

Of Songs and Men: a Model for Multiple Choice with Herding

Christian Borghesi¹, Jean-Philippe Bouchaud^{1,2*}

¹ *Service de Physique de l'État Condensé, Orme des Merisiers,
CEA Saclay, 91191 Gif sur Yvette Cedex, France.*

² *Science & Finance, Capital Fund Management, 6-8 Bd Haussmann, 75009 Paris, France.*

(Dated: November 3, 2009)

We propose a generic model for multiple choice situations in the presence of herding and compare it with recent empirical results from a Web-based music market experiment. The model predicts a phase transition between a weak imitation phase and a strong imitation, ‘fashion’ phase, where choices are driven by peer pressure and the ranking of individual preferences is strongly distorted at the aggregate level. The model can be calibrated to reproduce the main experimental results of Salganik et al. (*Science*, 311, pp. 854-856 (2006)); we show in particular that the value of the social influence parameter can be estimated from the data. In one of the experimental situation, this value is found to be close to the critical value of the model.

I. INTRODUCTION

Making decisions is part of everyday life. Some situations require a binary choice (i.e. to vote yes or no in a referendum, to buy or not to buy a cell phone, to join or not to join a riot, etc. [1, 2]). Many others involve multiple options, for example in the first round of French presidential elections (where the number of candidates is typically 15), in portfolio management where very many stocks are eligible, in supermarkets where the number of possible products to buy is large, etc. In most cases, the choice is constrained by some generalized *budget constraint*, either strictly (at most one candidate in the French presidential election) or softly (the total spending in a supermarket should on average be smaller than some amount). It is common experience that people generally do not determine their action in isolation. Quite on the contrary, interactions and herding effects often strongly distort individual preferences, and are clearly responsible for the appearance of trends, fashions and bubbles that would be difficult to understand if agents were insensitive to the behaviour of their peers. Catastrophic events (such as crashes, or sudden opinion shifts) can occur at the macro level, induced by imitation, whereas the aggregate behaviour of *independent* agents would be perfectly smooth.

A relevant challenge in the present era of information economy is to be able to extract faithfully individual opinions/tastes from the publicly expressed preferences under the influence of the crowd. For example, book reviewers on Amazon may be biased by the opinion expressed by previous reviews; if imitation effects are too strong, overwhelmingly positive (or negative) reviews cannot be trusted (see [3]), as a result of “information cascades” [4]. In the case of financial markets, strong herding effects in the earning forecasts of financial analysts have been reported – the dispersion of these forecasts is typically ten time smaller than the *ex post* difference between the forecast and the actual earning (see [5] and refs. therein). These herding effects may lead to a complete divergence between the market price and any putative ‘rational’ price. In the context of scientific publications, the substitution of the present refereeing process by other assessment tools, such as number of downloads from a preprint web-page, or number of citations, is also prone to strong, winner-takes-all, distortions [6, 7]. More generally, it is plausible that such herding phenomena play a role in the appearance of Pareto-tails in the measure of success (wealth, income, book sales, movie attendance, etc.).

Despite their importance, already stressed long-ago by Keynes and more recently by Schelling [1], quantitative models of herding and interaction effects have only been explored, in different contexts, in a recent past, see [2, 4, 8, 9, 10, 11, 12, 13, 14, 15]. This category of models have in fact a long history in physics, where interaction is indeed at the root of genuinely collective effects in condensed matter, such as ferromagnetism, superconductivity, etc. One particular model, that appears to be particularly interesting and generic, is the so-called ‘Random Field Ising Model’ (RFIM) [16], which models the dynamics of magnets under the influence of a slowly evolving external solicitation. This model can be transposed in a socio-economics context [9, 17, 18, 19] to represent a *binary* decision situation under social pressure. A robust feature of the model is that discontinuities appear in aggregate quantities when imitation effects exceed a certain threshold, even if the external solicitation varies smoothly with time. Below this threshold, the behaviour of demand, or of the average opinion, is smooth, but the natural trends can be substantially amplified by

*Electronic address: christian.borghesi@cea.fr;jean-philippe.bouchaud@cea.fr

peer pressure. The predictions of the RFIM can be confronted, with some success, to empirical observations concerning sales of cell phones, birth rates and the terminal phase of clapping in concert halls [19].

Here, we want to generalize the RFIM to *multiple* choice situations. One motivation is that, as mentioned above, these situations are extremely common. A more precise incentive for such a generalization is however the recent publication of a remarkable experimental paper by Salganik, Dodds and Watts [20]. In order to detect and quantify social influence effects, the authors have conducted a careful Web-based experiment (described below) with several quantitative results. Their detailed interpretation begs for a specific model, which we introduce and discuss in this paper and compare with these empirical results. The model is found to fare quite well and allows one to extract from the data a quantitative estimate of the imitation strength, called J below. Interestingly, one of the situations corresponds to a value of J close to the critical point of the model, where collective effects become dominant and strongly distort individual preferences.

II. THE MODEL

We consider N agents indexed by roman labels $i = 1, \dots, N$, and M items indexed by Greek labels $\alpha = 1, \dots, M$. Each agent can construct his ‘shopping list’ or portfolio of items, for simplicity, we restrict here to cases where the quantity of item α is either zero or unity (in the example of movies, we neglect the possibility of going twice to see the same movie). The portfolio of agent i is therefore a vector of size M : $\{n_i^\alpha\}$ with $n_i^\alpha = 0, 1$. The “budget constraint” can in general be written as:

$$B_i^- \leq \sum_{\alpha=1}^M n_i^\alpha \leq B_i^+, \quad (1)$$

where the budget might be different for different agents.

The choices made by agent i are assumed to be determined by three different factors:

- a piece of public information affecting all agents equally, measuring the intrinsic attractivity of item α . This is modeled by a real variable F^α , which may contain, for example, the price of the product (low price means large F^α ’s), or its technological performances, past reputation, etc.
- an idiosyncratic part describing the preferences/tastes of agent i , in the absence of any social pressure or imitation effects. This part is again modeled by a real variable h_i^α , which is positive and large if agent i is particularly fond of item α .
- a social pressure/imitation term which describes how the choices made by *others* affect the perception of item α by agent i . In full generality, we can write this term as:

$$\sum_{j \neq i} \sum_{\beta} J_{j,i}^{\beta,\alpha} n_j^\beta \quad (2)$$

where $J_{j,i}^{\beta,\alpha}$ measures the influence of the consumption of product β by agent j . Positive $J_{j,i}^{\beta,\alpha}$ ’s describe herding-like effects (which could exist across different products), whereas negative $J_{j,i}^{\beta,\alpha}$ ’s are related to contrarian effects (for example, agent j buying item β might push the price of item α up). We will consider in this paper a simplified version of the model where only the aggregate consumption of item α itself influences the value of n_i^α , i.e.:

$$J_{j,i}^{\beta,\alpha} = \frac{JM}{\mathcal{C}} \delta_{\alpha,\beta}, \quad (3)$$

where the factor M is introduced for convenience and \mathcal{C} is the total consumption, defined as:

$$\mathcal{C} = \sum_i \sum_{\alpha} n_i^\alpha. \quad (4)$$

We will also introduce the total consumption of item α as $\mathcal{C}^\alpha = \sum_i n_i^\alpha$, and the relative consumption (or success rate) $\phi^\alpha = \mathcal{C}^\alpha / \mathcal{C}$, with $\sum_{\alpha} \phi^\alpha = 1$.

We assume that the consumption of item α by agent i is effective if the sum of these three determining factors exceed a certain *threshold*, and consider the following update rule for the n_i^α ’s:[27]

$$n_i^\alpha(t+1) = \Theta \left[F^\alpha + h_i^\alpha + JM \left(\phi^\alpha(t) - \frac{1}{M} \right) - b_i(t) \right], \quad (5)$$

where Θ is the Heaviside function, $\Theta(u > 0) = 1$ and $\Theta(u \leq 0) = 0$. In the above equation, we have added a ‘chemical potential’ b_i (borrowing from the statistical physics jargon) which allows the budget constraint to be satisfied at all times [21]. The $-1/M$ term was added for convenience, and makes explicit that it is the consumption of item α *in comparison with its expected average* $1/M$ that generates a signal (see also [22]). It is easy to check that the case $M = 1$, with $J \rightarrow \mathcal{J}C/N$, corresponds to the standard RFIM considered in [19]. Note also that the Θ function describes a deterministic situation: as soon as the total ‘utility’ of item α is positive for agent i , consumption is effective. One could choose a probabilistic situation where $\Theta(u)$ is replaced by a smoothed step function, for example:

$$\Theta_\beta(u) = \frac{1}{1 + e^{-\beta u}}. \quad (6)$$

The limit $\beta \rightarrow \infty$ corresponds to the deterministic rule, to which we will restrict throughout this paper.

In the following, we assume that both F 's and h 's are time independent, and taken from some statistical distributions which we have to specify. Here again, the number of possibilities is very large, and correspond to different situations. We choose the F^α 's as IID random variables (for example Gaussian), with mean m_F and variance Σ_F^2 . The mean m_F describes the average intrinsic attractiveness of items – for example, a large overall inflation would lead to a negative m_F . The dispersion in quality of the different items is captured by Σ_F . More realistic models should include some sort of ‘sectorial’ correlations between the F^α 's.

As for h_i^α 's, we posit that they can be decomposed as $h_i^\alpha = \bar{h}_i + \delta h_i^\alpha$, where \bar{h}_i describes the propensity of agent i for consumption (‘compulsive buyers’ correspond to large positive \bar{h}_i 's), whereas δh_i^α correspond to the idiosyncratic tastes of agent i , defined to have zero mean. For simplicity, we again assume that both \bar{h}_i 's and δh_i^α are IID; without loss of generality we can assume that the average (over i) of \bar{h}_i is zero (a non zero value could be reabsorbed into m_F). The variance of \bar{h}_i is Σ^2 and that of δh_i^α is σ^2 . Since in the limit $\beta \rightarrow \infty$ considered in this paper the overall scale of the fields is irrelevant, we can choose to set $\sigma \equiv 1$. One could also add explicit time dependence, for example choosing m_F to be an increasing function of time, to describe a situation where the average propensity for consumption increases with time.

The model as defined above is extremely rich and its detailed investigation as a function of the different parameters and budget constraints will be reported in a forthcoming publication. The most interesting question about such a model is to know whether the realized consumption is *faithful*, i.e. whether or not the actual choice of the different items reflects the ‘true’ preferences of individual agents, as would be the case in the absence of interactions ($J = 0$). Based on the RFIM, we expect that this will not be the case when J is sufficiently large, in which case strong distortions will occur, meaning that the realized consumption will (i) violate the natural ordering of individual preferences and (ii) become history dependent: a particular initial condition determines the ‘winners’ in an irreproducible and unpredictable way. In order to characterize the inhomogeneity of choices, the authors of [20] have proposed and measured different observables, in particular:

- The Gini coefficient G , defined as:

$$G = \frac{1}{2M} \sum_{\alpha, \beta} |\phi_\alpha - \phi_\beta|, \quad (7)$$

which is zero if all items are equally chosen, and equal to $1 - 1/M$ if a unique item is chosen. The Gini coefficient is a classic measure of inequality. In fact, a more relevant measure of interaction effects is the ratio G/G_0 , where G_0 is the Gini coefficient for $J = 0$.

- The unpredictability coefficient U , defined as:

$$U = \frac{1}{M \binom{W}{2}} \sum_{\alpha=1}^M \sum_{k=1}^W \sum_{\ell < k}^W |\phi_{(k)}^\alpha - \phi_{(\ell)}^\alpha| \quad (8)$$

where the indices k, ℓ refer to W different ‘worlds’, i.e. different realizations of the model *with the very same* F^α 's but a different set h_i^α 's (chosen with the same distribution) or different initial conditions. In the limit of a large population ($N \rightarrow \infty$), it is easy to show that $U = 0$ when $J = 0$, since the ϕ^α only depends on the F^α 's. A non zero value of U , on the other hand, reveals that it is impossible to infer from the intrinsic quality of the items the aggregate consumption profile (strong distortion).

- A more detailed information is provided by the scatter plot of ϕ^α versus $\phi^\alpha(J = 0)$; for J small one expects a nearly linear relation, whereas for larger J the points acquire a larger dispersion and the average relation becomes non-linear, indicating a substantial ‘exaggeration’ of the consumption of slightly better items.

We have studied these quantities both numerically and analytically within the above model. We present below some of our numerical results, and compare them with the empirical results of [20]. Our most important analytical result is the existence of a critical value J_c , below which the unpredictability U is strictly zero in the limit $N \rightarrow \infty$, and becomes positive for $J > J_c$, growing as $U \sim J - J_c$ close to the transition. The fluctuations of U diverge close to J_c , as for standard second order phase transitions. The value of J_c can be computed exactly in the limit of a large number of items $M \gg 1$, and depends on the detailed shape of the distribution of the fields F and h . More precisely, J_c is given by:

$$J_c = \int_{-\infty}^{\infty} dF P_F(F) \gamma(F), \quad (9)$$

where $\gamma(F)$ is the solution of:

$$\gamma = \int_{J_c - F - \gamma}^{\infty} du \frac{P_h(u)}{P_h(0)}, \quad (10)$$

and P_F and P_h are the distributions of the fields F and h .

III. THE WEB-BASED EXPERIMENT OF SALGANIK ET AL.

Here we describe the beautiful experimental set-up of M. J. Salganik, P. S. Dodds and D. J. Watts [20], which allows them to conclude that social influence has a determinant effect on the choices of individual agents. In the next section, we will in fact use their quantitative results to measure, within the above theoretical framework, the strength of the social influence factor J . Salganik et al. have [20] created an artificial ‘‘music market’’ on the web with $M = 48$ songs from essentially unknown bands in which 14,341 (mostly teen-agers) participated. Songs are presented in a screen and participants make decisions about which songs to listen to, and in a second step, whether they want to download the song they listened to. Participants are randomly assigned to one of the three following situations:

- an independent (zero-influence) situation where the list of songs carries no mention of the songs downloaded by other participants. This situation allows to define a benchmark, where an ‘intrinsic’ mix between the quality of the songs and the preference of the participants can be measured. This situation corresponds to $J = 0$ in the model above;
- a ‘weak’ social influence situation. In this case, the number of times a given song has been downloaded by other participants is shown. However, the songs are presented in random order so that the ranking of the preference of other participants is not obvious at first glance. This situation corresponds to a certain small value $J_1 > 0$ in the model above;
- a ‘strong’ social influence situation. In this case, the list of songs is presented by decreasing number of downloads, such as to emphasize the preferences expressed by previous participants. This situation corresponds to a certain value $J_2 > J_1 > 0$ in the model above.

Furthermore, in both social influence conditions participants are randomly assigned to $W = 8$ different worlds, each one with its own history and evolving independently from one another, but with the same initial conditions, i.e. zero downloads. For each of the two influence conditions, the outcomes (i.e. the number of downloads of all songs) are compared to the independent, zero-influence situation. In this way, the authors are able to conclude that increasing the strength of social influence increases both the inequality G and the unpredictability of success U [20].

Because these experiments look very much like those in physical laboratories, we believe that they could play an important role in the development of scientific investigations of collective human behavior. The Web gives the opportunity to devise and perform large scale experimentation (see also [23]), with a number of participants that allows one to extract meaningful statistical information. We expect that many other experiments of the same type will be conducted in the future. In the present case, the experiment is very carefully thought through to remove many artefacts: for example, download is free (no consideration of the wealth of participants is required – no ‘budget constraint’) and anonymous (no direct social pressure is involved); participants are not rewarded to have made a ‘good’ or ‘useful’ choice, songs and bands are not well known (avoiding strong a priori biases), etc.

IV. MODEL CALIBRATION: TOWARDS A MEASUREMENT OF SOCIAL INFLUENCE?

We now turn to a semi-quantitative analysis of the empirical data collected by Salganik et al. [20]. Once the distribution of F^{α} ’s and h_i^{α} ’s are fixed (we chose them to be Gaussian for simplicity), the model depends on four

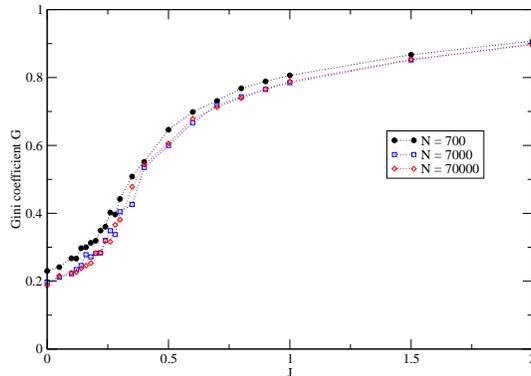


FIG. 1: Gini coefficient as a function of J for the choice $m_F \approx -2$, $\Sigma_F \approx 0.2$, $\Sigma = 1$, and for different number of agents $N = 700, 7000$ and 70000 . Note the rather weak dependence on N of this quantity. The empirical values of G in the three different situations are: $G_0 \approx 0.22$ (no imitation), $G_1 \approx 0.35$ (weak imitation) and $G_2 \approx 0.5$ (strong imitation).

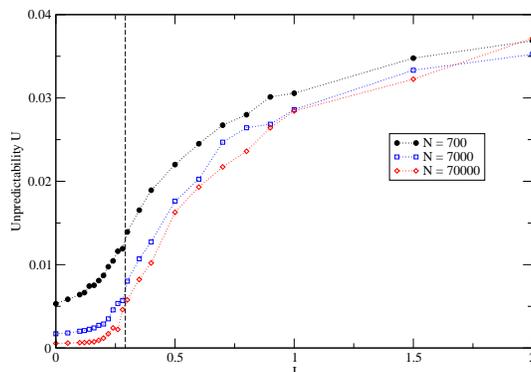


FIG. 2: Unpredictability U as a function of J for the choice $m_F \approx -2$, $\Sigma_F \approx 0.2$, $\Sigma = 1$, and for different number of agents $N = 700, 7000$ and 70000 . In this case, the finite size effects are strong; one in fact expects U to be zero for $J < J_c \approx 0.29$ (dashed vertical line), and to grow linearly for small $J - J_c > 0$. The empirical values of U for $N = 700$ and in the three different situations are: $U_0 \approx 0.0045$ (no imitation), $U_1 \approx 0.008$ (weak imitation) and $G_2 \approx 0.013$ (strong imitation). This last case corresponds, for $N = 700$, to $J_2 \sim J_c$.

parameters: m_F, Σ_F, Σ and the social influence J . These values must be chosen as to reproduce the observations reported in [20], namely:

- The Gini coefficient G_0 , the unpredictability U_0 and the qualitative shape of the distribution of ϕ_0^α in the independent situation, corresponding to $J = 0$.
- The Gini coefficient G , the unpredictability U and the qualitative shape of the relation between ϕ^α and ϕ_0^α in the social influence conditions

Quite a lot more data is reported in the supplementary material of [20], for example the average number of downloaded songs per participant $d = \mathcal{C}/N$. In fact, the situation of [20] is slightly more complicated than assumed in the above model because each participant makes a two-step decision. Participants, before possibly downloading a song, first choose to listen to it. These two decisions may be correlated and both influenced by the choice of other participants. The authors of [20] report separate statistics for the number of downloaded songs and the number of ‘tested’ songs. In order to reproduce these results in full detail, one must generalize the above model, for example by assuming that the number of downloads of song α by agent i can be written as:

$$n_i^\alpha(t+1) = \Psi_i^\alpha \Theta \left[F^\alpha + \bar{h}_i + \delta h_i^\alpha + JM \left(\phi^\alpha(t) - \frac{1}{M} \right) \right], \quad (11)$$

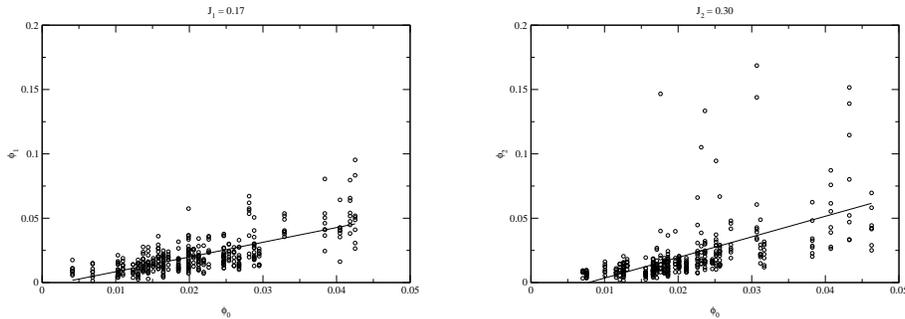


FIG. 3: Scatter plot of the realized preferences $\phi^\alpha(J)$ as a function of the ‘intrinsic’ preferences ϕ_0^α , in the weak social influence condition ($J_1 = 0.17$, left), and in the strong social influence condition ($J_2 = 0.30$, right), all for $m_F \approx -2$, $\Sigma_F \approx 0.2$, $\Sigma = 1$. Lines are linear regressions. These plots compare well with the corresponding plots of [20] (Figs. 3-A and 3-C; we use here the same scale as in [20]).

where $\Psi_i^\alpha = 1$ with probability p^α and 0 otherwise describing the decision of actually downloading a song after listening to it. Although the inclusion of this second decision step is crucial to account fully for the results of [20], we neglect this aspect altogether in the present paper and refer the reader to a later, more detailed publication [21]. Here we want to show that the main empirical features can indeed be reproduced by the model.

Different choices of m_F, Σ_F, Σ are in fact compatible with the observations corresponding to $J = 0$, for which Salganik et al. find $G_0 \approx 0.22$ and $U_0 \approx 0.0045$ (for a number of participants in each ‘world’ of $N = 700$, the value we also use in our numerical simulations). A possible choice (further justified in [21]) is: $m_F \approx -2$, $\Sigma_F \approx 0.2$, $\Sigma = 1$. The resulting shape of the distribution of ϕ_0^α is found to be compatible with the data of [20]. Note that $\Sigma_F^2 = 0.04 < \Sigma^2 + \sigma^2 = 2$, suggesting that the intrinsic quality of songs is less dispersed than the preference of agents. This is expected in a situation where songs and bands are unknown, leading to very small *a priori* information on their intrinsic quality.

Now, it is interesting to see how G and U are affected by a non zero value of J – cf. Figs. 1 and 2. From these plots, one sees that the ‘weak’ social influence situation, characterized by $G_1 \approx 0.35$ and $U_1 \approx 0.008$ [20], corresponds to $J_1 \approx 0.17$. On the other hand, the ‘strong’ influence situation yields $G_2 \approx 0.5$ and $U_2 \approx 0.013$ [20], which we can account for by setting $J_2 \approx 0.30$. The scatter plots of ϕ^α vs. ϕ_0^α are shown in Figs 3-a and 3-b and can be satisfactorily compared to Figs. 3-A and 3-C of [20].

It is of particular interest to compare the above values of J_1 and J_2 to the critical value J_c of the model, which can be determined exactly as a function of m_F, Σ_F, Σ in the limit $M \rightarrow \infty$ [21]. In the present case, we find $J_c \approx 0.29$, such that, in the limit $N \rightarrow \infty$, $U(J < J_c)$ should be strictly zero. As expected on general grounds and shown in Fig. 2, the value of U at finite N suffers from large finite size effects. Only a careful extrapolation for $N \rightarrow \infty$ allows one to confirm the existence of a critical value J_c [21]. But in any case, the value J_2 accounting for the data in the ‘strong’ influence situation is indeed quite large, since it corresponds to the critical region where imitation effects become dominant.

Another effect worth noticing is the dependence of the average number of downloaded songs d (or consumption $\mathcal{C} = Nd$) on the imitation parameter J , predicted by the model and reported in Fig 4. We see that this quantity has a clear *maximum* as a function of J : at first, imitation effects tend to increase the total consumption until $J \sim 1$, beyond which over-polarisation on a small number of items become such that the total consumption goes back down. This might have interesting consequences for marketing policies, for example (see e.g. [24, 25]). The increase of the d with J is actually not observed in [20]; see [21] for a further discussion of this point.

V. CONCLUSIONS

We have proposed a generic model for *multiple* choice situations with imitation effects and compared it with recent empirical results from a Web-based cultural market experiment. Our model predicts to a phase transition between a weak imitation phase, in which expressed individual preferences are close to their value in the absence of any direct social pressure, and a strong imitation, ‘fashion’ phase, where choices are driven by peer pressure and the ranking of individual preferences is strongly distorted at the aggregate level. The model can be calibrated to reproduce the main experimental results of Salganik et al. [20]; we show in particular that the value of the social influence parameter can be estimated from the data. In one of the experimental situation, this value is found to be close to the critical value

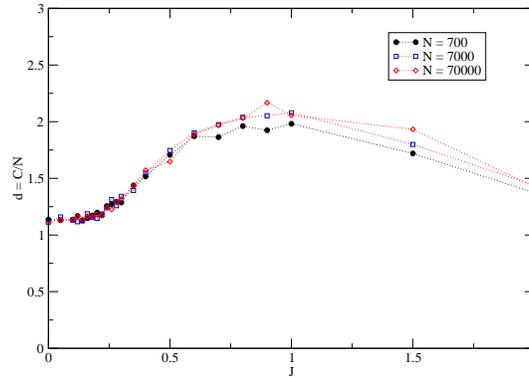


FIG. 4: Average number of downloaded songs d (or consumption $C = Nd$) as a function of J for the choice $m_F \approx -2$, $\Sigma_F \approx 0.2$, $\Sigma = 1$, and for different number of agents $N = 700, 7000$ and 70000 . Finite size effects are quite small in this case. Note the clear maximum of this quantity as a function of the imitation strength J .

of the model, confirming quantitatively that social pressure are strong in that case. This concurs with the conclusions of [19], who also found near critical values of the social influence parameter.

Our model can be transposed to many interesting situations, for example that of industrial production, for which one expects a transition between an archaic economy dominated by very few products and a fully diversified economy as the dispersion of individual needs becomes larger. We leave the investigation of these questions, and the detailed analytical investigation of our model, for a further publication. We believe that the simultaneous development of theoretical models and detailed, rigorous experiments in the vein of [20] or [23, 26], will help promoting a quantitative understanding of collective human (and animal) behaviour.

Acknowledgments

C.B. thanks Bertrand Roehner for useful conversations and for his enlightening efforts to extend physical intuition towards social sciences. J.P.B thanks Raphael Hitier for bringing ref. [20] to his attention. We also warmly thank Matteo Marsili for many discussions over the years, and for hospitality during completion of this work.

-
- [1] T. Schelling, *Micromotives and Macrobehaviour*, W W Norton & Co Ltd (1978)
 - [2] M. Granovetter, *Threshold models of collective behaviour*, Am. J. Sociol. **83**, 1420 (1978); M. Granovetter, R. Soong, *Threshold models of diffusion and collective behaviour*, J. Math. Socio. **9**, 165 (1983); M. Granovetter, R. Soong, *Threshold models of interpersonal effects in consumer demand*, J. Econ. Behavior & Organization, **7** 83 (1986); M. Granovetter, R. Soong, *Threshold models of diversity: Chinese Restaurants, Residential Segregation and the spiral of silence*, in Sociological Methodology, C. Clogg Edt. (1988), p. 69-104.
 - [3] F. Slanina, Y. C. Zhang, *Referee Networks and their spectral properties*, Acta Physica Polonica B **36**, 2797 (2005)
 - [4] S. Bikhchandani, D. Hirshleifer, I. Welch, *A theory of fads, fashions, custom and cultural changes as informational cascades*, Journal of Political Economy, **100** 992 (1992).
 - [5] O. Guedj, J.P. Bouchaud, *Experts Earning forecasts, bias, herding and gossamer information*, J. Theor. Appl. Finance **8**, 933-946 (2005).
 - [6] S. Redner, *Citation Statistics From More Than a Century of Physical Review*, Physics Today, p. 49, June 2005
 - [7] M. V. Simkin and V. P. Roychowdhury, *Copied citations create renowned papers?* cond-mat/0310150; *Stochastic modeling of citation slips*, cond-mat/0401529.
 - [8] H. Föllmer, *Random Economies with many interacting agents*, J. Math. Econ. **1**, 51 (1974).
 - [9] S. Galam, *Majority rule, hierarchical structure and democratic totalitarianism*, J. Math. Psychol. **30**, 426 (1986); S. Galam, S. Moscovici, *Towards a theory of collective phenomena: Consensus and attitude changes in groups*, Euro. J. Social Psy. **21**, 49 (1991); S. Galam, *Rational group decision making: a RFIM at $T = 0$* , Physica A **238**, 66 (1997).
 - [10] A. Kirman, *Ants, rationality and recruitment*, Quarterly Journal of Economics, **108** 137 (1993)
 - [11] A. Orléan, *Bayesian interactions and collective dynamics of opinions*, J. Econ. Behav. Organ. **28**, 257 (1995)
 - [12] R. Cont, J.P. Bouchaud, *Herd behaviour and aggregate fluctuations in financial markets*, Macroeconomic Dynamics **4**, 139 (2000).

- [13] D. Challet, M. Marsili, Y.C. Zhang, *Minority Games*, Oxford University Press, 2005.
- [14] Ph. Curty, M. Marsili, *Phase coexistence in a forecasting game*, J. Stat. Mech. (2006) P03013
- [15] for a recent review: A. de Martino, M. Marsili, *Statistical mechanics of socio-economic systems with heterogeneous agents*, physics/0606107.
- [16] J. Sethna, K. Dahmen, C. Myers, *Crackling Noise*, Nature, **410**, 242 (2001); J. Sethna, K. Dahmen, O. Perkovic, *Random Field Ising Models of Hysteresis*, cond-mat/0406320
- [17] J.-P. Bouchaud, *Power-laws in economics and finance: some ideas from physics*, Quantitative Finance, **1**, 105 (2001); J.-P. Bouchaud and M. Potters, *Theory of Financial Risks and Derivative Pricing*, Cambridge University Press, 2003.
- [18] J.-P. Nadal, D. Phan, M. B. Gordon and J. Vannimenus, *Multiple equilibria in a monopoly market with heterogeneous agents and externalities*, Quantitative Finance **5** 557 (2005); M. B. Gordon, J.-P. Nadal, D. Phan and J. Vannimenus, *Seller's dilemma due to social interactions between customers*, Physica A **356** 628 (2005).
- [19] Q. Michard, J.-P. Bouchaud, *Theory of collective opinion shifts: from smooth trends to abrupt swings*, Eur. J. Phys. B **47**, 151-159 (2005)
- [20] M. J. Salganik, P. S. Dodds, D. J. Watts, *Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market*, Science **311**, 854-856 (2006)
- [21] Ch. Borghesi, J. P. Bouchaud, in preparation.
- [22] Ch. Borghesi, S. Galam, *Chaotic, staggered and polarized dynamics in opinion forming: the contrarian effect*. Phys. Rev. E **73**, 066118 (2006).
- [23] P. Laureti, P. Ruch, J. Wakeling, Y. C. Zhang, *The Interactive Minority Game: a Web-based investigation of human market interactions*, Physica A **331**, 651 (2004)
- [24] F. M. Bass, *A new product growth model for consumer durables*, Management Science, **15**, 215 (1969).
- [25] A. Steyer, *Géométrie des interactions sociales et modèles de diffusion des innovations*, working paper, GREQAM (2002).
- [26] A. Cavagna, et al., *The starflag project*, in preparation.
- [27] We neglect $1/N$ corrections here.