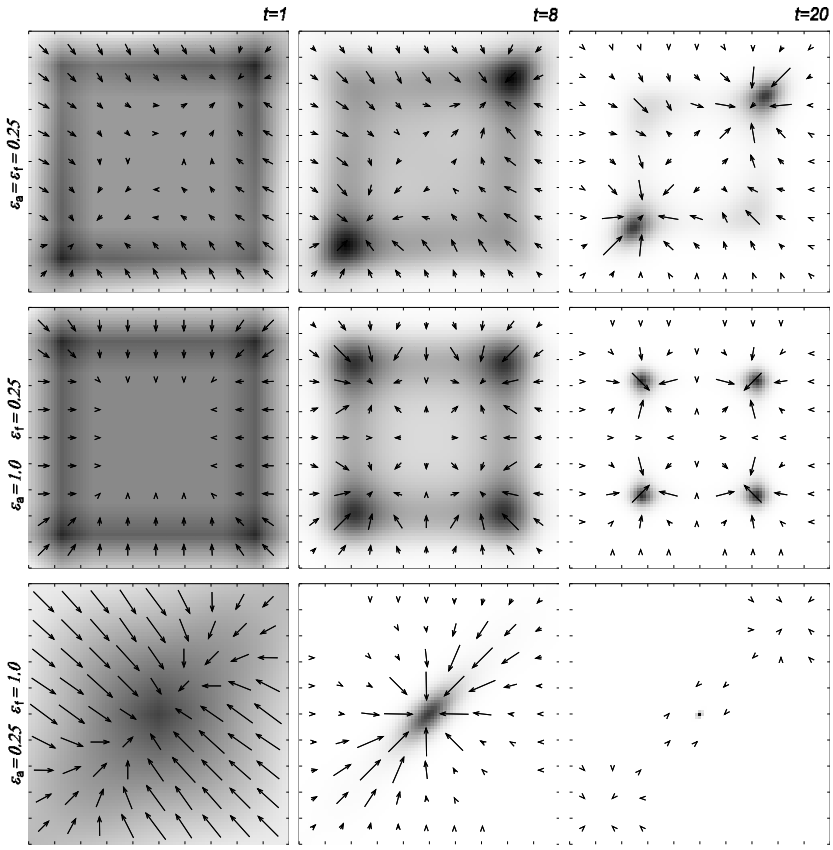

Fig. 2. Evolution of the distribution of attitudes in different scenarios: A: $\varepsilon_a = \varepsilon_f = 0.25$, B: $\varepsilon_a = 1.0$, $\varepsilon_f = 0.25$, and C: $\varepsilon_a = 0.25$, $\varepsilon_f = 1.0$

equally distributed. For the case of no binding restriction on the attitude level we hypothesize that if equally distributed attitudes are assumed, we always end up with equally sized majorities.

In our last analysis (see Figure 4) we take a closer look at the case that the selective exposure is equal on both levels, at the attitude level and at the feature evaluation level, i.e. $\varepsilon_a = \varepsilon_f = \varepsilon$. We vary ε between 0.1 and 0.8. This gives us a much better insight regarding the range of the parameter for selective exposure, where polarization or consensus can be expected. For creating the plot we did some transformations. For each distribution we normalized the values by dividing all cells by the maximum value of all cells. From Ben-Naim et al. (2003) we know that minorities are typically $3 \cdot 10^{-4}$, while majorities are much larger. We use the threshold 0.001 to discriminate between majorities and minorities. Doing this one should keep in mind that majorities can be rather small. And in fact, in some scenarios they get rather small, e.g. the central cluster or majorities

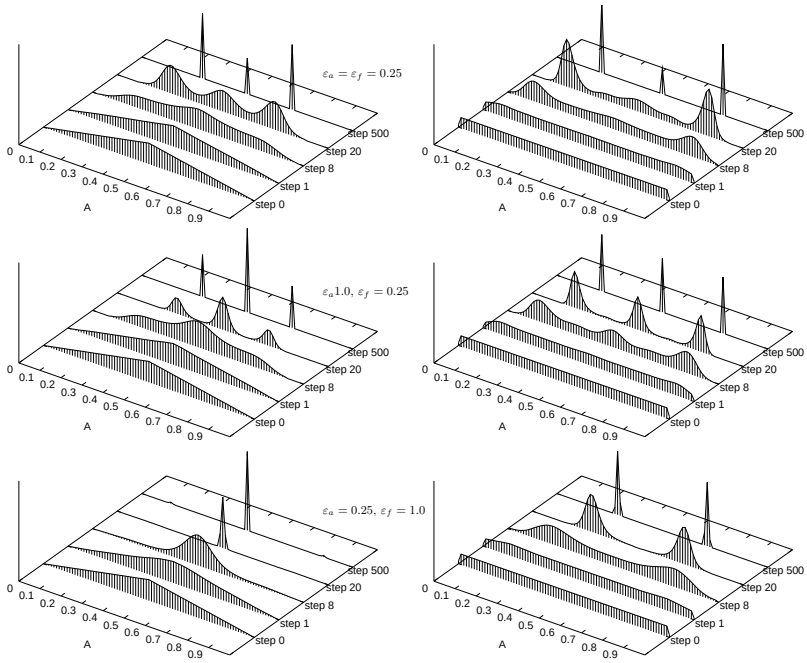

Fig. 3. Evolution of the distribution of attitudes in different scenarios. Left column represents cases where initial evaluations of features are equally distributed and thus attitudes follow a triangular distribution, while second column represents cases with equally distributed initial attitudes.

that are close to the extremes, where already at the beginning of the dynamics there are only small probability masses due to he triangular distribution. Since Figure 4 marks every majority with a dot, these very small majorities look like large majorities. In Figure 4 we observe that for the scenario where selective exposure works on all levels, polarization indeed appears for larger parameter values, $\varepsilon > 0.25$. If we consider equally distributed initial attitudes we even get polarization for some $\varepsilon > 0.35$.

5 Conclusion

In this paper we extended simple models of the social formation of attitude by the aspect of multidimensionality of single attitudes. We introduced the idea of hierarchical selective exposure, which means that when selecting interaction partners and issues first select at a rather general level and in case of selection they step-wise go down in the hierarchy to more specific partners or issues. This model removes a rather heroic assumption of previous models of multidimensional opinion dynamic models, which assume that before selecting a partner both partners have perfect knowledge about all of the other's beliefs and atti-

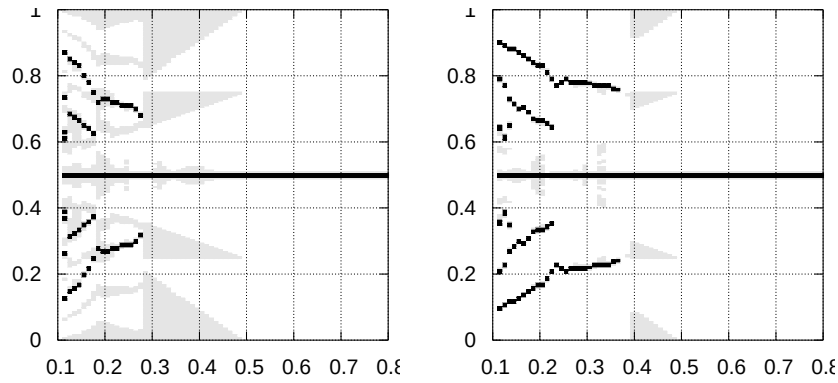

Fig. 4. Bifurcation diagrams for varying $\varepsilon_a = \varepsilon_f$ from 0.1 to 0.8. Left: initially equally distributed feature evaluation, Right: initially equally distributed attitudes

tudes. We have demonstrated that our extension can explain a systematic drift to more extreme attitudes. Polarization that was previously usually not expected for parameter of selective exposure beyond 0.25 can occur systematically for values way beyond this threshold.

For future work one could split the concept of impressions regarding features into beliefs about the presence of features and evaluations of these features. Agents might interact regarding both of these aspects. Agents could affect their beliefs without necessarily affecting their values assigned to specific attributes. Furthermore, applying this idea to more complex interaction rules, e.g. relative agreement model, would probably be of interest.

Acknowledgments. Many thanks to Sylvie Huet and Guillaume Deffuant for their very helpful suggestions regarding the model and its analysis.

References

Ajzen, I. (1991). The theory of planned behavior. *Organizational behavior and Human Decision Processes*, *50*, 179–211.

Ajzen, I. (2002). Attitudes. In R. F. Ballesteros (Ed.), *Encyclopedia of psychological assessment* (Vol. 1, pp. 110–115). London: Sage Publications.

Amblard, F., Weisbuch, G., & Deffuant, G. (2003). *The drift to single extreme appears only beyond a critical connectivity of the social networks.*

Ben-Naim, E., Krapivsky, P., & Redner, S. (2003). Bifurcations and patterns in compromise processes. *Physica D*, *183*(190).

Ben-Naim, E., Krapivsky, P., Vazquez, F., & Redner, S. (2003). Unity and discord in opinion dynamics. *Physica A*, *330*, 99–106.

Deffuant, G., & Huet, S. (2007). Propagation effects of filtering incongruent information. *Journal of Business Research*, *60*(8), 816–825.

Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems*, *3*, 87–98.

Eagle, A. H., & Chaiken, S. (1993). *The psychology of attitude.* Orlando, FL: Harcourt Brace Jovanovich.

Festinger, L. (1950). Informal social communication. *Psychological Review*, *57*, 271–282.

Fortunato, S., Latora, V., Pluchino, A., & Rapisarda, A. (2005). Vector opinion dynamics in a bounded confidence consensus model. *International Journal of Modern Physics C.*

French, J. R. P. (1956). A formal theory of social power. *Psychological Review*, *63*(3), 181–194.

Galam, S. (2002). Minority opinion spreading in random geometry. *The European Physical Journal B*, *25*, 403–406.

Hegselmann, R. (2004). Opinion dynamics: Insights by radically simplifying models. In D. Gilles (Ed.), *Laws and models in science.* London.

Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence models: Analysis, and simulation. *Journal of Artifical Societies and Social Simulation (JASSS)*, *5*(3).

Jager, W. (2000). *Modelling consumer behaviour.* Unpublished doctoral dissertation, University of Groningen, Centre for Environmental and Traffic psychology, Kurt Lewin Instituut, Groningen. (<http://www.ub.rug.nl/eldoc/dis/ppsw/w.jager/>)

Kuran, T. (1987, September). Preference falsification, policy continuity and collective conservatism. *The Economic Journal*, *97*(387), 642–665.

Latané, B., Nowak, A., & Liu, J. H. (1994). Measuring emergent social phenomena: dynamism, polarization, and clustering as order parameters of social systems. *Behavioral Science*, *39*, 1-23.

Lorenz, J. (2003, May). *Multidimensional opinion dynamics when confidence changes.*

Lorenz, J. (2005). A stabilization theorem for continuous opinion dynamics. *Physica A*, *355*(1), 217–223.

Lorenz, J. (2006). Consensus strikes back in the hegselmann-krause model of continuous opinion dynamics under bounded confidence. *Journal of Artificial Societies and Social Simulation*, *9*(1). (<http://jasss.soc.surrey.ac.uk/9/1/8.html>)

Nowak, A., Szamrej, J., & Latané, B. (1990). From private attitude to public opinion: A dynamics theory of social power. *Psychological Review*, *97*(3), 362–376.

Trommsdorff, V. (1998). *Konsumentenverhalten* (third ed.). Stuttgart; Berlin; Köln: Kohlhammer.

Urbig, D. (2003). Attitude dynamics with limited verbalisation capabilities. *Journal of Artificial Societies and Social Simulation (JASSS)*, *6*(1).

Urbig, D., & Malitz, R. (2005, Sep). Dynamics of structured attitudes and opinions. In *Representing social reality. Pre-proceedings of the third conference of the european social simulation association (ESSA '05)* (pp. 206–212).

Weisbuch, G., Deffuant, G., Amblard, F., & Nadal, J. P. (2002). Meet, discuss, and segregate! *Complexity*, *7*(3).