Physica A 400 (2014) 124–138

Contents lists available at ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

Moral foundations in an interacting neural networks society: A statistical mechanics analysis

R. Vicente [a,*], A. Susemihl [b], J.P. Jericó [c], N. Caticha [c]

[a] Department of Applied Mathematics, Instituto de Matemática e Estatística, Universidade de São Paulo, 05508-090, São Paulo-SP, Brazil

[b] Artificial Intelligence Group, Technical University Berlin, Franklinstraße, 28/29, D-10587 Berlin, Germany

[c] Dep. de Física Geral, Instituto de Física, Universidade de São Paulo, CP 66318, 05315-970, São Paulo-SP, Brazil

HIGHLIGHTS

- Moral foundations statistics depend on peer pressure and cognitive style.
- An order–disorder transition is found in the peer pressure-novelty seeking plane.
- A mean field theory is constructed and confirms numerical simulations.
- Dynamical properties are consistent with cognitive/political affiliation groups.

ARTICLE INFO

Article history:
Received 24 July 2013
Received in revised form 22 October 2013
Available online 20 January 2014

Keywords:
Sociophysics
Social interactions
Opinion dynamics
Neural networks
Moral foundations theory

ABSTRACT

The moral foundations theory supports that people, across cultures, tend to consider a small number of dimensions when classifying issues on a moral basis. The data also show that the statistics of weights attributed to each moral dimension is related to self-declared political affiliation, which in turn has been connected to cognitive learning styles by the recent literature in neuroscience and psychology. Inspired by these data, we propose a simple statistical mechanics model with interacting neural networks classifying vectors and learning from members of their social neighbourhood about their average opinion on a large set of issues. The purpose of learning is to reduce dissension among agents when disagreeing. We consider a family of learning algorithms parametrized by δ, that represents the importance given to corroborating (same sign) opinions. We define an order parameter that quantifies the diversity of opinions in a group with homogeneous learning style. Using Monte Carlo simulations and a mean field approximation we find the relation between the order parameter and the learning parameter δ at a temperature we associate with the importance of social influence in a given group. In concordance with data, groups that rely more strongly on corroborating evidence sustain less opinion diversity. We discuss predictions of the model and propose possible experimental tests.

1. Introduction

Sociophysics, the approach to mathematical modelling of social science, is still maturing as a scientific field [1]. Opinion dynamics, voting, social influence and contagion models have been thoroughly studied [2,3], patterns in social data have been identified (e.g. Ref. [4] or Ref. [5] and references therein) and some successful predictions have been achieved (e.g. Ref. [6]).

In this paper our aim is to present a data driven statistical mechanics model for the formation of opinions about morality. We would like to verify if we can explain features of social data by considering a stylized model for neurocognitive processes

* Corresponding author. Tel.: +55 11 22943492.
E-mail addresses: rvicente@ime.usp.br (R. Vicente), nestor@if.usp.br (N. Caticha).

0378-4371/$ – see front matter © 2014 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.physa.2014.01.013

that are well described in the literature. Clearly practical limits to such a goal have to be considered. At the scale of individuals, neurocognitive data inspiring any modelling are always exposed to ecological validity issues with multiple uncontrolled causes. At the social scale, we also have to keep in mind the sheer complexity of human nature and human relationships. By stylized model we here mean a model to be used mainly to connect pieces of empirical evidence, to help the identification of important variables and as an aid to formulate new empirical questions. Furthermore, we would also like to have a testable model, namely, a model that makes some predictions after fitting a few key parameters to empirical data.

We argue here and in our previous work [7] that the evidence available on the moral classification problem can be accommodated by assuming agents that are conformist classifiers adapting to their social neighbourhood by reinforcement learning. Empirical evidence regarding different cognitive styles can then be represented in the model as distinct learning algorithms following the now established tradition of the statistical mechanics of learning [8]. Studies on social psychology [9] allow the further simplification of assuming that social influence only takes place between individuals perceived as similar. As a first approximation we thus assume that the social network can be partitioned into homogeneous social influence subnetworks, each one with a given cognitive style or learning algorithm.

But what do these conformist agents classify? We assume that any issue under debate can be parsed into a discrete set of independent attributes or dimensions. The modern theory of moral foundations [10] suggests that, as far as morality is concerned, these dimensions are not many more than five, namely: (a) harm/violence; (b) justice/fairness; (c) in-group loyalty; (d) respect for authority; and (e) purity or sanctity. For our modelling effort it is, however, sufficient that morality can be parsed into a discrete number of identifiable dimensions. As a starting point we do not consider the origin of these dimensions, its particular meanings or the practical issues that may be involved in trying to parse a given subject into these dimensions. These five dimensions have been found empirically to be sufficient to characterize political orientations along the liberal–conservative spectrum. The need for a sixth dimension, (f) liberty/oppression, has also been discussed to extend the description to include libertarians, but this is outside the scope of the data we have analyzed.

We thus consider that the moral content of an issue may be represented by a direction in a unit radius D-dimensional hypersphere $\mathbf{x} \in \mathbb{S}^D$. In the course of daily social relationships an individual j will be exposed to a variety of issues of diverse moral content parsed as $\mathbf{x}_j^\mu$ with $\mu = 1, 2, \ldots$. For each of these issues an opinion $\sigma_j^\mu \in [-1, 1]$ with a sign and an amplitude $|\sigma_j^\mu|$ is displayed. The sign can be interpreted as providing a for/against information and the amplitude as carrying information on how convict individual j is. A way to describe a classification task of this sort is by assuming that $\sigma_j^\mu = \mathbf{x}_j^\mu \cdot \mathbb{J}_j$, where $\mathbb{J}_j$ is an adaptive internal representation, inaccessible to other individuals, used by individual j to perform moral classification tasks. For simplicity we will study the case where all moral vectors are normalized to unit length $\mathbb{J}_j \in \mathbb{S}^D$. This also implies that differences in moral values are not interpreted as any type of moral superiority and that no moral shallowness is implied by the differences. Thus only the direction the moral vector points is considered as important, removing a layer of complexity in the interpretation of the model.

A conformist individual will then seek agreement with social neighbours in moral classifications by adjusting the internal representation $\mathbb{J}_j$. Employing the statistical mechanics of learning jargon, we are supposing that model agents are interacting *normalized linear perceptrons* [8] (for previous studies of interacting neural networks see Refs. [11–13]).

We suppose that in their daily lives agents interact within a social neighbourhood and exchange opinions about a large set of P issues. We assume that every issue has a *subjective* representation in the space of moral foundations $\mathbf{x}_j^\mu = \mathbf{x}^\mu + \epsilon_j^\mu$, with an idiosyncratic component ϵ_j^μ and an *objective* component $\mathbf{x}^\mu$. To further simplify the model we assume that agents do not adapt to each issue separately, instead, agents decrease their cognitive load by reducing the whole set of opinions of a neighbour to a single opinion about an average *objective* issue $\mathbb{Z}$ (that we call the *Zeitgeist* vector). Another layer of information compression can be added by assuming that agents are consistent in their moral classifications, to say, knowing the opinion about $\mathbf{x}_j^\mu$ determines the opinion about $-\mathbf{x}_j^\mu$ and agents only consider one of these alternatives at a time. With this restriction the opinion field would then be given by

$$h_j = \left(\frac{\sum_{\mu=1}^{P} \mathbf{x}_j^\mu}{\left\| \sum_{\mu=1}^{P} \mathbf{x}_j^\mu \right\|} \right) \cdot \mathbb{J}_j. \tag{1}$$

We also assume that there are no relevant biases or correlations in the individual parsing through the social network, and, by the law of large numbers, that idiosyncratic components cancel out as the number of issues discussed grows. We then write the opinion field as $h_j = \mathbb{Z} \cdot \mathbb{J}_j$, where the mean issue

$$\mathbb{Z} = \frac{\sum_{\mu=1}^{P} \mathbf{x}_j^\mu}{\left\| \sum_{\mu=1}^{P} \mathbf{x}_j^\mu \right\|}, \tag{2}$$

is supposed to be *objective* (or independent of the index j). We observe that our use of the idea of *objectivity* should be understood in a very restricted sense. We take it to mean that two agents, despite their differences in moral dimensions, will have access to a vector $\mathbb{Z}$ that is the same.

In summary, we suppose that conformist agents classify the average issue represented by the vector $\mathbb{Z}$ and exchange information about their classifications in the form of opinion fields $h_j = \cos\theta_j$, where θ_j represents the angle between the internal (moral) representation $\mathbb{J}_j$ and a symmetry breaking direction $\mathbb{Z}$ given by the mean issue parsed into moral dimensions. Additionally, we also assume that the objective mean issue $\mathbb{Z}$ changes very slowly in such a manner that we can concentrate our analysis on equilibrium properties.

For the sake of brevity we here only provide a short summary of empirical evidence and focus on the statistical mechanics model. To the reader interested in knowing more about relevant empirical sources we suggest reading our previous work on the subject [7]. Empirical evidence suggests that individuals are conformist agents that adapt to each other by reinforcement learning [14], that cognitive styles are diverse [15] and that agents are more strongly influenced by other agents with similar style [9]. Learning styles can be parametrized by $\delta \in [0, 1]$ that represents the ratio between how an agent modulates her learning in response to corroborating information (agreement) and how intensely she learns from novelty (disagreement). Agents with δ closer to one, weigh disagreement and agreement more similarly, while those with δ closer to zero give more weight to novelty than to corroboration. Additionally, psychological and neurocognitive data suggest a positive correlation between cognitive style and self-declared political affiliation (*p.a.*) [15].

The aggregate behaviour, represented by statistics of the opinion fields h, can be derived using statistical mechanics and then compared to social data on moral foundations. Given δ, the model *predicts* the shape of histograms $p(h|\delta)$. If we postulate that cognitive style, e.g. δ positively correlates with political affiliation, the model also predicts certain aspects of the behaviour of $p(h|p.a.)$. Alternatively, similarity between $p(h|\delta)$ and $p(h|p.a.)$ leads to a confirmation that cognitive style is positively correlated to political affiliation.

The simultaneous comparison of six predicted histograms $p(h|p.a.)$ ($p.a. = 1, 2, \ldots, 6$) with data requires the selection of two phenomenological parameters: the average node degree in the social influence subnetwork $\bar{k}$ and the average social pressure per social neighbour α. A mean field approximation predicts that histograms depend on the total social pressure, namely, on these two parameters combined as $\bar{k}\alpha$. An optimization procedure, designed to maximize the similarity between predicted and empirical histograms, can then be used to estimate $\bar{k}\alpha$.

This paper is organized as follows. In the next section we describe the available data. Section 3 describes the statistical mechanics model in detail and an analysis via Monte Carlo simulations. Section 4 treats analytically a mean field version of the model. In Section 5 we simulate the model in a real-world social graph provided by the Facebook. The dynamical behaviour of the model is discussed in Section 6. We discuss the meaning of diverse cognitive styles in the light of the model in Section 7. Finally, a section with conclusions and perspectives is provided.

2. Data on moral foundations

In a series of papers [10,16–20] Jonathan Haidt and coworkers have described moral foundation theory (MFT), a heuristically driven theory dealing with the foundations of moral psychology. Its aim is to understand statistically significant differences in moral valuations of social issues and their association to coordinates of a political spectrum.

Following Kohlberg [21] and Gilligan [22], work in moral psychology in the western world tradition dealt with the representation of moral issues in a two dimensional space. The first historically identified dimension is related to whether an action leads to harm and violence or not. Later, the existence of a second dimension, associated to justice and fairness was introduced. By the analysis of literature extending across time, geography and scientific disciplines, Haidt and coworkers introduced the main ingredient that yields a foundation theory [10]: humans when classifying issues as either moral or immoral, navigate not in a two dimensional space, but in one that is at least five dimensional. Valuations along these dimensions, called foundations in the literature of moral psychology, are necessary to characterize the moral content of a given issue.

Their striking quantitative result is extracted from massive amounts of data: as the political spectrum is traversed from very liberal ($p.a. = 1$) to conservative ($p.a. = 6$), there is an increase, from two to five, in the number of moral dimensions considered relevant by an average individual to form opinions (see Fig. 1). Liberals regard (a) harm/violence and (b) justice/fairness, the previously identified dimensions, as the most relevant foundations. In addition, conservatives hold (c) in-group loyalty, (d) respect for authority and considerations about (e) purity or sanctity in a considerable higher position than liberals. That means that, independent of the semantic role of the attributes, it can be asserted that liberals rely on a different subset of moral foundations than conservatives. We here employ liberal in the manner defined in the USA as social liberal.

The data[1] we have analyzed were furnished by Jonathan Haidt [10,23]. They were collected from the answers to a specially designed questionnaire aimed at probing opinions about morally relevant situations. Respondents were $N = 14\,250$ US nationals. For each respondent, Haidt and coworkers extracted five dimensional moral vectors with components related to five Moral Foundations. Each vector was labelled by the subject's self-declared political affiliation (*p.a.* ranges from $= 1$

1 Raw data can be requested to Haidt's team at http://www.yourmorals.org/. Preprocessed data can be downloaded at https://github.com/renatovicente/statmech_mft.git.

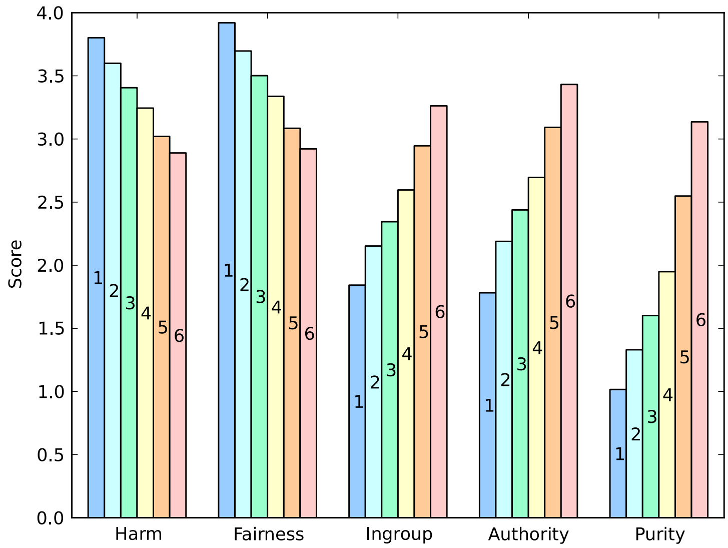

Fig. 1. *Experimental data—mean scores*: five dimensional vectors with components (scores) in the interval [0, 5] are extracted from questionnaires consisting of 30 questions. The figure represents mean scores for each moral foundation probed and for self-declared political affiliation ranging from (left to right) "very liberal" ($p.a. = 1$, blue) to "conservative" ($p.a. = 6$, red). While self-declared liberals in average rely more on harm and fairness foundations, conservatives rely equally in the five foundations.

(very liberal) to 7 (very conservative)). As the statistics for $p.a. = 6$ and $p.a. = 7$ are found to be indistinguishable (see Ref. [7]) these two classes are merged into $p.a. = 6$.

The questionnaires consisted of 30 questions probing the subject in the 5 moral dimensions. From the set of answers the five dimensional vectors with components in the interval [0, 5] are extracted. Thus a subject can either be represented as a point in the 30 dimensional space of questions, or in the reduced moral foundation space. It is interesting to see if the cloud of data points have a similar structure in both spaces. A negative answer would be indicative that the reduction has either deleted or invented some structure. We stress that we are not looking for clusters of different political affiliation in this analysis. The data should, if the questionnaires are relevant, characterize the relation between the complex moral valuation systems and the simple one dimensional continuous political affiliations.

The consistence of the five factors model has been already probed in Ref. [23]. We here confirm that the data reduction is significant employing a visualization technique known as SPIN [24] used in the analysis of large dimensional datasets in bioinformatics. It is a dimensional reduction technique that identifies a nonlinear one dimensional manifold irrespective of the embedding dimension of space. In both spaces we use a Euclidean distance to measure how different are any two individuals. A permutation of a set of individuals is done in order to give close labels to pairs that are close in the original space and try to give them far apart labels in case their distance is large. It has the advantage that the shape of one dimensional structures can be identified. Fig. 2 shows distance matrices for balanced sets of subjects. A random subsample of the data were constructed keeping the same sample size for each category of political affiliation. In the left column we show the matrices before the permutation, in the right column, after the permutation. In the top row, we show the bare data from the 30 dimensional space. In the lower row, the data of the reduced 5 dimensional space. The fact that the one dimensional structure that can be seen embedded in both spaces is similar, gives further support to the hypothesis that the 5 dimensional reduction to the moral foundations matrix from the 30 dimensional questionnaires preserve a one dimensional geometrical structure in the cloud of data points which is associated to political affiliation.

Fig. 3 depicts three components for $p.a. = 1$ (very liberal) and $p.a. = 6$ (conservative). For comparison purposes moral vectors $\mathbb{J}_i$ of the subjects were normalized to unit length as in the statistical mechanics model. In the model the vector $\mathbb{Z}$ is a symmetry breaking direction determined by the set of issues under discussion in a society. This set is a complicated thing to define. In particular, we have no access to the parsing that would permit its representation in five dimensions. To make possible a verification of the model, we have to identify the analogous of the direction $\mathbb{Z}$ within the data. Looking at Fig. 3 a reasonable choice, further justified in the next section, consists on identifying $\mathbb{Z}$ to the average vector within the conservative ($p.a. = 6$) and very conservative ($p.a. = 7$) classes.

We then calculate empirical histograms H_E for $h_j = \mathbb{J}_j \cdot \mathbb{Z}$ that characterize the different political groupings in a semantic free manner and will be compared to similar statistics obtained using analytical methods and numerical simulations.

3. Statistical mechanics model

Agents exchange information in the form of fields $h_j = \mathbb{J}_j \cdot \mathbb{Z} \in [-1, 1]$ that represent the mean opinion of agent j on a large set of issues or, considering that the information exchange is much faster than the adaptation dynamics, the opinion agent j has about the mean issue. The mean issue is objective and it is described by a set of D numbers $\mathbb{Z} \in \mathbb{S}^D$.

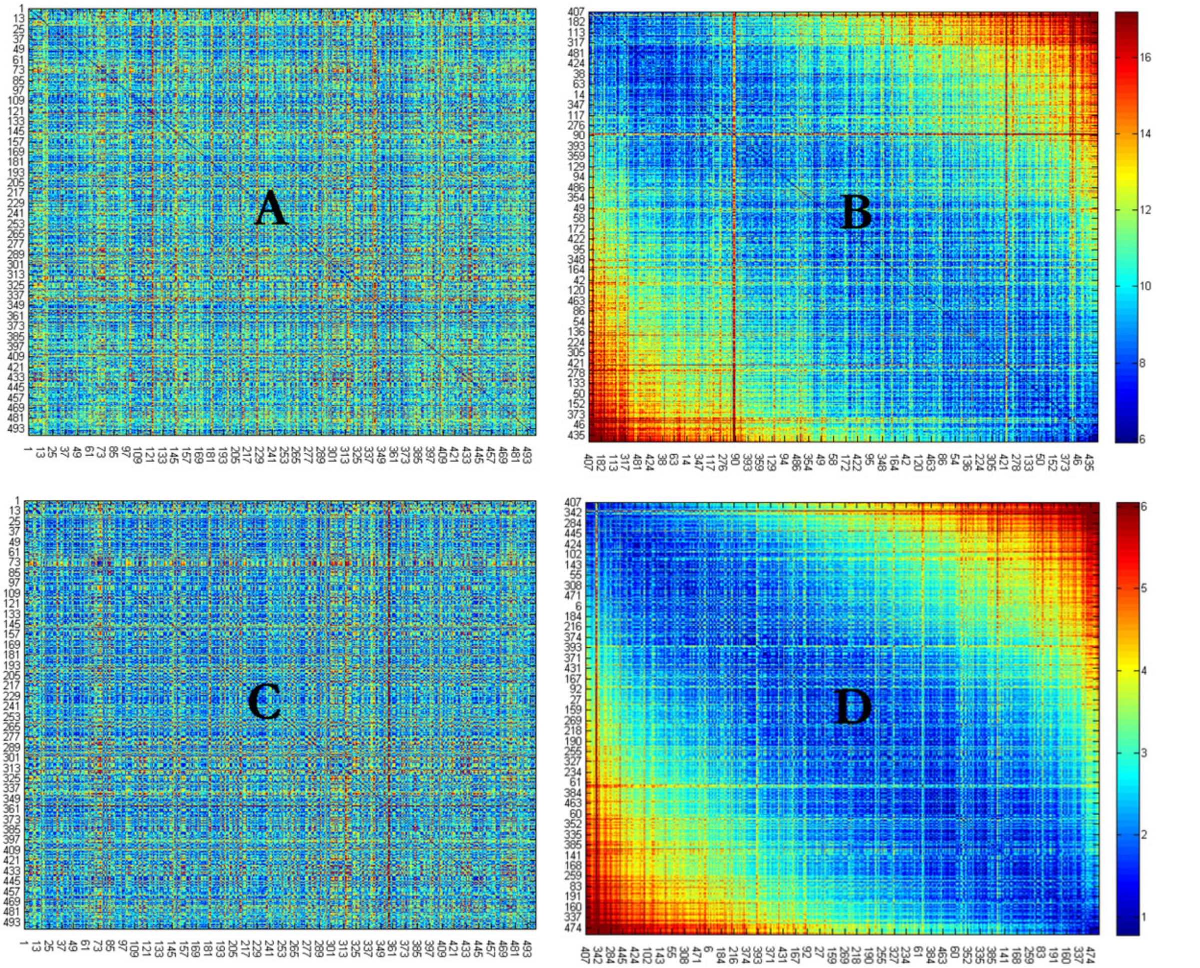

Fig. 2. *Experimental data—matrices of distances*: top row, (A) 30 dimensional representation of subjects, bottom row, (C) 5 dimensional space of moral foundations. Left column: before, Right column: (B) and (D), respectively, after application of SPIN algorithm. Blue means near and red far. The colour pattern in matrices B and D is typical of a one dimensional structure embedded in the respectively 30 and 5 dimensional spaces.

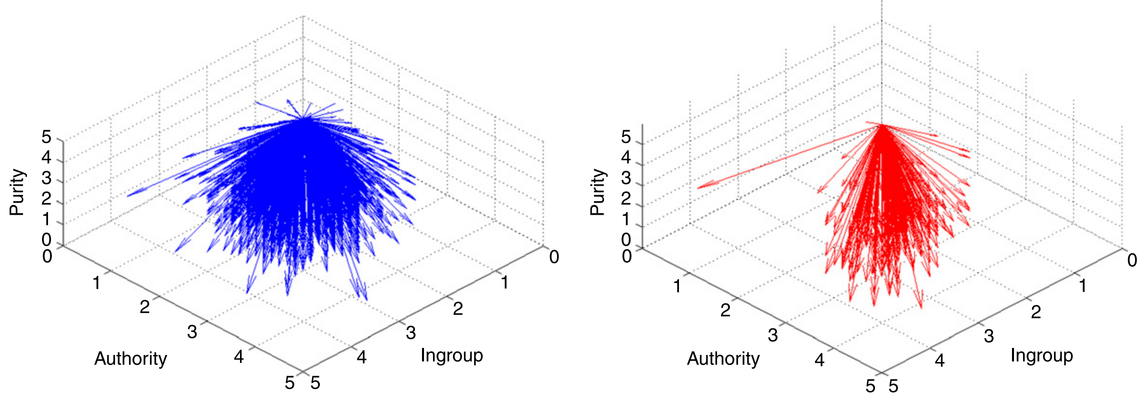

Fig. 3. *Experimental data—moral vectors*: Three-dimensional projections of moral vectors for very liberal ($p.a. = 1$, blue) and conservative ($p.a. = 6$, red) subjects. Axes are labelled according to the associated moral foundation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The relevant variables, representing the society, are the internal variables of the agents. Every agent j has two main properties. (A) Its internal state is determined by a set of D weights $\mathbb{J}_j$ (*moral vector*), which is invisible to other agents. (B) The main hypothesis in this work is that while weights jointly code for prior experience, they are subject to change due to the social interactions through a learning mechanism.

The vector $\mathbb{Z}$ changes in time reflecting social changes in moral parsing or values. We, however, consider that the adaptation dynamics of $\mathbb{J}_j$ is much faster than the dynamics of $\mathbb{Z}$ and suppose the latter as being fixed. We also concentrate on $D = 5$, but it might be interesting to explore the consequences of using different values.

It is interesting at this point to discuss a bit further what are the implications of parsing both issues and moral vectors. We could certainly find issues $\mathbf{x}$ with few or no moral content. For instance "Some cars are blue" is basically free of moral content. This particular issue (or assertion) could be parsed into some other dimensions, but its projection onto the moral subspace would be null. Some other issues are intrinsically filled with moral content. For example: "There should be gender equality". How this issue is to be parsed into five (or more) moral foundations depends on culture and time. For our model, however, it is enough to be able to use that there is objective moral content in every issue debated (objectivity approximation). Notice that the moral vector $\mathbb{J}$ also measures moral content, however, it is understood as the subjective rule used by an agent to classify an issue. For instance, an average liberal would judge "There should be gender equality" in terms of fairness and harm dimensions while someone more conservative would also take into account loyalty, purity and respect for authority. It is important to keep in mind that the model is necessarily a cartoon of reality designed to organize our thought and help us in asking meaningful and productive new questions.

The only interaction among agents comes from learning about the opinion fields of other agents in their social neighbourhood. Learning occurs in order to decrease the psychological discomfort due to dissent. Learning is described by a noisy gradient descent dynamics on a potential function describing a psychological cost of disagreement with each of its social neighbours. In Ref. [7] we have introduced a potential to model this and called it the psychological cost, which depends on a parameter δ, taking values between 0 and 1. δ represents an attempt to model different cognitive strategies with respect to how novel or corroborating information is used in the learning process. For $\delta = 0$, as will be seen below, the agents can be called error correctors. They only learn from social neighbours from which they disagree by bringing a different average opinion on the set of issues under discussion. Thus we also refer to these agents as *novelty* seekers, for they do not learn unless the information carries a new and different point of view on the issues. For $\delta = 1$, learning occurs by extracting correlations. These agents learn from neighbours always, independently of agreement or disagreement. This led us to call them *corroboration* seekers.

The family of psychological costs or interaction potentials, indexed by δ (see Fig. 4) is defined by:

$$V_\delta(h_i, h_j) = -\frac{1+\delta}{2}h_i h_j + \frac{1-\delta}{2}|h_i h_j| \tag{3}$$

which can be written as $V_\delta = -\delta h_i h_j$ for same sign opinions and $V_\delta = -h_i h_j$ for opinions of a different sign. We can also consider the total cost for a group homogeneous in δ leaving on a social graph $\mathcal{G}$ as

$$\mathcal{H} = \sum_{(i,j)\in\mathcal{G}} V_\delta(h_i, h_j). \tag{4}$$

There are a few reasons that justify using the same δ for both agents in each interaction. First there is evidence [9] that people tend to interact more with those of similar cognitive styles. Second we have tried in simulations with different δ's in the same population and in different social networks and the qualitative results are similar as long as δ's are independently and uniformly distributed throughout the social network, showing the robustness of this approximation. Finally a third reason is that it simplifies the analytical mean field calculations which we discuss herein.

The specific form of the potential is inspired in learning algorithms for linear classifiers [8,13]. A Hebbian algorithm can be considered to lead to learning from the use of correlations in the input and output units. Information from a pair (issue, opinion) will be embedded with a strength independent of whether a prediction was correct or not. A Perceptron algorithm, on the other hand works by error corrections. If the prediction, on an example was correct, it will not do any changes. Changes will be made only when the prediction was incorrect. In this sense we can say that a Hebbian algorithm learns both from corroborating and from new information. A Perceptron algorithm will only learn from new information.

We consider a discrete time dynamics with information exchanges that are asynchronous as parallel updating [25] would not be reasonable for these exchanges of information, because they would require some sort of an external clock. We also assume that there are no relevant delays [26] in information exchanges.

Learning proceeds in the following way. At each time step an agent is chosen and its weights are updated, if there is no noise in the communication, using a gradient descent dynamics:

$$\mathbb{J}_i(t+1) = \frac{\mathbb{J}_i(t) - \epsilon\nabla_{\mathbb{J}_i(t)}\mathcal{H}}{\|\mathbb{J}_i(t) - \epsilon\nabla_{\mathbb{J}_i(t)}\mathcal{H}\|}, \tag{5}$$

where ϵ defines the time scale.

We can also consider the case where noisy exchange of opinions might drive the update uphill and to describe this scenario we introduce an inverse temperature α and a Monte Carlo Metropolis dynamics [27]. As usual, then choose a D

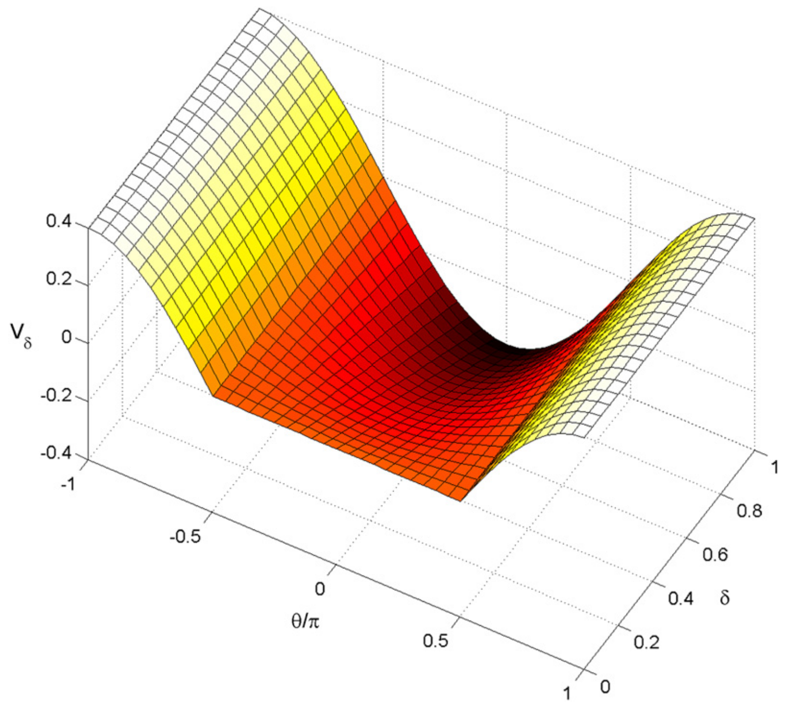

Fig. 4. *Psychological cost function* $V_\delta(h_i, h_j)$ for different values of δ, as a function of θ_i/π, where $h_i = \cos\theta_i$ for $h_j = 0.4$ fixed. For fixed δ, the slope of the potential determines the scale of changes of the moral vector. For $\delta = 0$ changes only occur if there is a difference in the signs of opinion fields h_i and h_j. For $h_j > 0$ this occurs when h_i becomes negative at $|\theta_i| > \pi/2$. On the other extreme of cognitive styles, for $\delta = 1$, any difference in magnitude of opinions has an associated slope of the cost function.

dimensional vector $\mathbf{u}$ drawn uniformly on a ball of radius ϵ. A trial weight vector is defined by

$$\mathbb{T} = \frac{\mathbb{J}_i(t) + \mathbf{u}}{\|\mathbb{J}_i(t) + \mathbf{u}\|} \tag{6}$$

and accepted as the new weight vector, $\mathbb{J}_i(t+1) = \mathbb{T}$ if the social cost decreases: $\Delta\mathcal{H} := \mathcal{H}(\mathbb{T}) - \mathcal{H}(\mathbb{J}_i(t)) \leq 0$. If $\Delta\mathcal{H} > 0$ the change is accepted with probability $\exp(-\alpha\Delta\mathcal{H})$. This leads, after a transient, to a distribution of states given by the Boltzmann distribution $P_B(\{\mathbb{J}_i\}) \propto \exp(-\alpha\mathcal{H})$.

Alternatively we can proceed by making explicit the hypothesis that the average social cost characterizes macroscopic states and suppose that the expected value $\mathbb{E}[\mathcal{H}]$ has a certain value. The distribution of probabilities for the moral vectors of the agents has to be chosen from among those that satisfy the information constraint and makes the least amount of additional hypotheses. This is the natural framework of maximum entropy, which leads again to the Boltzmann distribution. The parameter α, which characterizes the noise level in the first approach, appears now as a Lagrange multiplier. It can be seen to determine the scale in which changes in social cost, brought about by differences in opinion, are important. This justifies being called the *scale of peer pressure*. Thus the scale of peer pressure is analogous to the noise amplitude of the exchange of information and to an inverse temperature in statistical mechanics. If there is a high level of noise, then the opinions of others will not be very influential, thus a low peer pressure. As the dependence of fluctuations on temperature permits measurements of one in terms of the other, in statistical mechanics, fluctuations in opinions may be empirically used to characterize the level of peer pressure in a society.

This system has been studied using Monte Carlo methods and Mean Field methods. The main empirical finding we focus on is the difference in the statistics of moral foundations between self-declared liberals and conservatives. To study this we need to introduce an appropriate order parameter. Our model, however, has no semantics. Concepts like "pure", "harmless", "loyal" in our model are just represented by indistinguishable dimensions of a vector space. Our goal is to model the statistics of moral vectors around a symmetry breaking direction. Notice that the Hamiltonian (3), that describes the social learning process, has ground states at a region around $\mathbb{J}_j \cdot \mathbb{Z} = 1$ for all j, where the symmetry breaking direction is introduced by the objective part of the issues under discussion. This symmetry breaking direction may be regarded as the simplest vector to characterize the society and what its members are discussing. But how can we find $\mathbb{Z}$ empirically?

We know that the variance of the angle among moral vectors $\mathbb{J}_k$ of a homogeneous δ sub-population decreases with $p.a.$ We thus assume that the average moral vector in conservative and very conservative ($p.a. = 6$ and 7) sub-populations provides a good estimate of $\mathbb{Z}$. Empirically we also know that this estimate is close to $\mathbf{1} = (1, 1, 1, 1, 1)$ for our dataset, we thus fix $\mathbb{Z} = \mathbf{1}$. Our strategy then is to characterize both the state of the agents model and the empirical questionnaire data

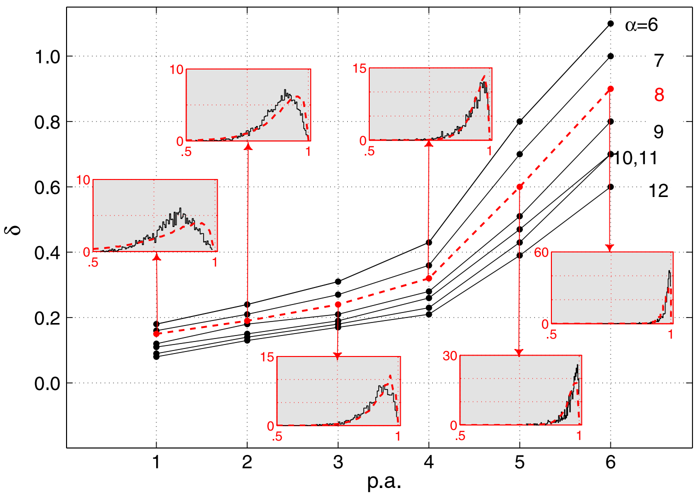

Fig. 5. *Cognitive style and political affiliation*. The similarity of experimental and simulated histograms suggests a link between cognitive style and political affiliation. In this figure δ is chosen at a fixed peer pressure α to yield a simulated histogram that matches empirical histograms. The insets show empirical normalized histograms H_E (black full line) and simulated normalized histograms H_S as a function of opinion fields h (dashed red line) for the case $\alpha = 8$, that yields a good fit but must be regarded as only illustrative. The social graph is a Barabási–Albert (BA) network [28] with average degree $\bar{k} = 22$ (branching parameter $M = 11$) and size $N = 400$. The horizontal location of the insets is indicative of the political affiliation group to which the empirical histogram pertains. The best δ are depicted by the lines for several α. The case $\alpha = 8$ is illustrative and is indicated as a red dashed line. The relation suggests that liberals have smaller δ than conservatives, meaning that liberals rely less on corroborating information than conservatives.

by introducing rotationally invariant order parameters which are semantically free as far as possible. Further analysis of the semantics of the model would be outside of the scope of this paper.

Simulated histograms H_S for opinion fields h are calculated as equilibrium distributions in a Monte Carlo Metropolis dynamics run in a social graph that we here choose, as a representative illustration of typical results, to be a Barabási–Albert (BA) network [28] with average degree $\bar{k} = 22$ (branching parameter $M = 11$) and size $N = 400$. As the relation between the cognitive parameter δ and empirical affiliation $p.a.$ is unknown we proceed by fixing the also unknown (but robust) peer pressure parameter α and finding for each $p.a. = 1$–6, by using simple incremental search, the δ that minimizes a Euclidean distance between simulated histograms $H_S(\ell \mid \delta, \alpha)$ and empirical histograms $H_E(\ell)$ defined as

$$\mathcal{D}\left[H_S \mid H_E\right] = \sum_{\ell=-L}^{L} \left[H_S(\ell \mid \delta, \alpha) - H_E(\ell)\right]^2, \tag{7}$$

where the interval $[-1, +1]$ for h is appropriately binned such that $h^{(\ell)} = \ell/L$ for $\ell = -L, \ldots, L$. This procedure results in the curves depicted in Fig. 5. The model behaviour is consistent with empirical evidence [15] in its general features, namely, δ is an increasing function of the empirical political affiliation. We also notice the robustness in the qualitative behaviour as we vary the peer pressure α. For the fits depicted as insets in the figure we use for illustration purposes $\alpha = 8$.

We also calculate thermodynamic quantities by employing the Wang–Landau technique [7,29]. By doing that we are able to compute the phase diagram of Fig. 6. From the point of view of ordering, the resulting diagram is straightforward exhibiting an ordered $m = \langle h \rangle > 0$ phase and a disordered phase with $m = \langle h \rangle = 0$ separated by a continuous transition line that is well-fitted, in the case of a Barabási–Albert network, by a simple power law $\alpha \propto 1/\delta$ [7].

4. Mean field analysis

This section aims at providing theoretical support to the numerical results previously presented. Our understanding of the model can be enriched by studying a tractable approximation with qualitatively similar behaviour. Let us consider the set of issues to be fixed (quenched disorder). In the analysis of this section we will not deal with the difficult task of averaging over quenched disorder, since it would draw attention and direct energy to technical issues beyond our current purpose. We fix a set of issues and study the resulting thermodynamics. The problem is still not simple and an exact solution for the statistical mechanics problem is not known. Here we present mean field results obtained from information theory considerations in the form of a Maximum Entropy argument. We introduce a space of tractable probability distributions, which factor over groups of agents. The first and the simplest choice is to consider a tractable family that factors over the individual agents: $P_0 = \prod_i P_i$. The parametrization of P_i will be done in terms of the order parameters which we still do not know. An advantage of the mean field approach along these lines is that it tells the relevant order parameters. Our problem

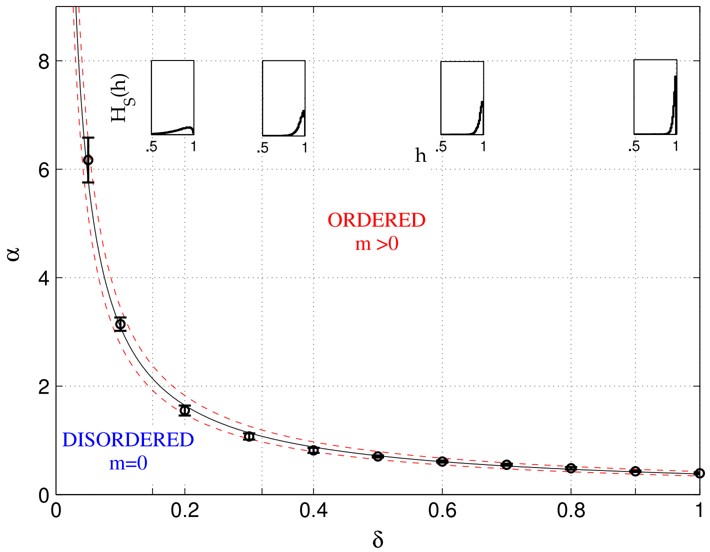

Fig. 6. *Phase diagram in the space δ vs. α*. The case depicted corresponds to a BA network with average degree $\bar{k} = 22$ and size $N = 400$. Points correspond to 20 runs of a Wang–Landau algorithm. The insets show the histogram $H_S(h)$ obtained as the equilibrium of a Monte Carlo Metropolis dynamics with peer pressure $\alpha = 8$ and δ provided by the optimization process that yields Fig. 5. The location of the insets is indicative of the parameters δ and α used in the simulation. Full line indicates a fit $\alpha \propto 1/\delta$, with red dashed lines corresponding to 95% confidence error bars.

is reduced to minimization of the relative entropy

$$S\left[P_0 \parallel P_B\right] = -\int \left(\prod_i \mathrm{d}\mu(\mathbb{J}_i)\right) P_0 \ln \frac{P_0}{P_B} - \lambda(\langle P_0\rangle_\mu - 1) \tag{8}$$

where P_B is the Boltzmann distribution for the above Hamiltonian, $P_B = \exp(-\alpha\mathcal{H})/Z$ and $\mathrm{d}\mu(\mathbb{J}_i)$ is the uniform measure over the surface of the D sphere. The only relevant constraint that have to be taken into account is the normalization. The fact that the expected value $\langle\mathcal{H}\rangle_\mu$ has a given fixed value $\mathcal{E}$, which might even be unknown, but is important in characterizing the state of the agent society at least with respect to the opinions about the issues, is taken into consideration by the choice of P_B as a Boltzmann distribution.

We can drop the logarithm of the original partition function $\ln Z$ without changing the variational problem to obtain, from Eqs. (3) and (8)

$$S\left[P_0 \parallel P_B\right] = -\sum_i \int \mathrm{d}\mu(\mathbb{J}_i)P_i \ln P_i - \lambda \int \mathrm{d}\mu(\mathbb{J}_i)P_i - \alpha \sum_{(i,j)} \int \mathrm{d}\mu(\mathbb{J}_i)\mathrm{d}\mu(\mathbb{J}_j)P_iP_jV_\delta(h_i, h_j)$$

and considering variations of the set of P_i, $\frac{\delta S[P_0\|P_B]}{\delta P_i} = 0$, leads to

$$0 = -1 - \lambda - \ln P_i - \alpha \sum_{(j),\nu} \int \mathrm{d}\mu(\mathbb{J}_j)P_jV_\delta(h_i, h_j).$$

This is an expression relating the probability density of an agent to those of the social neighbours:

$$P_i \propto \exp\left(-\alpha \sum_j \int \mathrm{d}\mu(\mathbb{J}_j)P_jV_\delta(h_i, h_j)\right). \tag{9}$$

Now we go back to the problem of choosing the family of distributions P_i. The main reason to call a family tractable is that the set of equations above is closed. Depending on the structure of the Hamiltonian, different families can be used.

The form of the Hamiltonian imposes the use of two order parameters for each issue, which in order to close the set, we take to be independent of the agent,

$$\int \mathrm{d}\mu(\mathbb{J}_j)P_jV_\delta(h_i, h_j) = -\frac{1+\delta}{2}h_i m + \frac{1-\delta}{2}|h_i| r \tag{10}$$

where we have introduced

$$m = \int \mathrm{d}\mu(\mathbb{J}_j)P_j\ \mathbb{J}_j \cdot \mathbf{x} \tag{11}$$

$$r = \int \mathrm{d}\mu(\mathbb{J}_j)P_j\ |\mathbb{J}_j \cdot \mathbf{x}|. \tag{12}$$

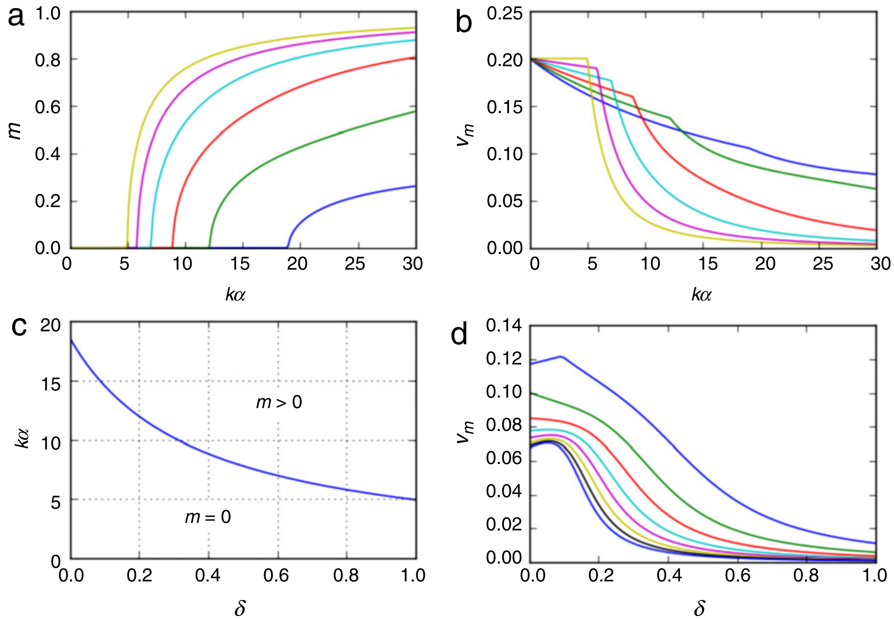

Fig. 7. *Mean field theory*: (a) $m = \langle h \rangle$ as a function of total peer pressure $k\alpha$ for $\delta = 0, 0.2, 0.4, 0.6, 0.8, 1.0$ (from top to bottom) (b) $v_m = \langle h^2 \rangle - \langle h \rangle^2$ as a function of $k\alpha$ (δ from bottom up in the right side of the picture). (c) Phase diagram. (d) Width of distribution v_m as a function of δ, for fixed $k\alpha = 15, 20, 25, 30, 35, 40, 45, 50$ (from top to bottom).

In principle the order parameters m and r could have an index j identifying the agent, but we make a reasonable assumption of homogeneity. This does not mean that all agents are equal, but that they will present values of the moral vector $\mathbb{J}_i$ drawn from the same probability distribution. Then the mean field probability distribution is given by

$$P_{\mathrm{MF}}(\{\mathbb{J}\}|k\alpha, \delta, m, r) = \prod_i P_{\mathrm{MF}}(\mathbb{J}_i|k\alpha, \delta, m, r) = \prod_i \frac{\exp\left\{k\alpha\left(\frac{1+\delta}{2}h_i m - \frac{1-\delta}{2}|h_i r|\right)\right\}}{\mathcal{Z}_i}, \tag{13}$$

where the denominators $\prod_i \mathcal{Z}_i$ ensure normalization and k is the number of social neighbours. Now Eqs. (11) and (12) can be seen not as the definitions of m and r, but as the self consistent mean field theory equations from which their values can be calculated.

The model can be studied for any value of the dimension of the internal space, D. We use $D = 5$ and our problem is reduced to doing some integrals of up to five dimensions. Since there is only one symmetry breaking direction $\mathbb{Z}$, we rotate the coordinate system such that $\mathbb{Z}$ is in the $\hat{e}_5$ direction. The polar angle $\theta_3 := \theta$ with this direction is the only non trivial integration variable since the other angular variables (θ_0, θ_1 and θ_2) drop out and are trivial.

Call $B(\theta|k\alpha, \delta, m, r) := \exp\{k\alpha\,(am\cos\theta - br|\cos\theta|)\}$ where $a := \frac{1+\delta}{2}$ and $b := \frac{1-\delta}{2}$, then

$$\begin{aligned} m &= \frac{1}{\mathcal{Z}}\int_0^\pi \mathrm{d}\theta\, \sin^3\theta\cos\theta B(\theta|k\alpha, \delta, m, r) \\ r &= \frac{1}{\mathcal{Z}}\int_0^\pi \mathrm{d}\theta\, \sin^3\theta|\cos\theta|B(\theta|k\alpha, \delta, m, r) \\ \mathcal{Z} &= \int_0^\pi \mathrm{d}\theta\, \sin^3\theta B(\theta|k\alpha, \delta, m, r). \end{aligned} \tag{14}$$

Eqs. (14) can be solved numerically self consistently. Results in Fig. 7(a) show the fixed points m as a function of the total peer pressure, showing the existence of a phase transition as the critical line of total peer pressure $k\alpha_c(\delta)$ depicted in 7(c) is crossed. The critical total peer pressure $k\alpha_c$ decreases with larger values of δ. Fig. 7(b) shows the dispersion of the distribution of fields $\mathbb{J}\cdot\mathbf{x}$ (denoted v_m). An important prediction of the theory is that this depends strongly on the corroboration parameter δ as it can be seen in Fig. 7(d).

We can use Eq. (13) to calculate the distribution of opinions about the symmetry breaking direction $\mathbb{Z}$

$$P(h|k\alpha, \delta) = \int \mathrm{d}\mu(\mathbb{J})\delta(\mathbb{J}\cdot\mathbb{Z} - h)P_{\mathrm{MF}}(\mathbb{J}|k\alpha, \delta, m, r). \tag{15}$$

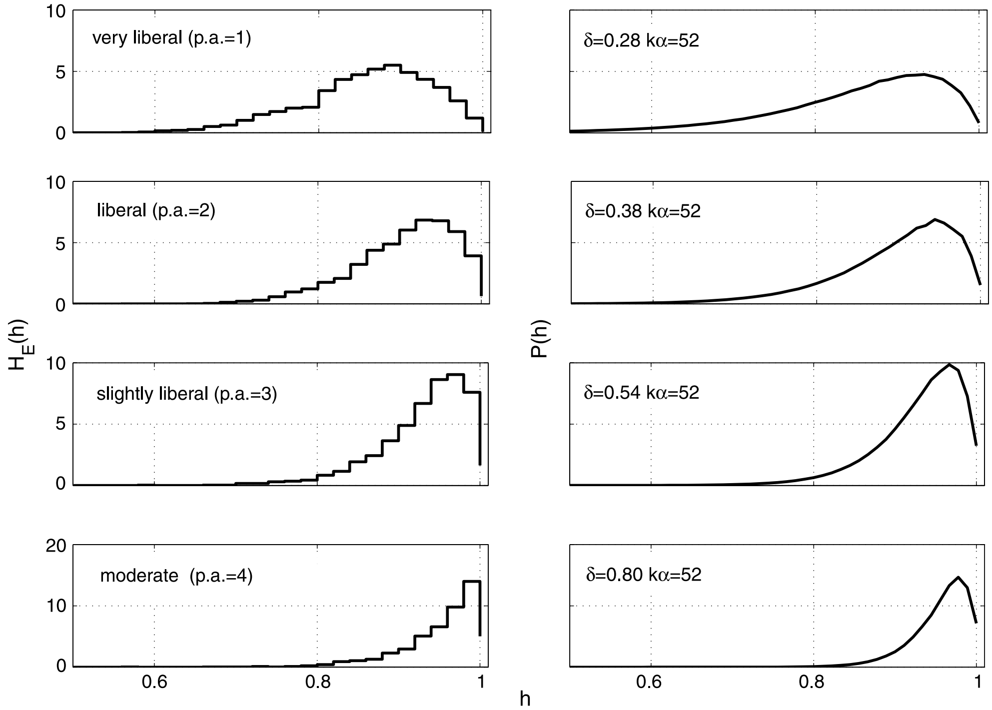

Fig. 8. *Mean field theory histograms*. Left column: empirical histograms. Right column: mean field results. To recover histograms for $p.a. > 4$ a larger $k\alpha$ is required. Vertical axes are shared in each row.

This is a mean field prediction that can be confronted to Monte Carlo simulations and, more importantly, to experimental data. The result is

$$P(h|k\alpha,\delta) = \frac{1}{C}(1-h^2)\exp\{k\alpha\,(ahm - br|h|)\} \tag{16}$$

where $C = \int_{-1}^{1}(1-z^2)\exp\{k\alpha\,(azm - br|z|)\}\,\mathrm{d}z$, is given to good approximation by

$$C = \frac{2}{\delta^2\tilde{m}^2}\left(1-\frac{1}{\delta\tilde{m}}\right)\mathrm{e}^{\delta\tilde{m}} - \frac{1-\delta}{\delta\tilde{m}} + \frac{2}{\tilde{m}^3}(\tilde{m}-1)\mathrm{e}^{-\tilde{m}} \tag{17}$$

where $\tilde{m} = k\alpha m$ and we used that in the experimentally relevant region $m = r$. Approximately

$$P(h|k\alpha,\delta) = \frac{(\delta\tilde{m})^2}{2}(1-h^2)\mathrm{e}^{-\delta\tilde{m}(1-h)}. \tag{18}$$

This comparison is shown in Fig. 8, it hints that the similarity of the data and theory point to a relation between δ for the agents and political affiliation for the experimental subjects. Histograms for the more conservative groups resemble more the histograms for agents with higher δ's; i.e. conservative behaviour is more likely to be identified with larger reliance on corroboration and alternatively liberal behaviour, with smaller reliance.

5. Facebook network

In the previous sections we have based our discussion on simulations run on regular and "synthetic" BA networks. Characteristically real social networks are modular what may have important dynamical consequences [30]. To check the robustness of our conclusions, in this section we present simulation results on a realistic network extracted from Facebook network data[2] [31,32].

In Fig. 9 we show a comparison between empirical $H_E(h)$ histograms and the best fit, in terms of a Euclidean metric, for the Princeton graph (size $N = 6596$ and average degree $\bar{k} = 88.9$) and for a Barabási–Albert construction with $N = 800$ and $\bar{k} = 22$ [28]. To build theoretical histograms we have run Metropolis simulations[3] [7] fixing $\bar{k}\alpha = 176$ in both scenarios and choosing, for each $p.a.$, a δ that minimizes the metric defined by Eq. (7).

2 Data can be downloaded at http://archive.org/details/oxford-2005-facebook-matrix.

3 Scripts and results are available at https://github.com/renatovicente/statmech_mft.git.

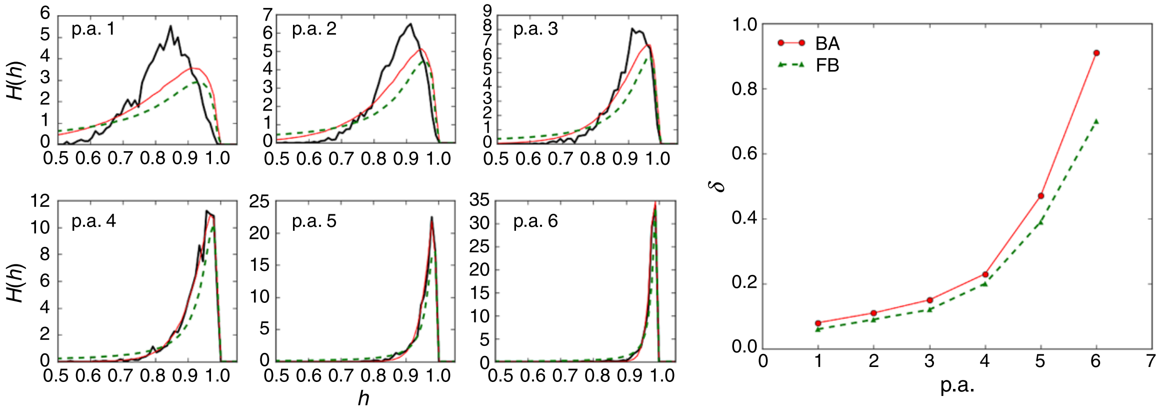

Fig. 9. *Histograms for empirical data, synthetic and Facebook networks*. Left panel: the thick black lines are the empirical histograms $H_E(h)$, red lines are the simulated histograms $H_S(h)$ with a Barabási–Albert construction of size $N = 800$ and $\bar{k} = 22$. The peer pressure is fixed to $\alpha = 8$. Dashed green lines are the simulated histograms $H_S(h)$ using Princeton's Facebook network ($N = 6596$ and $\bar{k} = 89$) with $\alpha = 1.98$. Right panel: best δ as a function of *p.a.* for the BA network (dashed red lines) and for Princeton's Facebook network (thin green line).

As it is suggested by the mean field theory of the previous section, for homogeneous δ, histograms, in good approximation, only depend on the social topology through the average total peer pressure $\bar{k}\alpha$. As a consequence, a BA network and a Facebook network have very similar phase diagrams in the plane $(\delta, \bar{k}\alpha)$. Also we see that to find from the data the peer pressure per social neighbour α we have first to measure the average degree independently.

6. Dynamics: what do *conservative* agents conserve?

In addition to obtaining that novelty seeker agents are identified with liberals and that corroboration seeker agents are more similar to conservatives, the model can be studied to determine dynamic collective properties. In particular we study in this section how groups of agents identified with conservative or liberal differ in time scales to adopt new positions. Given the relation between political affiliations and the corroboration parameter suggested by the model, it would be a contradiction if characteristic reaction times to changes in the symmetry breaking direction $\mathbb{Z}$ turn out to decrease with increasing δ. So, putting the theory to the test, we now turn to study the response to changes of the issues and how the group accommodates to such changes. Once the MC simulation has equilibrated, we change the $\mathbb{Z}$. The new direction and the old one have an overlap $\mathbb{Z}_{\text{old}} \cdot \mathbb{Z}_{\text{new}} = \cos\zeta$. We continue the Metropolis simulation and characterize, as a function of simulation time, the distance to the equilibrium distribution. A natural distance from equilibrium would be a measure of the Kullback–Leibler divergence. However, we do not have access to a theoretical form of the out-of-equilibrium distribution. A simpler procedure is to calculate a distance directly from the histograms. After a MC step, which includes a learning sweep over the whole population, we obtain $H_t(h)$ the histogram of opinions $h_{\text{new}} = \mathbb{J} \cdot \mathbb{Z}_{\text{new}}$ about the new symmetry breaking direction, giving the fraction of agents with opinion in a given range. Define the Euclidean distance by

$$\mathcal{D}\left[H_t|H_{\text{eq}}\right] = \sum_{h=-1}^{1} \left(H_t(h) - H_{\text{eq}}(h)\right)^2 \tag{19}$$

where the range of the variable h has been discretized into 20 bins. The distance from equilibrium as a function of time can be parametrized as $\mathcal{D}(t) = F(\zeta)e^{-t/\tau}$, where the measured $\tau = \tau(\zeta, \alpha, \delta)$ appears in Fig. 10. The valley, shown in blue, shows the region where the agents are faster to re-equilibrate adapting to the new conditions. It occurs inside the ordered phase, not in the high δ region of the conservatives, nor at the border of the phase transition. This is to be expected, since at the border there is critical slowing down. The interesting thing is that the group of agents that re-adapts to equilibrium the fastest is the one which has been identified with the most liberal subjects of the data.

The surprise lies not in that ultra-liberals adapt the fastest, but that our simple model is also consistent in this respect. This result is central to our proposal for interpreting what conservative and liberal mean in a group of agents. Agents with high δ are conservative and with lower δ more liberal, also in terms of their time to adapt to changes. Based on this we can attribute political labels to the agents which are consistent with the attribution based on the data since they show the same dependence on δ.

7. Discussion: diversity of cognitive styles

Statistics and reaction times permit identifying different cognitive styles with different statistics of opinions about a mean issue (that we call *Zeitgeist* in Ref. [7]). Therefore given that people do present different cognitive styles and that they

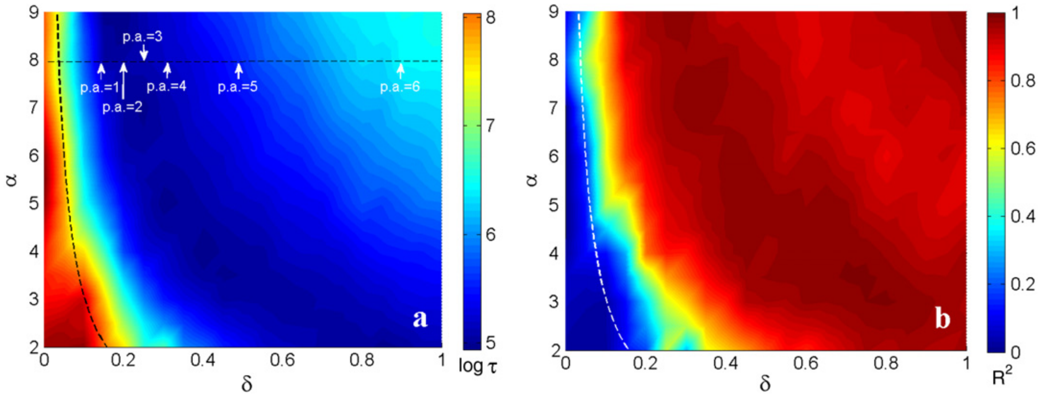

Fig. 10. *Characteristic re-adaptation times*. The topology is the same used to produce Fig. 6. Dashed lines indicate the phase transition. A rotation by an angle of $\zeta = 0.6\pi$ rad has been employed as a perturbation to the vector $\mathbb{Z}$ all over, however, conclusions are general. (a) Re-adaptation times τ as a function of α and δ. Arrows indicate political affiliation scores as identified in Fig. 5, from $p.a. = 1$ (very liberal) at left to $p.a. = 6$ at right. Note that the minimum time occurs at values of δ inside the ordered phase (see Fig. 6) and that this value of δ corresponds to liberal and very liberal scores (Fig. 5). (b) Re-adaptation times τ are inferred by regressing $\log \mathcal{D}$ against t. R^2 statistics are indicative that the data fits employed are only relevant inside the ordered phase.

interact and learn from each other we expect that there will be people who hold different sets of moral values. Cultural wars will follow from diversity of cognitive styles.

This is a semantic free conclusion. The model cannot distinguish between the different moral foundations. There must be another reason why some of the foundations are always present while others may be absent. Evolutionary arguments by Haidt go a long way in explaining why the harm/care and justice dimensions are more uniformly common. They may be found, to some extent in other primates [33,34]. The emergence of the other dimensions, which seem to be present only in humans, are supposed to foster higher cooperation levels and ultra-social behaviour. Suppose, as we do, these facts to be reasonable, then the question that emerges is why society has kept all types of cognitive styles and not only those that lead to a more cohesive society? A possibility is that different strategies within a society are useful to cope collectively with different challenges. A higher level of agreement yields higher fitness during times where current opinions lead to correct answers from a survival point of view. Conservative behaviour would then be the fittest when maintaining current behaviour is beneficial to the society. However, during times when current opinions are not guiding in the finding of useful answers, in a survival enhancing sense, a different perspective is needed. A larger menu of choices may permit finding better alternatives in an efficiently distributed manner.

When the Zeitgeist changes, due to external conditions, or due to a new issue being introduced to the debate, a more liberal approach seems reasonable. The question, from the current perspective, is then translated to whether this behaviour can be seen within our model. In Fig. 5 we presented a connection between agents characterized with a given δ parameter and the political affiliation of the questionnaire respondents. A question that arises is why a lower but non-zero value of δ was found for a peer pressure around $\alpha = 8$ which is well inside the ordered phase? Why were not the ultra-liberals associated to a δ right on the edge of the phase transition? Maybe the reason can be found in the dynamics of adaptation to the new Zeitgeist. We find from the simulations, within the appropriate α region, that at that value, $\delta \approx 0.20$–0.35 the characteristic time of re-adaptation to a new Zeitgeist has a minimum (see Fig. 10). The lower the conservatism of a population the less cohesiveness it will present in responding to external challenges as a group. There is no benefit in being more liberal than necessary. Letting the analogy to lead us to its ultimate consequences, ultra-liberals are not on the disordered phase, but in the ordered phase. They even are not at the border of the transition, they are in a way prevented from being on the disordered phase by critical slowing down. Closer to the border the system is softer but takes longer to rearrange. And from our results it seems that even the ultra-liberals observed in the data rely on some corroboration in order to construct their moral vector.

8. Conclusions and perspectives

The modelling of social systems has a long (and well-fought [1]) history. It might be surprising to some that a mathematical model can be constructed and directly confronted to data, replicating some statistical findings and making predictions borne out by observations. We believe that this is possible by setting the problem in the context of information theory. After relevant variables were identified, information about neurocognitive, psychological and social science was used to attribute a probability distribution for the variables, which is finally used to estimate relevant experimental signatures from order parameters. This is, ultimately, what is done in traditional areas of Physics.

We try not to take sides in the study of morality, what is acceptable or unacceptable to us personally should not matter. We are seeking a theory that describes the system (moral agents) but not a theory that "proves" that our personal views about morality are correct. We do not dispute that there are other layers of complexity that would be relevant to a unified and applicable theory. For instance, we understand that moral judgements may decouple from actions by coalitions and political calculation, we, however, insist that simplification is necessary in order to make progress towards any quantitatively testable model. We also find it relevant to observe that the data we use is collected anonymously using the web, with opinions devoid of any immediate consequences. This alone considerably weakens any issues regarding inconsistencies due to conformity effects or political calculation.

We have presented results from Monte Carlo numerical methods and analytical approximation schemes such as mean field for a model of interacting agents. These techniques are suited to study the collective or aggregate properties of our model of agents. Drastic changes in collective properties signal phase transitions and the emergence of different regimes of behaviour.

The neural networks of the agents are quite simple. The only way to know if exaggerated simplifications have been made is to compare with data. Even if not useful for heuristic confrontation, models may be useful in their own right as laboratories where we develop intuition about the different methodologies needed to extract information from possible datasets of more complex phenomena of the future. They help in formulating a set of questions that can be addressed experimentally and theoretically. By pointing out their own limitations, current models can bring us closer to more useful models in the future. The neural networks are not supposed to model the brain networks of individuals, but rather the fact that people integrate the different moral dimensions of an issue, weighted by their own views about the importance of each dimension, in order to reach conclusions in an intuitionist way rather than by using a rationalist *if–then* set of rules.

A summary of conclusions about our results should first of all mention what we have not attempted to do. No mention of any evolutionary perspective of how the moral foundations came to be was presented. In particular it seems reasonable to agree about the possible enhanced fitness that may derive from increased cooperation of a more conformist society but this should be central in future discussions. If this is granted, then we have to answer why less conformity promoting cognitive strategies should be present and not have been eradicated by selection. Are liberals just free riders invading a society of authority/loyalty/purity respecting conservatives? Trying to stay aside of semantic interpretations, we suggest an evolutionary reason why cognitive styles compatible with liberal behaviour are found in modern times and have not been purged. Reaction times of the society of agents show that it is consistent to call large δ agents as conservatives, since they have a large equilibration time under changes of the Zeitgeist. On the flip side, small δ is expected to correlate with liberal fast adapting behaviour under the same changes. But this was common knowledge. What is the novelty of conservatives taking more time to re-adapt than liberals? We found that liberals do not correspond to a $\delta = 0$ cognitive style. Not only conservatives, but also liberals are on the ordered side of the phase diagram. But as we approach the disordered phase, critical slowing down sets in. So agents with δ too small will also have large equilibration times and these have probably been eliminated. A compromise between being fast to re-adapt and being more conformist shapes the societies of agents that live in an ever changing environment.

We have shown that different cognitive styles give rise through social interactions to different statistics about the opinion field h. The interactions are represented by a potential that, although it was never intended to claim precise realism, captures several stylized features of human cognitive styles. We have been cautious to allow agents to learn from opposing views. While this may not always occur in human interactions, there are certain windows of time where people acquire their moral values from their social network of interactions. Qualitative information from fMRI and psychological tests about cognitive properties have been used to construct the interactions but future work will have to refine the learning algorithms. Agents in social networks have shown a better agreement with the experimental data than simulations in e.g. square lattices and the model successfully predicts what sort of connectivity is to be expected if the subjacent social network is complex. We feel that it is rather remarkable that from the answers of questionnaires and by postulating certain information exchange mechanisms something about the topology of the social interactions of a society can be inferred.

Many questions are raised. While we are aware of the previous use of the term peer pressure, we have introduced a quantitative definition that might lead to experimental characterization and measurement. This might help deciding whether our concept is useless or not, but it is the nature of experimental work to help decide relevance. We are now looking at a new and richer dataset that also contains data from different countries. With this new data we will be able to control the effects of other variables such as gender and nationality. We also want to infer the α parameter for different countries and to see if we can come with a better interpretation of it. Maybe it will need another name and not peer pressure. At this point we believe that it is still empirically consistent to call it peer pressure.

An interesting consequence of our approach and of the idea of peer pressure is that histograms of opinion fields might change after external threats to a society are detected. From the properties of the model we can predict that the mean of histograms $H(h)$ will increase and variances of the distribution will be reduced after external threats are detected.

Several methodological problems are raised and should be analyzed, among them, the measurement of the peer pressure, the parsing of moral discourse into 5 or more dimensional vectors, the determination of the Zeitgeist vector and time scales of change. Among the theoretical topics, we should approach the problem of semantics and the problem of dressing the different moral dimensions of the model with an interpretation in the language of moral philosophers. Evolutionary considerations will probably guide the process and break the remaining symmetries. We have also neither addressed a possible role of genetic factors influencing cognitive styles nor if the value of δ depends on the agent's environment. For the latter we will have to consider more sophisticated learning algorithms. Understanding evolutionary and cognitive influences behind cultural wars and their mathematical modelling seems to be a reachable goal.

Acknowledgements

We thank Jonathan Haidt for kindly permitting us to have access to his data on moral foundations. We also thank Eytan Domany for the permission to use the SPIN software. This work received financial support from FAPESP (grant 2007/06122-0), CNPq (grant 550981/07-1) and The Center for Natural and Artificial Information Processing Systems at The University of São Paulo (Núcleo de Apoio à Pesquisa CNAIPS-USP). AS's research was partially funded by the DFG training group 1589/1.

References

[1] S. Galam, Sociophysics: A Physicist's Modeling of Psycho-Political Phenomena, Springer, 2012.
[2] S. Galam, Sociophysics: a review of Galam models, Internat. J. Modern Phys. C 19 (2008) 409–440. http://dx.doi.org/10.1142/S0129183108012297.
[3] C. Castellano, S. Fortunato, V. Loreto, Statistical physics of social dynamics, Rev. Modern Phys. 81 (2009) 591–646. http://dx.doi.org/10.1103/RevModPhys.81.591.
[4] C. Borghesi, J.-P. Bouchaud, Of songs and men: a model for multiple choice with herding, Qual. Quant. 41 (2007) 557–568. http://dx.doi.org/10.1126/science.1137651.
[5] A. Chatterjee, M. Mitrović, S. Fortunato, Universality in voting behavior: an empirical analysis, Sci. Rep. 3 (2013) 1049. http://dx.doi.org/10.1038/srep01049.
[6] S. Galam, From 2000 Bush-Gore to 2006 Italian elections: voting at fifty-fifty and the contrarian effect, Qual. Quant. 41 (2007) 579–589. http://dx.doi.org/10.1007/s11135-007-9072-8.
[7] N. Caticha, R. Vicente, Agent-based social psychology: from neurocognitive processes to social data, Adv. Complex Syst. 14 (2011) 711–731. http://dx.doi.org/10.1142/S0219525911003190.
[8] A. Engel, C. van den Broeck, Statistical Mechanics of Learning, Cambridge University Press, 2001.
[9] R. Spears, M. Lea, S. Lee, Deindividuation and group polarization in computer-mediated communication, Brit. J. Soc. Psychol. 29 (1990) 121–134.
[10] J. Haidt, The new synthesis in moral psychology, Science 316 (2007) 998–1002. http://dx.doi.org/10.1126/science.1137651.
[11] W. Kinzel, R. Metzler, I. Kanter, Dynamics of interacting neural networks, J. Phys. A 33 (2000) L141–L147. http://dx.doi.org/10.1088/0305-4470/33/14/101.
[12] R. Metzler, W. Kinzel, I. Kanter, Interacting neural networks, Phys. Rev. E 62 (2000) 2555–2565. http://dx.doi.org/10.1103/PhysRevE.62.2555.
[13] R. Vicente, A. Martins, N. Caticha, Opinion dynamics of learning agents: does seeking consensus lead to disagreement? J. Stat. Mech. Theory Exp. (2009) P03015. http://dx.doi.org/10.1088/1742-5468/2009/03/P03015.
[14] V. Klucharev, K. Hytonen, M. Rijpkema, A. Smidts, G. Fernandez, Reinforcement learning signal predicts social conformity, Neuron 61 (2009) 140–151. http://dx.doi.org/10.1016/j.neuron.2008.11.027.
[15] D.M. Amodio, J.T. Jost, S.L. Master, C.M. Yee, Neurocognitive correlates of liberalism and conservatism, Nat. Neurosci. 10 (2007) 1246–1247. http://dx.doi.org/10.1038/nn1979.
[16] J. Haidt, C. Joseph, Intuitive ethics: how innately prepared intuitions generate culturally variable virtues, Daedalus 133 (2004) 55–66. http://dx.doi.org/10.1162/0011526042365555.
[17] J. Haidt, S. Koller, M. Dias, Affect, culture, and morality, or is it wrong to eat your dog? J. Pers. Soc. Psychol. 65 (1993) 613–628.
[18] J. Haidt, The emotional dog and its rational tail: a social intuitionist approach to moral judgment, Psychol. Rev. 108 (2001) 814–834. http://dx.doi.org/10.1037//0033-295X.108.4.814.
[19] J. Haidt, J. Graham, When morality opposes justice: conservatives have moral intuitions that liberals may not recognize, Soc. Justice Res. 20 (2007) 98–116. http://dx.doi.org/10.1007/s11211-007-0034-z.
[20] J. Haidt, J. Graham, Planet of the durkheimians, where community, authority, and sacredness are foundations of morality, in: J. Jost, A. Kay, H. Thorisdottir (Eds.), Social and Psychological Bases of Ideology and System Justification, Oxford University Press, 2009, http://dx.doi.org/10.2139/ssrn.980844.
[21] L. Kohlberg, Moral Stages: A Current Formulation and a Response to Critics, S Karger Pub, 1984.
[22] C. Gilligan, In A Different Voice: Psychological Theory and Women's Development, Harvard University Press, 1982.
[23] J. Graham, J. Haidt, B. Nosek, Liberals and conservatives use different sets of moral foundations, J. Pers. Soc. Psychol. 96 (2009) 1029–1046. http://dx.doi.org/10.1037/a0015141.
[24] D. Tsafrir, I. Tsafrir, L. Ein-Dor, O. Zuk, D. Notterman, E. Domany, Sorting points into neighborhoods (SPIN): data analysis and visualization by ordering distance matrices, Bioinformatics 21 (2005) 2301–2308. http://dx.doi.org/10.1093/bioinformatics/bti329.
[25] A. Sousa, T. Yu-Song, M. Ausloos, Effects of agents' mobility on opinion spreading in Sznajd model, Eur. Phys. J. B 66 (2008) 115–124. http://dx.doi.org/10.1140/epjb/e2008-00391-6.
[26] J. Miskiewicz, M. Ausloos, Delayed information flow effect in economic systems. An ACP model study, Physica A 382 (2007) 179–186. http://dx.doi.org/10.1016/j.physa.2007.02.018.
[27] M. Newman, G. Barkema, Monte Carlo Methods in Statistical Physics, Oxford University Press, 1999.
[28] R. Albert, A.-L. Barabási, Statistical mechanics of complex networks, Rev. Modern Phys. 74 (2002) 47–97. http://dx.doi.org/10.1103/RevModPhys.74.47.
[29] F. Wang, D.P. Landau, Determining the density of states for classical statistical models: a random walk algorithm to produce a flat histogram, Phys. Rev. E 64 (2001) 056101. http://dx.doi.org/10.1103/PhysRevE.64.056101.
[30] R. Lambiotte, M. Ausloos, J.A. Holyst, Majority model on a network with communities, Phys. Rev. E 75 (2007) 030101. http://dx.doi.org/10.1103/PhysRevE.75.030101.
[31] A.L. Traud, E.D. Kelsic, P.J. Mucha, M.A. Porter, Comparing community structure to characteristics in online collegiate social networks, SIAM Rev. 53 (2011) 526–543. http://dx.doi.org/10.1137/080734315.
[32] A.L. Traud, P.J. Mucha, M.A. Porter, Social structure of Facebook networks, Physica 391 (2012) 4165–4180. http://dx.doi.org/10.1016/j.physa.2011.12.021.
[33] F. de Waal, Good Natured, Harvard University Press, 1996.
[34] L. Katz (Ed.), Evolutionary Origins of Morality, Imprint Academic, 2002.