Agent-Based Modeling of Cascades on Social Networks

By
Anton Tayanovskyy

Submitted to
Central European University
Department of Economics

In partial fulfillment of the requirements for the degree of
Master of Economics

Supervisor: Professor Péter Medvegyev

Budapest, Hungary
2007

Abstract

Patterns of conformity in agent behavior, known as cascades, are observed in many areas including technology diffusion, evolution of social norms, fashion, politics, and multimedia. These patterns are consistently found to be unpredictable and fragile. Traditional explanations in the economics literature focus on the micro-level and include rational herding and observational learning. In contrast, statistical mechanics has a macro-level methodology that deals with similar phenomena. This paper seeks to integrate the two approaches. Cellular automata are adjusted to become closer to real social networks in their topology and equipped with a probabilistic update rule compatible with economic cascade models. The resulting flexible agent-based framework is used to simulate cascades and reproduce their key features, including fragility.

Contents

List of Figures

Introduction

Social cascades are patterns of conformity in populations of decision makers. Examples of cascades are consumers buying the same product, firms adopting the same technology, citizen voting for the same political party or investors holding the same asset. Cascades have both a space and a time dimension. Agents who are close according to such measures as geographic location, industry, income, or religion are more likely to choose similarly, therefore cascades appear as cohesive clusters on a social map. Every cascade also evolves in time, with some agents joining and some abandoning it.

It is precisely the dynamics that make cascades fascinating. In one word, they are unpredictable. Typically a large cascade appears unexpectedly, grows rapidly and appears impregnable during its lifetime, but is in fact brittle and can suddenly and unpredictably dissolve. Frequently cited examples include stock market crashes and financial bubbles, but in fact a lot more phenomena display similar regularities at different size and time scales. Cascades of interest to economics are observed not only in financial markets, but also in the evolution of norms, the diffusion of innovations, fashion cycles, and generally in all markets with differentiated products.

What makes cascades so fragile? One of the first papers to pose this question is the now seminal work by Sushil Bikhchandani, David Hirshleifer, and Ivo Welch (1992). In their model imperfectly informed agents facing a decision can gain by observing and imitating the choices of their predecessors. Cascades, when they form, are fragile because rational agents understand that the observed conformity results from imitation rather than from multiple independent signals that the prevalent option is the best. Because of the ubiquity of cascade phenomena and the elegance of this explanation the paper quickly attracted the attention of the economic community, and, together with two related works (Abhijit Banerjee, 1992; David Scharfstein and Jeremy Stein, 1990), stimulated research in what is now known as the rational herding literature.

Unfortunately, the rational herding literature has focused on studying efficiency of decisions in cascades and neglected another prominent aspect explained by Bikhchandani et al. (1992), the localized conformity. Their original explanation hinges on the assumption that agents decide sequentially and observe all predecessors. A much more appealing approach is to model interaction and observation as happening between neighbors in a social network, but because of the lack of understanding and data of the social structure this has not been feasible until recently. But access to larger network databases such as author collaborations in scientific journals or the Internet infrastructure has led to groundbreaking discoveries in network topology (Réka Albert and Albert-László Barabási, 2002). A decade after the rational herding literature was born, Duncan J. Watts (2002) utilized these discoveries in a simple social network model where a basic local interaction rule produced cascades with both fragility and local conformity.

The network explanation differs from rational herding in that it traces cascade formation, local conformity and fragility to the macroscopic topology of interactions rather than to their microscopic specifics. These are, in fact, two different approaches. The questions critical to understanding cascades are how these approaches relate, whether they can be integrated, and what the minimal set of assumptions needed to produce cascades is.

This paper argues that for cascades to happen it is sufficient that a social system has three characteristics: local interactions, incentives to imitate, and random deviation. If these conditions are satisfied, social cascades can be understood and modeled within the analytical framework of the frustrated systems theory from physics. In this way the two approaches can be integrated: the micro approach, such as rational herding, explains the origin of imitation and deviation, while the macro approach uses these factors together with localized interactions to explain cascade formation and dynamics. Originally developed to study magnets, spin-glasses, and other systems with a phase transition, frustrated systems theory provides excellent macroscopic explanations for fragility and local conformity. In addition, it predicts that cascade sizes near a phase transition have a scale-free distribution, which is indeed observed in many economic applications.

The research methodology adapts computational models from physics. In particular, frustrated spin-glass models from statistical mechanics are found to be particularly suited to social cascade analysis, especially when their interaction structure is corrected. The correction amounts to using social networks instead of lattices, for which purpose the Watts-Strogatz network generation model is employed (Duncan J. Watts and Steven H. Strogatz, 1998). The simulations of the resulting model demonstrate that it is capable of illustrating cascade formation, fragility, local conformity, and other regularities. Because of the limitations of the software implementation the simulations are carried out on a modest scale, but the results nevertheless suggest that some of the classical results obtained on lattices can indeed be generalized to social networks.

The rest of this paper is organized as follows. Chapter 1 reviews the economic research on social cascades, highlighting those dynamic regularities that are discovered across applications. These are explained in Chapter 2 within the analytical framework of the frustrated systems theory. The chapter develops arguments for applying this theory to economics, discusses prior applications, and presents findings from social network topology research. These findings are used in Chapter 3 to construct an agent-based model which is in fact an adapted spin-glass model. The same chapter presents the simulation results and discusses the limitations of the implementation. The paper concludes with the broad implications of this study and suggestions for further research.

Chapter 1: Literature Review

The purpose of this chapter is to review the economic studies of cascades, emphasizing two findings that justify the modeling approach taken by this paper. The first finding is the existence of incentives to imitate, which the approach assumes; the second one is a set of stylized facts about cascades, which it is able to explain.

The message of the economic literature on cascades is that conformity in behavior arises when agents have incentives to behave similarly. There are two explanations of such incentives. The first explanation, found in the rational herding literature, is that imperfectly informed agents can make better choices by observing the actions of others. The second one is network effects, which are studied in the technology adoption literature and arise when technologies increase in value as their user base grows.

Besides explaining the origin of imitation, research on cascades also studies their dynamics. Although the perspectives vary, several aggregate regularities are consistently discovered. Cascades are found to be:

1. formed by the interplay of imitation and deviation;
2. fragile and tending to dissolve suddenly;
3. difficult to predict;
4. influenced by the topology of agent interactions;
5. clustered according to that topology;
6. growing along an S-curve;
7. unequal in size, with the size distribution following a power law;
8. succeeding one another, which produces punctuated equilibria;

These regularities are found in both theoretical and empirical studies, experiments, and computational models. The present chapter reviews this research, and the following one shows how these stylized facts connect social cascades to a wider class of phenomena.

1.1 Rational Herding and Information Cascades

Rational herding models date back to the three seminal papers by Bikhchandani et al. (1992); Banerjee (1992); Scharfstein and Stein (1990). Imitation arises in a queue of imperfectly informed individuals making a decision and trying to learn from the observed choices of predecessors. A cascade obtains when agents rationally choose to ignore their private information and follow the crowd. The motivation behind these models is mixed: they are used to study efficiency of social learning and fragility of conforming group behavior.

Rational herding outcomes are fragile. Specifically, although during a cascade most agents join the herd, a sufficiently confident agent with a strong private signal can deviate, possibly creating a different herd. Deviation is possible since an agent can deduce that observed conformity of choices does not imply a conformity of independent private signals, but simply imitation. Also, behavior is highly unpredictable as it depends on the choice of the first agents (Bikhchandani et al., 1998). This low information content of conforming behavior and initial unpredictability are used to explain such phenomena as financial market bubbles and diffusion of technologies.

Much effort has been made to generalize rational herding by lifting the assumptions about sequential decision making, full history observability, and agent homogeneity. These efforts revealed that herding probability, dynamic stability, efficiency (correct learning) and predictability are all highly sensitive to these assumptions.

The easiest assumption to lift is the full history observability. Glenn Ellison and Drew Fudenberg (1995) substitute it with random sampling of history, arriving at a counter-intuitive result that increased social information leads to less efficient outcomes. Lones Smith and Peter Sorensen (1998) also use random sampling, obtaining complex dynamics, where welfare may decrease as well as increase over time. In contrast, Bogachan Celen and Shachar Kariv (2004) study a model where agents can observe only the immediate predecessor. They find the stochastic process of choices to be unstable, generating increasingly longer periods of conformity with punctuated equilibria.

As to agent heterogeneity, a highly sophisticated analysis of the issue is found in

Smith and Sorensen (2000). The authors discover an inefficient stable equilibrium they call *confounded learning*, where incoming signals cease being informative. In a different framework motivated by behavioral studies, Kariv (2005) includes agent overconfidence into the model, showing that the presence of overconfident agents makes the cascades more likely and less fragile.

Other assumptions are important as well. Complications are found by Ivan Pastine and Tuvana Pastine (2005), who show that even small departures from symmetry of signals greatly diversify the system dynamics. And as an alternative to sequential decision making, Nicolas Melissas (2005) allows agents to wait before others decide, concluding that herding, though inefficient in itself, is a remedy to the greater evil of strategic delay.

The rational herding literature was a great contribution to economics by pioneering the study of social cascades, and it provided a plausible explanation for incentives to imitate. However, because of the sensitivity to assumptions, the original setup is too specific to offer a general explanation for cascades, while later extensions lead to complex dynamics that are difficult to interpret.

1.2 Observational Learning in Social Networks

There is another relevant extension to the rational herding setup which lifts the full observability assumption in favor of a social network where only the neighbors' decisions can be observed. This research direction is reviewed in Sanjeev Goyal (2003). While there is agreement on the fact that network topology matters, the extreme complication of the problem led to a divergence in methodology. Asymptotic action convergence is derived for general networks under bounded (Venkatesh Bala and Sanjeev Goyal, 1998) and full (Douglas Gale and Shachar Kariv, 2003) rationality, with finite-time dynamics obtained by numerical simulations (discussed below). A similar result is derived by H. Peyton Young (2000), who studies a more general context and shows that sufficiently clustered networks stabilize in the limit. There are attempts to explain network formation by the strategic interaction of agents. Troy Tassier and Filippo Menczer (2001) build a network model of a labor market that generates a highly connected and clustered network, and

Sanjeev Goyal and Fernando Vega-Redondo (2004) in a more general study suggest that the long-run equilibrium shape is a star network.

These models are relevant to the present work because they choose a network approach and therefore encounter the same difficulties. But as these models relate to the social learning literature (Gale, 1996) they tend to focus on outcome efficiency, which is out of scope of this paper. Focusing on cascades, using a simpler interaction model, borrowing results from social network topology and adapting dynamical models from physics allows the present paper to avoid many of the complications of this literature.

1.3 Experimental and Empirical Evidence

The experimental and empirical results demonstrate that imitative behavior is not uncommon in practice, suggesting that it often is the rational response to uncertainty. Specifically, despite the sensitivity to their assumptions, rational herding models correctly describe behavior when these assumptions hold. An example from a large body of experimental literature is found in Mathias Drehmann, Jorg Oechssler, and Andreas Roider (2004). In this remarkably large experiment, 6,000 subjects participated in a cascade simulated over the Internet, and their behavior was found to generally confirm the theoretical predictions. The authors also find evidence that network effects amplify herding. In another experiment Clemens Oberhammer and Andreas Stiehler (2001) allowed users to buy information and deduced their belief structure from the prices. Although the revealed preferences are not consistent with Bayesian updating assumed by the model, the behavior is.

Rational herding was also used to explain empirical patterns in software adoption over the Internet (Wenjing Duan, Bin Gu, and Andrew B. Whinston, 2006). The situation is close to the original model setup because users have access to software download ratings. Using data from Download.com (CNETD), the authors show that it is the ratings and not user reviews that determine future downloads. A similar research (Xinxin Li and Lorin Hitt, 2004) analyzes the effect of user ratings on sales in a dataset obtained for a random sample of books on Amazon, finding that user ratings do affect sales positively.

Cascades on social networks were also studied empirically, primarily because artificial cascade creation, or viral marketing, has received attention from the business world as a means of targeted advertising. For instance, Jure Leskovec, Lada A. Adamic, and Bernardo A. Huberman (2006) analyze a huge dataset to study the effect of product recommendations on sales. Their research reveals that care should be taken in interpreting network data in such models. In particular, it is found that repeated use weakens the links, probability of converting has a saturation point, and highly connected individuals may not be highly influential. These insights are taken into account in the present paper.

1.4 Technology Adoption

Technology adoption is one of the best studied instances of behavior cascades. It demonstrates the stylized facts of S-shaped growth, high uncertainty, and switching equilibria. The technology adoption literature is also important in providing a complementary explanation for imitative behavior. While in rational herding imitation arises solely from uncertainty about the choice being made, in technology adoption this motive may be coupled with network effects, which arise from technical complementarities and imply that the utility of a technology is increasing in the number of other users.

Traditional diffusion models (Frank M. Bass, 1969) assume random interaction and a contagion metaphor, where agents are "infected," or adopt a new technology, as soon as they learn about it. This fits into the cascade framework as a limit case of observational learning in a random network. The predicted S-shaped growth pattern has been consistently observed in most empirical studies. However, this evidence does not prove the Bass model to be correct. For example, Young (2006) suggests an alternative model that allows for heterogeneity of agents, and predicts a similar (though steeper) growth pattern. The accumulated evidence simply implies that any plausible explanation of technology adoption must account for S-shaped growth (and cascades do, as will be shown later).

Including uncertainty and network effects (Joseph Farrell and Garth Saloner, 1985) greatly complicates the model predictions. For example, Christoph H. Loch and Bernardo A. Huberman (1999) study a dynamic model of binary technological choice with these two elements, and obtain unpredictability and switching equilibria. They also resort to numerical simulations in the full version of their model, for it is too complicated for an analytical solution. Jay Pil Choi (1997) studies technological adoption in the rational herding framework, allowing for strategic delay of adoption. He finds that network effects amplify herding in the simple case and lead to delay in the strategic case.

While most models assume random or sequential interaction of agents, some empirical work has suggested that the social network structure is significant. Catherine Tucker (2005) analyzes a dataset on the adoption of a videoconferencing technology in a large firm, finding network effects that are purely local, influencing employees' decisions only when these employees do communicate (are neighboring nodes in the network representation). In a very different setting, Oriana Bandiera and Imran Rasul (2002) study the adoption of a new crop, sunflower, in Northern Mozambique, finding that social information is a highly significant factor in this process. They also discover a peculiar inverse-U relationship between the fraction of adopters in the social neighborhood of a farmer and the probability that she will adopt the new crop, claiming that it arises from strategic delay.

1.5 Computational Models

Agent-based computational models allow researchers to study the dynamics of those cascade and technology adoption systems that have no closed-form solutions. Lacking the elegance of analytical results, computational models nevertheless capture the stylized facts of cascade unpredictability, dynamic instability, and power law size distributions. They also stress the importance of network topology. Typically agents are arranged on cellular grids, but where this assumption is scrutinized it is found that topology significantly influences the model dynamics.

Bala and Goyal (1998) simulate a technology adoption model and obtain S-curve growth and spatial clusters on a 2d lattice. Their paper discusses two topological factors important for the limiting learning outcomes. These are the existence of a "royal family," a set of agents who are observed by everyone else, and local independence of agents, where two agents are defined as locally independent unless they observe the same set of neighbors.

An example of a larger model is found in Robert Axtell (1999), who studies firm formation as a choice of agents to join a firm or start their own. The outcomes most relevant to the present discussion are the fast growth and then almost a vertical collapse of the largest firm, followed by a stable equilibrium with a power law distribution of firm sizes. Being aware that power law distributions are ubiquitous in many simpler physical systems, the author substitutes the complex optimization rules with a randomized decision making heuristic, and shows that the power law does not arise in this setting.

Arthur De Vany and Cassey Lee (2001) develop an agent-based cascade model of motion picture box office revenues. The traditional model is extended to multi-product settings, and the agents are prescribed a behavior heuristic of choosing between no movie, the most popular movie, and the movie the preceding agent has seen. The resulting distribution of box office revenues follows a power law and well fits the actual data.

Joshua M. Epstein (2001) builds an ingenuously simple model of the social norm evolution based on a ring cellular automaton. The agents have a binary state which is subject to noise. Behavior heuristic is a simple conformance to a local majority, however the local neighborhood varies. Its radius is increased if the environment is diverse, and decreased if it is homogeneous, capturing the idea of thinking about a norm. Despite the simplicity, the model is capable of displaying global diversity and punctuated equilibria.

Tassier (2004) develops a model to explain fad and fashion cycles. He uses a cellular ring, but extends it with random transient weak ties. In this model agents buy goods to identify with social groups, however identifying with the majority is perceived as undesirable. Depending on the parameters, the model can generate conformity, diversity, and cycles. Turning to the effect of network topology on the dynamics, the author finds that group formation is facilitated by both connectedness and clustering of the network.

These models influenced the present research in a number of ways. They provided practical ideas, such as the evolution panel device from Epstein (2001). They stressed the importance of network topology. But most importantly, by deriving stylized facts about social cascades in seemingly different setups, these models motivated a search for an abstraction that includes all of them as special cases. A candidate for such an abstraction is the theory presented in the following chapter.

Chapter 2: Theory

The basic conclusion from the literature discussed above is that the general properties of social cascades, including fragility, unpredictability, and power law distributions are robust to specific modeling assumptions about individual behavior. A natural question to ask, then, is whether there is a conceptual framework that accounts for these regularities and is compatible with the former findings.

A good candidate is the frustrated systems theory found in physics. A system is called frustrated if it has several equilibria, or ground states, but cannot achieve any one of them because of excess energy. Physical systems of this kind, such as spin glasses, display phase behavior, with an organized and a disordered phase. Fragile cascades with power law size distributions arise near the phase transition point, when different clusters of the system adjust toward different equilibria.

As the idea that physical models can be adapted to economic phenomena is old and controversial, all new adaptation attempts have to be well justified. This paper seeks to adapt spin-glass models to social cascades because they are based on constrained optimization and reproduce fragility and other cascade phenomena. The adaptation posits two difficulties: interpreting the concept of temperature, which plays a central role in spin-glasses, and modeling the topology of agent interactions in a realistic way.

This chapter introduces the necessary basics of phase transitions and then proceeds to overcome the two difficulties by showing how the concept of temperature naturally emerges in the context of economic decision making, and by presenting recent findings in social network analysis that are later used to adapt spin-glass models to economics. It concludes by discussing related work.

2.1 Phase Transitions

Perhaps the simplest model that has a phase transition is a random graph (Paul Erdős and Alfréd Rényi, 1960). A graph is a set of objects called nodes, pairs of which can be either connected or disconnected with links called edges. For any pair, connect the two nodes with a probability p. Surprisingly, the result of this simple operation has two radically different outcomes, or phases. In the disconnected phase, $p < 1/N$ (N - number of nodes), nodes form little connected groups (clusters), while in the connected phase ($p > 1/N$) there is only one huge cluster. It is called infinite because it occupies a finite fraction of an arbitrarily large graph. The critical probability of connection $p_c = 1/N$ separates two radically different regimes. An equivalent problem is studied in percolation models (Albert and Barabási, 2002), but instead of random graphs the objects analyzed are d-dimensional lattices. Yet another example of a phase transition is found in the Ising model, the simplest of spin-glass models, which is introduced in the next section.

There is an important connection between phase transitions and power law distributions: the latter arise near the critical threshold. For instance, in random graphs as $p \to p_c$ in the subcritical phase, connected clusters form, and their distribution follows a power law (Albert and Barabási, 2002). A similar distribution is observed in the sizes of clusters of particles aligned in their spin in the Ising model near the critical temperature. This motivates an application of such models to social cascades; however, phase transitions are not the only source of power laws. The latter are commonly observed in nature, society and economics (including the Pareto distribution of incomes), and can be explained by a variety simple stochastic processes (Herbert A. Simon, 1955). What makes phase transitions an attractive explanation is that they explain not only the distribution, but also the fragility of cascades.

2.2 Interpreting Temperature

Physical and economic systems are similar as they are engaged in constrained optimization. Specifically, any closed thermodynamical system attempts to smoothly distribute a conserved amount of energy, maximizing entropy. In the Ising model (Roy J. Glauber, 1963), the analogy with optimizing individuals is even more pronounced. Structurally the model is a regular lattice of microscopic magnets. Each magnet decides between two states (spins), in a way that minimizes the energy of interaction with other magnets (bond energy) and with a magnetic field (potential energy). Thus the two factors influencing a magnet's decision are the decisions of neighbors and the intrinsic quality of an option, and they are modeled by bond and potential energy respectively.

However, magnets do not always choose the best option, because that may imply a change in energy. The Ising model has an equilibrium (ground state) where all magnets have the spin favored by the magnetic field, but it cannot achieve it because of the excess energy. For any given level of excess energy, there are many spin configurations, and dynamically the system is randomly alternating between them. At the microscopic level, magnets are randomly choosing a spin, in a way that the spin that minimizes the sum of bond and potential energies is the most probable, but the expected energy change is zero.

Choosing options other than the best one is a striking contradiction to rationality. However, this is an apparent contradiction. It can be resolved by admitting that the model does not account for all factors relevant to the decision in question and therefore includes noise, by considering dynamic learning under uncertainty about payoffs, or by applying game theory. The first approach simply states that agents sometimes behave unexpectedly because their motives are not fully described by the model; the second one makes greedy behavior irrational when payoffs are changing with time, and the agent can gain by sometimes deviating from her favorite option to check whether it is still the best; finally, the third approach models agents playing mixed strategies and adapting them by learning, or populations of agents evolving a strategy mix by natural selection.

Once the agents are allowed to randomize and make choices that are not myopically optimal, there is a problem of obtaining a decision probability distribution (mixed strategy) from option utilities. Physics offers an elegant solution, namely the Boltzmann formula:

$$P(x) = \frac{e^{-\frac{E(x)}{T}}}{\sum_i e^{-\frac{E(i)}{T}}}$$

This equation, used in thermodynamics as well as in the Ising model, gives the probability of a system being in state x as a function of its energy $E(x)$, the energies of all possible states $E(i)$ and the system temperature T. The latter is uniquely determined by the expected energy of the system. Indeed, given the state energies and the expected energy, the probability distribution described by the Boltzmann formula is in a sense the least biased among all admissible distributions, because it maximizes the information entropy (Claude E. Shannon, 1948). Instead of the expected energy, one can work with temperature directly. Note that it is a scaling factor: multiplying state energies and temperature by a common factor does not influence the probabilities.

The same functional form appears in economics as the multinomial logit model, and has often been used in theories of dynamic decision making and learning (Dirk Helbing, 1993; Dirk Helbing, 2004; Dale O. Stahl, 1997). In these theories the energies are replaced by utilities, temperature simply called a a scaling factor, and the exponent sign is now positive, for agents are maximizing utility rather than minimizing energy. But it is the same functional form, and this fact allows interpreting the interactions in the Ising model as decisions of optimizing agents. This is how the Boltzmann formula is used and interpreted in this paper.

2.3 Interaction Topology

The second difficulty in adapting spin-glass methods to economics is developing a realistic model of agent interactions. There are various ways to do this: the agents can interact with everyone via the market, or only with their predecessors as in rational herding; in evolutionary game theory they are randomly matched in pairs, and in the Ising model interactions happen with neighbors on a lattice. However, none of the above approaches captures the idea that the society is a complex web. Therefore, it is more reasonable to view interactions as happening between connected nodes of a network. This section introduces recent results from mathematical sociology that describe common topological features found in social networks, and describes a simple network construction algorithm that is used in Chapter 3.

2.3.1 Small-World Networks

The small-world network (Watts and Strogatz, 1998) phenomenon captured a lot of attention from the scientific community because networks of the specific form it prescribes are extremely prevalent in both nature and society[1]. Indeed, tho phenomenon appears to be almost as general and widely applicable as the normal distribution, and it is quite surprising that it was not discovered earlier. The three features that make small-world networks (SWN) stand out are:

1. high connectivity (short path lengths);
2. high clustering (probability of a node's neighbors being connected);
3. scale-free degree[2] distributions;

SWN are different from both totally random graphs and perfectly ordered systems like lattices. Unlike the former, small world networks maintain a lot of local structure: if two nodes share a neighbor, they are likely to be connected. Unlike the latter, SWN are small: traveling from one node to another takes only a few edges.

The third property means that social networks typically have a degree distribution that follows a power law, and therefore does not depend on the size of the network. Such distributions are found, to cite a few common examples, in scientific collaboration networks, the Internet, electric power grids, and neural networks of living organisms (Albert and Barabási, 2002). The power law contrasts the Poisson distribution of random graphs (Erdős and Rényi, 1960) where any pair of nodes has the same probability of being connected. Random graphs correspond to such common techniques as random pairing in evolutionary games, and it is therefore striking that they are not appropriate models of social networks. To get a further idea of how important this is, consider the results of Romualdo Pastor-Satorras and Alessandro Vespignani (2001), who show that epidemics in scale-free networks can start with any non-negative infection probability. In other words, there is no epidemic threshold. These theoretical predictions are well supported by empirical data on the spread of viruses over the Internet.

[1]An excellent book on small-world networks for the general reader is *Six Degrees* (Watts, 2003), while a good technical treatment of the subject from which the present discussion heavily borrows is found in Albert and Barabási (2002).

[2]Node degree is defined as the number of edges attached to a node.

2.3.2 Network Generation Algorithms

Because of their importance, scale-free networks have been the focus of a lot of research. Predictions of random graph theory have been corrected to allow for random networks with arbitrary degree distributions (Mark E. J. Newman, Duncan J. Watts, and Steven H. Strogatz, 2002). A number of network generating models were developed, in which an exponential degree distribution arises from preferential attachment of edges: the network is built sequentially, and new edges prefer (in the sense of an increasing probability) to attach to nodes with a larger degree (Barabási and Albert, 1999). However, the true generating mechanism remains unclear. Also, it is difficult to justify large scale-free networks as models of individual interaction for they do not have an upper limit for the degree, while such a limit exists for the number of people a particular individual can influence.

The second important departure from random graphs, which gave the name to SWN, is clustering. An individual's friends are likely to be acquainted, and articles cited in an academic paper are likely to cite each other. The robust measure for this phenomenon, the clustering coefficient, is defined on each node as the proportion of links between neighbors out of the maximum possible number of such links, and then defined on the graph as the average of node clustering coefficients. Empirical networks typically have a value of the clustering coefficient several orders of magnitude higher than random graphs of the same size.

One of the simplest models that generates a clustered SWN is the Watts-Strogatz model (Watts and Strogatz, 1998). It is used in this paper precisely because of its simplicity, even though it does not have preferential attachment and does not generate scale-free degree distributions. Start with a regular ring lattice and randomly rewire every edge with a probability β, distributing edges uniformly but excluding duplication and loops. The model's only parameter allows to interpolate between a regular lattice ($\beta = 0$) and an Erdos-Renyi graph ($\beta = 1$). Indeed, the model has been designed for this interpolation, with a view of comparing connectivity of random graphs and SWN. And the remarkable result is that a very low β is sufficient to radically improve the connectivity of the network, as measured by the mean path length. This means that in spite of clustering, distances within a SWN are not significantly different from those in a random graph.

Although the topology of social networks is far from being completely understood, there are two basic findings: real-world networks can be viewed as mixtures of random graphs and lattices, and their properties are often independent of their size, allowing to study large systems on small models.

2.4 Related Work

This concluding section briefly discusses the pioneering model of Watts (2002), the concept of self-organized criticality, and the Minority Game, to acknowledge the influence of these ideas, as well as to show what makes the present approach different from them.

2.4.1 Cascades on Social Networks

Watts (2002) was one of the first researchers to apply SWN to social cascades. In this paper he extends a linear threshold infection model (nodes get infected when the fraction of their infected neighbors exceeds a node-specific threshold) by using random graphs with arbitrary degree and threshold distributions. The resulting system has two phases, which differ by the existence of an infinite cluster, or global cascade. The author provides an elegant analytical solution that reduces the problem to statics, and derives a *cascade condition* formula that separates the two phases. The analysis is based on a structural phenomenon that is necessary for a global cascade, namely the *percolating vulnerable cluster*, a group of connected nodes that are activated by a single neighbor. If an infection hits such a group, the whole cluster is infected, and being sufficiently large it is able to infect the whole network. This paper approaches the same problem but with a different methodology: the approach is purely computational, the interaction model is more flexible, while the network model is simpler.

2.4.2 Self-Organized Criticality

The argument of self-organized criticality (SOC) is that if dynamical systems with phase transitions are attracted to the critical value, they will display scale-invariant features and criticality: unpredictable sudden changes as the relevant parameter passes above and below the threshold (Per Bak, Chao Tang, and Kurt Wiesenfeld, 1987). Frequently cited examples of such systems include sand piles, avalanches, forest fires, and since Bak (1999) attempts have been made to add the economy and social cascades to this list. The present work is similar to SOC in modeling cascades as phenomena that occur near the critical threshold, but does not require that threshold to be an attractor.

2.4.3 The Minority Game

The Minority Game (MG) provides a striking example of applicability of thermodynamic concepts (temperature, the Boltzmann formula, spin-glass models) to economics. MG was first introduced by Damien Challet and Yi-Cheng Zhang (1997) as a precise formulation of the El Farol Bar problem (Brian W. Arthur, 1994). Although it has a setup very different from cascades models, MG is briefly discussed here as an indirect justification for the methodology. The MG idea is extremely simple: agents make a binary choice (buy or sell, visit bar A or B) and are rewarded for being in the minority. In the original MG formulation, agents are limited in their analytical power, and remember only a short history (n bits of the system signal, where each bit encodes which option was the best in the corresponding period). They also have a randomly drawn strategy set, and play the strategy from this set that has the best performance on the recorded history.

The discovery that paved the way for thermodynamics into MG analysis is that the macroscopic dynamics do not change when memory is replaced with random noise, making the process Markovian (Andrea Cavagna, Juan P. Garrahan, Irene Giardina, and David Sherrington, 1999). Although this statement was later debated (Damien Challet and Matteo Marsili, 2000), thermal analysis introduced by Cavagna et al. (1999) was so fruitful that it is now the primary approach to MG. Further, Challet and Marsili (1999) related MG to spin systems, and found a phase transition and a counterpart to spontaneous magnetization. In another paper they characterize stationary MG states as ground states in a spin system (Marsili and Challet, 2001).

The results from MG analysis essentially indicate that thermodynamics provides models applicable to adaptive learning. This connection is explored by Giulio Bottazzi, Giovanna Devetag, and Giovanni Dosi (2001), who find that MG dynamics strongly depend on learning procedures, that the learning rule strongly resembles a replicator dynamic, connecting MG to evolutionary game theory, and finally that adding noise to individual decision rule can improve aggregate welfare. Although the precise connection between thermodynamics and behavior is not yet fully understood, it is a promising direction for research.

Now that the necessary building blocks are discussed, it is possible to assemble them into a computational model, which is the subject of the following chapter.

Chapter 3: Simulation

This chapter introduces a modeling framework for the study of social cascades. It first presents a general simulation methodology, then defines an interaction model compatible with most applications, and finally uses simulation examples to demonstrate how key facts about social cascades can be reproduced within this framework. It concludes by addressing the limitations of the software implementation employed to simulate the model.

3.1 Simulation Methodology

The simulation methodology employs a variation of cellular automata (CA) (Tommaso Toffoli and Norman Margolus, 1987; Stephen Wolfram, 2002). CA are well suited for studies of cascades because they explicitly model purely local interactions, produce complex macro behavior from simple micro rules, and complement numerical output with visual panels useful for quickly assessing qualitative features of the model. However, interactions in CA take place on a regular lattice which is not a good model of a social network.

This limitation is avoided by using an automata network, a one-dimensional CA where the traditional cellular ring is replaced by a random network constructed according to the undirected Watts-Strogatz model (Watts and Strogatz, 1998). Specifically, first a periodic lattice is built with n nodes and $\frac{nk}{2}$ edges, where every node is connected to $\frac{k}{2}$ neighbors at each side. Then every undirected edge of this lattice is replaced with probability β by a random edge, in such a way that loops and multiple edges are excluded, and the resulting graph is therefore strict. Varying β from 0 to 1 interpolates between a ring lattice and a random graph, producing structures comparable to social networks in the intermediate range.

Besides the underlying network, every CA requires an interaction model, which consists of a set of node states, an initial assignment of these states, and an update rule, a (possibly randomized) function that takes a node state and the states of the neighboring nodes and returns the new state. Once all parameters are specified, the CA simulation is run with Poisson updating: at time $t = 0$ every node is assigned an initial state, and then at each time step a random node is chosen and its state replaced by applying the update rule to it.

3.2 Visualization

The model dynamics are concisely represented by the evolution panel, an image where every horizontal line represents a snapshot of the system state at time step vt (typically $v = n$), with the top line marking the beginning of the evolution. Vertical lines in a panel represent the dynamics of individual nodes, with different states shown by different colors. Unlike the traditional CA ($\beta = 0$), network automata produce panels where neighboring locations are highly likely, rather than certain, to represent connected nodes, and distant locations can be connected as well.

3.3 Interaction Model

The interaction model seeks to be sufficiently general to include most applications as special cases. The set of states is a continuum $[0, 1]$, and agents can either imitate their neighbors (adopt a state present in the neighborhood), or innovate (adopt a random new state). Agents derive utility from the state itself, and from interacting with neighbors. Thus states can represent strategies, technologies or choices, and the utility function can capture payoffs, network effects, information externalities and the intrinsic quality of the available options. To transform utilities into a probabilistic update rule, a form of the Boltzmann equation is used:

$$P(x) = \frac{e^{\frac{U(x)}{T}}}{\sum_i e^{\frac{U(i)}{T}}}$$

The probability of choosing state x is proportional to $e^{U(x)/T}$, where $U(x)$ is the utility of state x, and T is the normalizing factor, or the (negative) temperature in the original Boltzmann formula. Cascades happen when imitation is rewarded, which can be modeled by making the utility higher for those states that are common in the neighborhood. Note that as $T \to 0$, the probability of choosing the state with the highest utility goes to 1, while as $T \to \infty$ the agent chooses all states with equal probability.

Obviously, T is very important. For high T, disorder swamps out all coordination, while for low T coordination may stabilize global cascades, preserving total conformity. Depending on the application, T can be interpreted as a measure of innovation frequency, availability of credit, or cultural conformity level. For some applications T need not be constant nor exogenous. A fluctuating T can produce punctuated equilibria as the system switches between the coordinated and the uncoordinated regimes. And an endogenous T can produce self-organized criticality (SOC) phenomena if T increases when low and decreases when high. For example, this argument may apply if T is the interest rate and innovation is the only source of demand for loans. While SOC models are appealing in some contexts, the general ubiquity of cascades is better explained by the fact that total order and randomness are not promising objects of study, and, possibly being even more frequent than cascades, they are simply ignored by scholars. For this reason, as well as for the sake of simplicity, the following simulations assume a constant exogenous T, studying how its level affects behavior.

3.4 Results

3.4.1 Phase Behavior

A simple model that illustrates phase behavior is shown in Figure 1. In this model at each time step the agent chooses between inventing a new state, or adopting a state from the neighborhood. Utility of a state is defined as the number of neighbors in that state. When T is small (left), only temporary and isolated clusters of innovation appear, while the majority of agents keeps the original state (black). The middle frame corresponds to a larger T, where innovation eventually leads to diversity with some local conformance. For very high T (right), the state distribution becomes random.

Such two-phase behavior is fairly general and robust to changes in the update rule, even though the critical temperature T_c often cannot be established. For example, qualitatively similar pictures can be obtained with a randomized majority rule (innovate with

Figure 1: **Phase transition near the critical temperature**

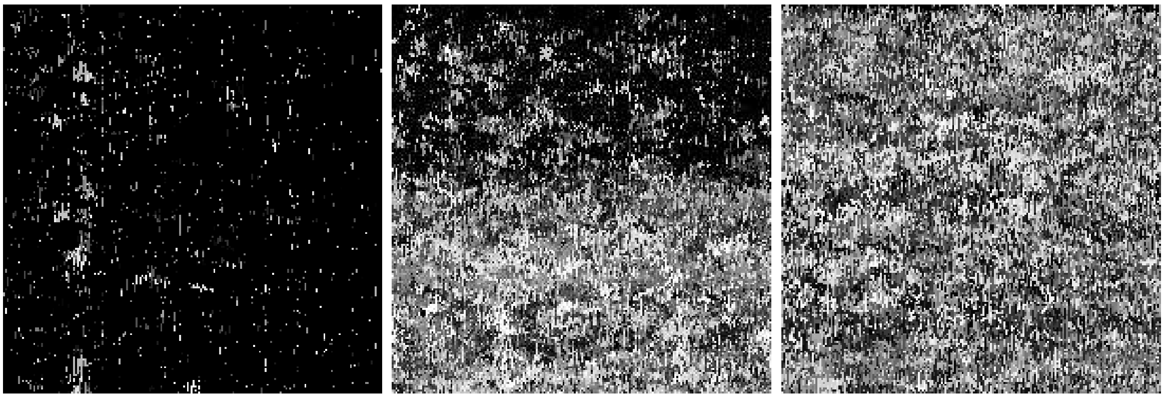

Simulations with $n = 200$, $k = 8$, $\beta = 0.3$, $t = 40000$ updates, $v = 200$ updates per line. $U(x)$ is the number of matching ($y = x$) states in the neighborhood. T is 2, 2.4 and 10, from left to right.

probability p, otherwise adopt the state that is most common among the neighbors, randomizing to break ties) or a simple infection rule (innovate with probability p, otherwise adopt the s state of a random neighbor). It is also observed in Ising models, which are a special case of the present model family with $\beta = 0$ and only two possible states. In Ising model simulations the fact that the lattice is finite blurs the phase transition (T_c is defined only in the thermodynamical limit). The same applies to the present model as well, and such blurring may be even stronger because of the departures from the Ising model.

The implication of phase behavior is that all cascade models should identify a critical parameter that corresponds to temperature, assess its stability and analyze factors (both endogenous and exogenous) that may influence it.

3.4.2 Effects of Network Topology

Network topology significantly influences cascade behavior. Local cascades are most readily observed on regular lattices ($\beta \to 0$) because their structure facilitates diversity in weakly connected clusters. Even small increases of β greatly improve the network connectivity, and local patterns cease to be sustainable. Then the system either synchronizes to a single state (global cascade), or displays random states, depending on the temperature.

Figure 2 displays three evolutions with the same T but three different levels of β, suggesting that T_c may increase with β. These results imply that a long average path length is necessary to preserve diversity of local cascades in systems that tend to coordinate. Such diversity is easily obtained with CA models, but does not carry over to dense social networks. Care should be taken therefore in applying CA to social processes, and simple one- or two-dimensional CA can only be justified for models of geographical interaction, such as urban development.

Figure 2: **Cascades on different network topologies**

Simulations with $n = 200$, $k = 8$, $T = 2.2$, $t = 40000$ updates, $v = 200$ updates per line. $U(x)$ is the number of matching ($y = x$) states in the neighborhood. β is 0, 0.1 and 0.3, from left to right.

Figure 3: **Fragility and unpredictability of cascades**

Simulations with $n = 200$, $k = 8$, $T = 1.75$, $t = 40000$ updates, $v = 200$ updates per line. $U(s) = 0$, where s is the initial state, otherwise $U(x)$ is the number of matching ($y = x$) states in the neighborhood. β is 0, 0.1 and 0.3, from left to right.

3.4.3 Fragility

Cascades are known to be unpredictable and fragile. To some extent this was visible in the simulations above, but unpredictability really stands out when the system starts out with an empty state rather than a global cascade. The former model is now slightly modified:

let the utility of the initial state be 0. Many agents start to innovate at $t = 0$ and then the different but identically useful innovations spread out. As Figure 3 illustrates, it is impossible to predict how successful any single innovation will be or when it will become extinct. It is also clear that better connected networks (higher β) converge to a single state much faster.

3.4.4 Cascade Size Distribution

In physical models with a phase transition, such as the percolation model and the Ising model, power law distributions appear near the critical point, and the exponents in these power laws are used to determine a universality class of the studied phenomenon. Similar exponents are indicative of similar dynamics. Since the present model is similar to the Ising model, the cascade size distribution may also be scale-free near the critical temperature.

Figure 4: **Cascade size distribution**

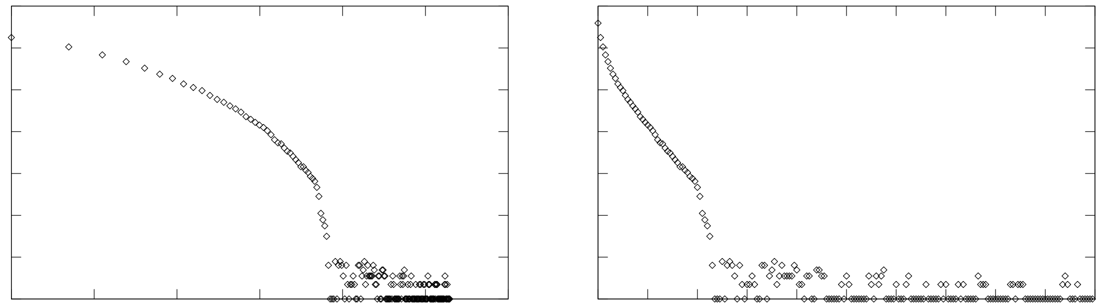

Simulation with $n = 200$, $k = 8$, $\beta = 0.3$, $T = 2.1$, $t = 40000$ updates, $v = 200$ updates per line. $U(s) = 0$, where s is the initial state, otherwise $U(x)$ is the number of matching ($y = x$) states in the neighborhood. The graph shows the time-averaged distribution of cascade sizes in log-log (left) and log-lin (right, X axis linear) scales.

To test this hypothesis, a simple simulation is carried out. T is tuned to the critical range, so that the different runs of the model alternate between a single global cascade and a diversity of local cascades. Then, a run that maintains diversity is selected, and the cascade sizes are measured at each time step.

The time-averaged cascade size distribution is shown in Figure 4, using log-log and log-lin plots. The former plot makes power laws appear linear, while the latter does the same to the exponential function. Note that the distribution can be reliably estimated only for small cascade sizes, as larger cascades are rare by design. Since the left part of the graph is more linear in the log-lin than the log-log plot, the exponential distribution is more consistent with the data. The failure to get a scale-free distribution in this experiment does not imply that scaling can never be obtained in the presented model, and may result from the use of time-averaging instead of independent experiments, or from the limited network size.

In general, cascade studies can always benefit from estimating the size distribution and, if possible, identifying a universality class.

3.4.5 Growth Patterns

So far the states were assumed to be identical. However, the model is flexible enough to have some states intrinsically more likely than others. This may reflect, for instance, technologies of different potential. An efficiency question then arises - can the best technology make its way and produce a global cascade?

The update rule is modified in such a way that the utility of a state depends on the state value (which is, as before, uniformly distributed on the $[0, 1]$ interval). Specifically, a state x gets x utility points for every neighbor that has the same state. This setup was selected for its simplicity, although it assumes zero utility for states that are unique in the neighborhood. These states are still selected with positive probability, which explains why simulations with a positive utility of unique states produce similar dynamics. Colors now represent states on a grayscale, from black (0) to white (1). Finally, two statistics are gathered during the system evolution: the average of the state values (average quality), and the share of the most prevalent state.

As Figure 5 shows, near the critical temperature the system converges to a global cascade. There are two things of note. First, although the system generally settles in a good (close to 1) state, it is possible for better states to be rejected because they failed to spread widely enough. Second, the growth pattern is S-shaped, or, more precisely, there are three periods - diversity, growth, and conformity. The small peak in the second graph in Figure 5 before the near-vertical growth stage corresponds to a state that used to be a favorite but was later rejected, illustrating the unpredictability of the process. Also, the S-shaped growth appears for reasons very different from those found in classical technology diffusion models (Bass, 1969), yet it appears.

Figure 5: **S-shaped cascade growth**

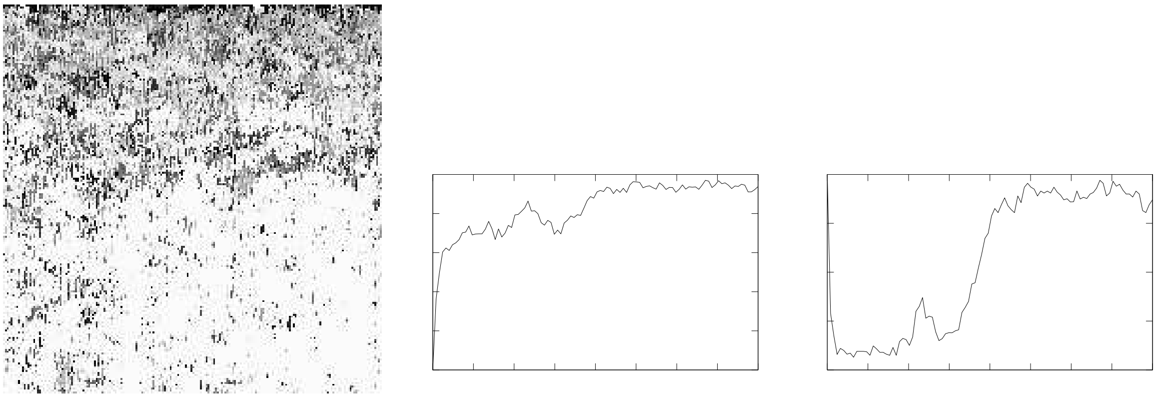

Simulation with $n = 200$, $k = 8$, $\beta = 0.3$, $T = 2.2$, $t = 40000$ updates, $v = 200$ updates per line. $U(x) = xn$, where n is the number of neighbors in state x. Colors encode states using gray-scale from black ($x = 0$) to white ($x = 1$). Graphs show time evolution of the state value average (middle) and the share of the most prevalent state (right).

The remaining figures show how the system behaves in different temperature regimes. At very low temperatures (Figure 6), local stable cascades are formed, and average state quality is low. When the temperature is slightly higher (Figure 7), efficiency is improved and there is convergence to a single cascade, however diversity may be maintained for a long time. Finally, temperatures above the critical range produce a diversity of local cascades that is not efficient (Figure 8).

To summarize, the states that are intrinsically more likely are also much more likely to produce a global cascade. However, the cascade formation process is unpredictable and, in terms of technology diffusion, there may be a trade-off between popularity and efficiency.

3.5 Limitations

While the presented results are sufficient to show that the use of automata networks is a promising direction in social cascade modeling, they have a number of limitations. First, the network size is small ($n = 200$), and therefore there is no attempt to model a scale-free degree distribution, which is unlikely to make a difference at this scale. Second, the

Figure 6: **Freezing at low temperatures**

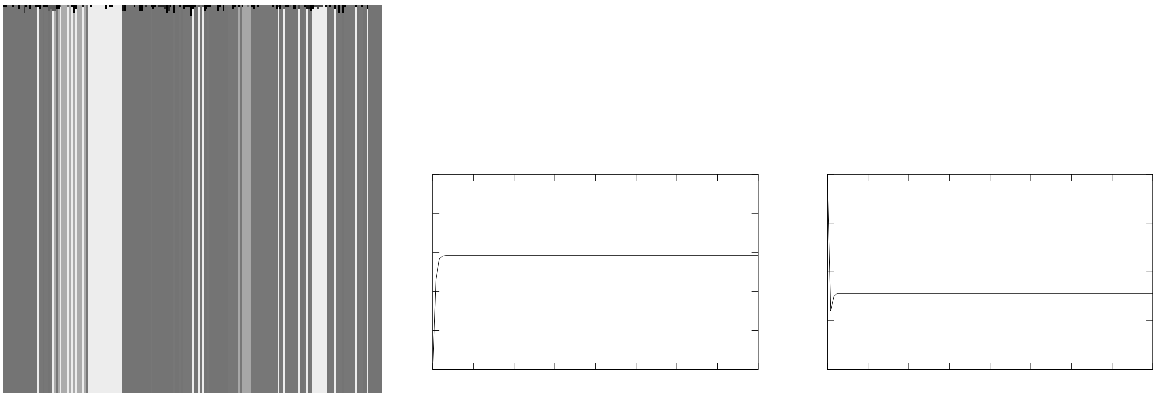

Simulation with the same setup as in Figure 5, but with $T = 0.01$.

Figure 7: **Co-existence of cascades**

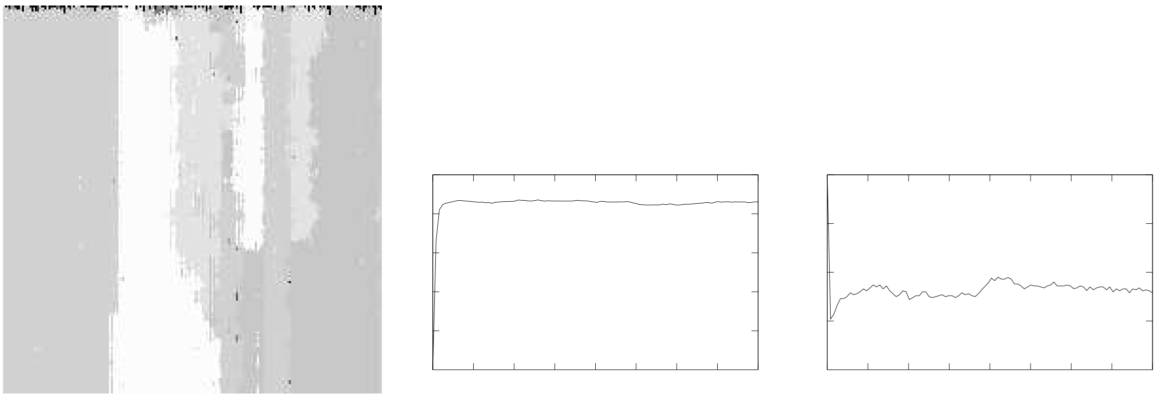

Simulation with the same setup as in Figure 5, but with $T = 0.5$.

Figure 8: **Local cascades at high temperatures**

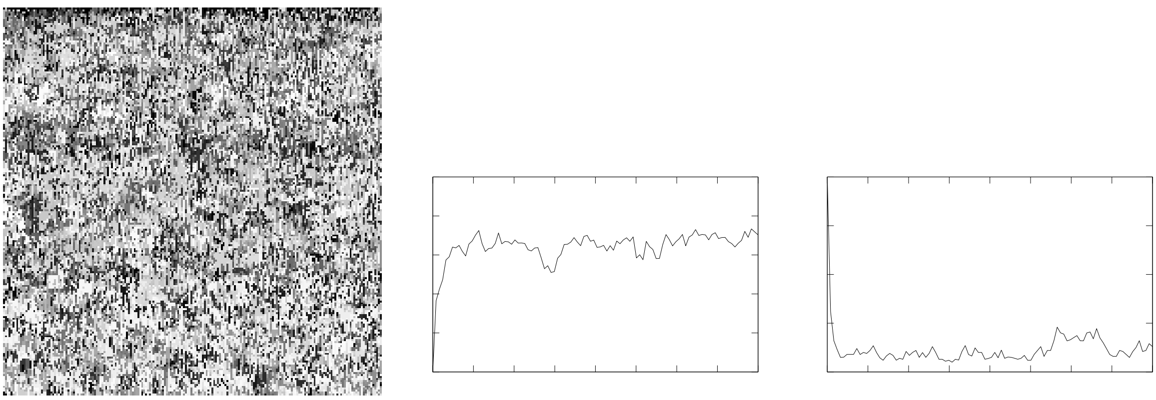

Simulation with the same setup as in Figure 5, but with $T = 3$.

cascade size analysis pools distributions across time in a single experiment, instead of collecting data at a common time point across independent experiments. These limitations result from the exploratory nature of the work and the low speed of simulation.

Many experiments had to be done with a variety of setups to develop an intuition of the model and to find a formulation that is not arbitrary. With limited time this meant that experiments had to be done quickly, and this, together with the speed of the software implementation, has put very tight size constraints. The typical experiment therefore had 200 nodes, 1600 edges and 40000 updates, which took $5-7$ seconds on a computer with a 1.8 GHz CPU. There are three ways to overcome these constraints: optimize the software, allow more time for an experiment, and use a faster computer.

Only the software optimization possibility deserves a brief discussion. The present implementation was done in Haskell, a purely functional language. Haskell excellently maps abstract mathematical concepts to code, has a good graph library, does rigorous type checking which eliminates many errors. It is also generally fast, comparing in performance to C++. However, the ease of programming comes at a price of difficult optimization. Also, because of difficulties in porting graphics packages, a highly suboptimal text-based PPM image format was selected for visualization. The present implementation therefore can be made considerably faster by either optimizing the existing code or translating the program to C.

Conclusion

Frustrated systems theory provides a plausible solution for the problem of explaining the dynamic regularities of social cascades. This solution does not depreciate the earlier findings of economic literature, as they explain the incentives for the individual to imitate her neighbors, and the theory relies on these explanations. Rather, these findings are complemented by a well-known and tested macroscopic approach which naturally reproduces fragility, unpredictability, and scale-free size distributions of social cascades. The approach is applicable to economics because it is based on constrained optimization, local interactions and randomized behavior strategies, all compatible with economic methodology and used in many models. Indeed, this paper is not the first to apply the frustrated systems theory to economics.

The specific finding of this paper is more narrow: it shows that agent-based models of social cascades based on local imitation with random deviation are frustrated systems sensitive to two assumptions, namely the amount of the built-in disorder and the interaction topology. The implied level of decision randomization, as measured by temperature, separates the organized and the disordered phases. If it is endogenous, the phase transition point may be an attractor for the system, leading to self-organized criticality manifested in punctuated equilibria. The second important assumption is the shape of interactions. While cellular automata capture clustering and random interactions reflect the interconnectedness of the society, none of these two common models adequately models the real structure of social networks which is both clustered and interconnected. Experiments with hybrid structures show that clustering facilitates diversity and interconnectedness leads to conformity, and that in general the interaction topology assumption is crucial to the resulting dynamics.

This finding contributes to the field of agent-based modeling of social cascades by providing an analytical framework that helps to design appropriate models and explain their results. In particular, it helps to better understand the computational models discussed in Section 1.5. Thus power law distributions of movie box office revenues found by De Vany and Lee (2001) are far from being unique to his model, and should be retained in a variety of formulations. Cascade clustering and S-shaped growth discovered by Bala and Goyal (1998) are also not surprising but must be sensitive to their assumption that interactions happen on a grid. The growth and collapse of the largest firm in Axtell (1999) can be seen as a phase transition produced by an increase in temperature. Finally, punctuated equilibria obtained by Tassier (2004) and Epstein (2001) are examples of self-organized criticality as both models are attracted to the critical temperature. Specifically, in Tassier (2004) a conforming phase encourages agents to deviate (it is not fashionable to associate with the majority), while in a totally disordered phase agents seek to conform more (it is fashionable to associate). Similarly, in the model introduced by Epstein (2001) conformity induces agents to reduce their search radius, making their decisions more susceptible to random changes of neighbors' decisions, which can be seen as an increase in a properly defined system temperature.

Besides helping to understand agent-based models, the frustrated systems theory is applicable to the real world. It illustrates an intuitive idea that social changes are difficult but possible to start, there being a critical level of effort toward change, beyond which stable conformity gives way to turbulent experimentation that in its turn ceases when a new consensus is reached. Most importantly, the theory accounts for the fact that social changes have a geography as they happen in clusters, and there is often a cluster that catalyzes change.

Simulations of generalized spin-glass models on social networks are a promising direction for further research. The simulations carried out in this paper are limited by being exploratory and using small network sizes, but even at this scale the results are sufficiently interesting. With more resources, larger networks can be explored. It is possible either to artificially construct networks or to use existing large network datasets. Larger sizes and faster software will allow to study models with scale-free degree distributions, and possibly other network features that the quickly developing field of mathematical sociology may discover.

Finally, the visualization technique used in this paper can be easily generalized. There exist algorithms that can arrange any network in an n-dimensional space in a way that facilitates visual perception by minimizing the distances between connected nodes. Although this arrangement was obtained by construction rather then by applying a special algorithm, this is exactly the way networks were represented in one dimension in this paper. Such algorithms make visualization tools developed for cellular automata applicable to the more general network models. Thus, with the second dimension showing time, evolution panels of one-dimensional arrangements can be drawn for presentation in printed material, while two-dimensional arrangements can be easily studied in real time on a computer screen.

References

Albert, Réka, and Barabási, Albert-László. 2002. "Statistical mechanics of complex networks." *Reviews of Modern Physics*, *74*, 47.

Arthur, Brian W. 1994. "Inductive reasoning and bounded rationality (the El Farol problem)." *American Economic Review*, *84*, 406–411.

Axtell, Robert. 1999. *The emergence of firms in a population of agents* (Working Papers No. 99-03-019). Santa Fe Institute. (available at http://ideas.repec.org/p/wop/safiwp/99-03-019.html)

Bak, Per. 1999. *How nature works: The science of self-organized criticality.* Springer-Verlag Telos.

Bak, Per, Tang, Chao, and Wiesenfeld, Kurt. 1987. "Self-organized criticality - An explanation of 1/f noise." *Physical Review Letters*, *59*, 381-384.

Bala, Venkatesh, and Goyal, Sanjeev. 1998, July. "Learning from neighbours." *Review of Economic Studies*, *65*(3), 595-621. (available at http://ideas.repec.org/a/bla/restud/v65y1998i3p595-621.html)

Bandiera, Oriana, and Rasul, Imran. 2002. *Social networks and technology adoption in Northern Mozambique* (STICERD - Development Economics Papers No. 35). Suntory and Toyota International Centres for Economics and Related Disciplines, LSE. (available at http://ideas.repec.org/p/cep/stidep/35.html)

Banerjee, Abhijit. 1992, August. "A simple model of herd behavior." *The Quarterly Journal of Economics*, *107*(3), 797-817. (available at http://ideas.repec.org/a/tpr/qjecon/v107y1992i3p797-817.html)

Barabási, Albert-László, and Albert, Réka. 1999, October. "Emergence of scaling in random networks." *Science*, *286*, 509-512.

Bass, Frank M. 1969, January. "A new product growth for model consumer durables." *Management Science*, *15*(5), 215-227.

Bikhchandani, Sushil, Hirshleifer, David, and Welch, Ivo. 1992, October. "A theory of fads, fashion, custom, and cultural change in informational cascades." *Journal of Political Economy*, *100*(5), 992-1026. (available at http://ideas.repec.org/a/ucp/jpolec/v100y1992i5p992-1026.html)

Bikhchandani, Sushil, Hirshleifer, David, and Welch, Ivo. 1998, Summer. "Learning from the behavior of others: Conformity, fads, and informational cascades." *Journal of Economic Perspectives*, *12*(3), 151-70. (available at http://ideas.repec.org/a/aea/jecper/v12y1998i3p151-70.html)

Bottazzi, Giulio, Devetag, Giovanna, and Dosi, Giovanni. 2001. *Adaptive learning and emergent coordination in minority games* (Computing in Economics and Finance 2001 No. 20). Society for Computational Economics. (available at http://ideas.repec.org/p/sce/scecf1/20.html)

Cavagna, Andrea, Garrahan, Juan P., Giardina, Irene, and Sherrington, David. 1999. "A thermal model for adaptive competition in a market." *Physical Review Letters*, *83*, 4429.

Celen, Bogachan, and Kariv, Shachar. 2004, April. "Observational learning under imperfect information." *Games and Economic Behavior*, *47*(1), 72-86. (available at http://ideas.repec.org/a/eee/gamebe/v47y2004i1p72-86.html)

Challet, Damien, and Marsili, Matteo. 1999. *Phase transition and symmetry breaking in the minority game.* (available at http://www.citebase.org/abstract?id=oai:arXiv.org:cond-mat/9904071)

Challet, Damien, and Marsili, Matteo. 2000. *Relevance of memory in minority games.* (available at http://www.citebase.org/abstract?id=oai:arXiv.org:cond-mat/0004196)

Challet, Damien, and Zhang, Yi-Cheng. 1997. *Emergence of coop-*

eration and organization in an evolutionary game. (available at http://www.citebase.org/abstract?id=oai:arXiv.org:adap-org/9708006)

Choi, Jay Pil. 1997, Autumn. "Herd behavior, the 'penguin effect,' and the suppression of informational diffusion: An analysis of informational externalities and payoff interdepend." *RAND Journal of Economics*, *28*(3), 407-425. (available at http://ideas.repec.org/a/rje/randje/v28y1997iautumnp407-425.html)

De Vany, Arthur, and Lee, Cassey. 2001, March. "Quality signals in information cascades and the dynamics of the distribution of motion picture box office revenues." *Journal of Economic Dynamics and Control*, *25*(3-4), 593-614. (available at http://ideas.repec.org/a/eee/dyncon/v25y2001i3-4p593-614.html)

Drehmann, Mathias, Oechssler, Jorg, and Roider, Andreas. 2004. *Herding with and without payoff externalities - an internet experiment* (Discussion Papers No. 51). SFB/TR 15 Governance and the Efficiency of Economic Systems, Free University of Berlin, Humboldt University of Berlin, University of Bonn, University. (available at http://ideas.repec.org/p/trf/wpaper/51.html)

Duan, Wenjing, Gu, Bin, and Whinston, Andrew B. 2006. "Herd Behavior and Software Adoption on the Internet: An Empirical Investigation" [Working Paper Series]. *SSRN eLibrary*.

Ellison, Glenn, and Fudenberg, Drew. 1995, February. "Word-of-mouth communication and social learning." *The Quarterly Journal of Economics*, *110*(1), 93-125. (available at http://ideas.repec.org/a/tpr/qjecon/v110y1995i1p93-125.html)

Epstein, Joshua M. 2001, August. "Learning to be thoughtless: Social norms and individual computation." *Computational Economics*, *18*(1), 9-24. (available at http://ideas.repec.org/a/kap/compec/v18y2001i1p9-24.html)

Erdős, Paul, and Rényi, Alfréd. 1960. "On the evolution of random graphs." *Published at Math. Inst. Hungarian Academy of Science*, *5*, 17-61.

Farrell, Joseph, and Saloner, Garth. 1985. "Standardization, compatibility, and innovation." *The RAND Journal of Economics*, *16*(1), 70-83.

Gale, Douglas. 1996, April. "What have we learned from social learning?" *European Economic Review*, *40*(3-5), 617-628. (available at http://ideas.repec.org/a/eee/eecrev/v40y1996i3-5p617-628.html)

Gale, Douglas, and Kariv, Shachar. 2003, November. "Bayesian learning in social networks." *Games and Economic Behavior*, *45*(2), 329-346. (available at http://ideas.repec.org/a/eee/gamebe/v45y2003i2p329-346.html)

Glauber, Roy J. 1963. "Time-dependent statistics of the ising model." *Journal of Mathematical Physics*, *4*(2), 294-307.

Goyal, Sanjeev. 2003. *Learning in networks: a survey* (Economics Discussion Papers No. 563). University of Essex, Department of Economics. (available at http://ideas.repec.org/p/esx/essedp/563.html)

Goyal, Sanjeev, and Vega-Redondo, Fernando. 2004, December. *Structural holes in social networks.* (available at http://eurequa.univ-paris1.fr/Colloques-s%E9minaires/CTN05/pdf/Goyal-Vega-Redondo.pdf)

Helbing, Dirk. 1993. "Boltzmann-like and Boltzmann-Fokker-Planck equations as a foundation of behavioral models." *Physica A*, *196*, 546.

Helbing, Dirk. 2004. "Dynamic decision behavior and optimal guidance through information services: Models and experiments." In Michael. Schreckenberg and Reinhard. Selten (Eds.), *Human behavior and traffic networks* (p. 47-95). Berlin: Springer. (available at http://ssrn.com/paper=960219)

Kariv, Shachar. 2005. *Overconfidence and informational cascades* (Tech. Rep.). UCLA Department of Economics. (available at http://ideas.repec.org/p/cla/levrem/122247000000000406.html)

Leskovec, Jure, Adamic, Lada A., and Huberman, Bernardo A. 2006. "The dynamics of viral marketing." In *Ec '06: Proceedings of the 7th acm conference on electronic commerce* (pp. 228–237). New York, NY, USA: ACM Press.

Li, Xinxin, and Hitt, Lorin. 2004. *Self selection and information role of online product reviews.* Washington DC.: Workshop on Information Systems and Economics (WISE).

Loch, Christoph H., and Huberman, Bernardo A. 1999. "A punctuated-equilibrium model of technology diffusion." *Manage. Sci.*, *45*(2), 160–177.

Marsili, Matteo, and Challet, Damien. 2001. *Continuum time limit and stationary states of the minority game.* (available at http://www.citebase.org/abstract?id=oai:arXiv.org:cond-mat/0102257)

Melissas, Nicolas. 2005, October. "Herd behaviour as an incentive scheme." *Economic Theory*, *26*(3), 517-536. (available at http://ideas.repec.org/a/spr/joecth/v26y2005i3p517-536.html)

Newman, Mark E. J., Watts, Duncan J., and Strogatz, Steven H. 2002. "Random graph models of social networks." *PNAS*, *99*(90001), 2566-2572.

Oberhammer, Clemens, and Stiehler, Andreas. 2001, May. *Does cascade behavior in information cascades reflect bayesian updating?* (Sonderforschungsbereich 373 No. 2001-32). Humboldt Universitaet Berlin. (available at http://ideas.repec.org/p/wop/humbsf/2001-32.html)

Pastine, Ivan, and Pastine, Tuvana. 2005. *Signal accuracy and informational cascades* (CEPR Discussion Papers No. 5219). C.E.P.R. Discussion Papers. (available at http://ideas.repec.org/p/cpr/ceprdp/5219.html)

Pastor-Satorras, Romualdo, and Vespignani, Alessandro. 2001. "Epidemic spreading in scale-free networks." *Physical Review Letters*, *86*, 3200.

Scharfstein, David, and Stein, Jeremy. 1990, June. "Herd behavior and investment." *American Economic Review*, *80*(3), 465-79. (available at http://ideas.repec.org/a/aea/aecrev/v80y1990i3p465-79.html)

Shannon, Claude E. 1948, July, October. "A mathematical theory of communication." *The Bell System Technical Journal*, *27*, 379–423, 623–656.

Simon, Herbert A. 1955, December. "On a class of skew distribution functions." *Biometrika*, *42*, 425-440.

Smith, Lones, and Sorensen, Peter. 1998, October. *Rational social learning with random sampling.* (available at http://www-personal.umich.edu/ lones/ftp/rational.pdf)

Smith, Lones, and Sorensen, Peter. 2000, March. "Pathological outcomes of observational learning." *Econometrica*, *68*(2), 371-398.

Stahl, Dale O. 1997. *Rule learning in symmetric normal-form games: Theory and evidence* (Economics, University of Texas at Austin No. 9710). Center for Applied Research in Economics. (available at http://ideas.repec.org/p/tex/carewp/9710.html)

Tassier, Troy. 2004. "A model of fads, fashions, and group formation." *Complexity*, *9*(5), 51–61.

Tassier, Troy, and Menczer, Filippo. 2001, October. "Emerging small-world referral networks in evolutionary labor markets." *IEEE-EC*, *5*, 482-492.

Toffoli, Tommaso, and Margolus, Norman. 1987. *Cellular automata machines: a new environment for modeling.* Cambridge, MA, USA: MIT Press.

Tucker, Catherine. 2005. *Networks, network effects and technology adoption.* PhD thesis, Stanford University.

Watts, Duncan J. 2002. "A simple model of global cascades on random networks." *Proceedings of the National Academy of Science*, *99*, 5766-5771.

Watts, Duncan J. 2003. *Six degrees: The science of a connected age.* W. W. Norton & Company.

Watts, Duncan J., and Strogatz, Steven H. 1998, April. "Collective dynamics of 'small-world' networks." *Nature*, *393*, 440-442.

Wolfram, Stephen. 2002. *A new kind of science.* Champaign, USA: Wolfram Media, Inc.

Young, H. Peyton. 2000, May. *The diffusion of innovations in social networks* (Economics

Working Paper Archive No. 437). The Johns Hopkins University,Department of Economics. (available at http://ideas.repec.org/p/jhu/papers/437.html)
Young, H. Peyton. 2006, March. *The spread of innovations by social learning.* (available at http://www.econ.jhu.edu/People/Young/Spread28march.pdf)