CHAPTER 5

On Information-Contagious Behavior★

NICOLAAS J. VRIEND

Queen Mary, Department of Economics, University of London, Mile End Road, London, E1 4NS, UK

Abstract

Alongside increasing returns, network externalities, and information cascades, information contagion has been presented in the literature as an explanation for particular patterns of macrobehavior that may seem at odds with the underlying micromotives. We present an agent-based computational economics model that could provide a microfoundation for information contagion. Our model exhibits self-organization through information contagious behavior, and the emergence of spontaneous orders, in which typically most agents choose the same superior item. But it turns out that this is not a simple monotonic process from disorder to order. Instead, the system continually moves back and forth between order and disorder as the self-organization is a continuing story in which the emerging order unravels time and again. In other words, information contagion is an inherently complex phenomenon.

Keywords: Agent-based computational economics, decentralized interaction, reinforcement learning, information aggregation, self-organization
JEL classifications: D11, D83, O33

★ The model presented here corresponds to the one used in Vriend (2002). Comments on either version by Pierre Barbaroux, Bruce Caldwell, Augustino Manduchi, Martin Posch, Jan Tuinstra, seminar and conference participants in Freiburg, Genova, London (QM and RH), Amsterdam, Aix-en-Provence, Vienna, Barcelona, Pisa (S. Anna), Essex, and Salerno are gratefully acknowledged. The usual disclaimer applies.

5.1. INTRODUCTION

Alongside increasing returns (Arthur, 1989), network externalities (Katz and Shapiro, 1985, 1986), information cascades (Bikhchandani et al., 1992), and herding behavior (Banerjee, 1992), information contagion (Arthur and Lane, 1991) has been presented in the literature as an explanation for particular patterns of macrobehavior (for example, path-dependence and lock-in effects) that may seem at odds with the underlying micromotives. But whereas these other explanations have been shown to have a proper microfoundation (either related to changing productivity or changing preferences, or to Bayesian updating in the face of incomplete information), information contagion has remained a phenomenon that occurs only when certain ad hoc rules of thumb for individual behavior are assumed.

Therefore we study the phenomenon of information contagion in a setup that is closely related to the one presented in Arthur and Lane (1991), in which individuals have to make repeatedly a choice between two previously unknown items while they can rely only on some information from previous adopters. We will present an agent-based computational economics (ACE) model that provides a microfoundation for information contagion, based on a simple model of adaptive behavior with agents trying to do the best they can, and without needing to assume that they use certain ad hoc rules of thumb. Our model exhibits self-organization through information contagious behavior, and the emergence of spontaneous orders, in which typically most agents choose the same, superior item. That is, through a self-organizing process the economy overcomes the problem of the division of knowledge (see, e.g., Hayek, 1948). But it turns out that this is not a simple monotonic process from disorder to order until the solution has been reached, with a happyending. Instead, the system continually moves back and forth between order and disorder. That is, the self-organization is a continuing, ongoing story, in which the emerging order unravels repeatedly. In other words, information contagion, unlike increasing returns to scale, network externalities, information cascades, and herding behavior, is an inherently complex phenomenon.

This chapter is organized as follows. Section 5.2 presents our ACE model of the emergence of information contagion, and Section 5.3 analyzes the properties of the model. In Section 5.4 we put our model into a somewhat wider perspective by discussing some related literature on information contagion and social learning, while Section 5.5 concludes.

5.2. THE AGENT-BASED COMPUTATIONAL MODEL

The basic choice problem we consider is that of a population of individual agents each of whom, sequentially face a decision problem between two items with uncertain qualities. We can think of these two items as new products, movies, technologies, services, financial gurus, or whatever binary choice agents might need to make frequently in every day life.

The only information the agents have is the choices plus the corresponding values experienced by a sample of other agents who had faced the same decision problem before them. This implies that there is an "information externality." That is, the choice of an agent does not only lead to utility for himself, but it will also be added to the pool of information from which other agents sample. The question, then, is what the consequences of this information externality are.

This basic choice problem has been considered in the literature. See in particular Arthur and Lane (1991), Dosi et al. (1994), Narduzzo and Warglien (1996), and Lane and Vescovini (1996). Basically what this literature shows, both theoretically and empirically, is that agents may behave in a way that the decision of a given agent positively affects the expected decisions of subsequent agents, leading to path-dependent lock-in effects. That is, there may be a diffusion process such that a certain choice once it starts being made by a certain number of people spreads quickly in a population (without the values actually experienced necessarily implying this). Since the only link between the decisions of the agents is the information externality, this contagious phenomenon is called "information contagion." What is missing in this literature is an *explanation* as to why we should expect people to behave in such a way that the information externality does indeed imply information contagion. In this respect, the literature on information contagion differs from the literature on increasing returns (Arthur, 1989), network externalities (Katz and Shapiro, 1985, 1986), information cascades (Bikhchandani et al., 1992), and herding behavior (Banerjee, 1992). All these models have been presented in the literature as an explanation for particular patterns of macrobehavior (for example, path-dependence and lock-in effects) that may seem at odds with the underlying micromotives. But whereas these other explanations have been shown to have a proper microfoundation (either related to changing productivity or changing preferences, or to Bayesian updating in the face of uncertainty), information contagion has remained a phenomenon that occurs only when certain ad hoc rules of thumb for individual behavior are assumed. Our ACE model will provide an explanation for information contagion.[1]

5.2.1. The basic choice situation

The model has a population of 100 decision makers. In a given period they face a choice between two items that were previously unknown to them. Each new item i is characterized by the expected value of the utility it will generate, EV_i. These expected values are unknown to the individual agents. Given an expected value, EV_i, the value that a specific agent will actually experience from an item will be a random draw from a uniform distribution with support from $EV_i - 0.25$ to $EV_i + 0.25$. Hence, if a given item i is characterized by an EV_i of, say, 0.40, the actual utility levels experienced by the individual agents choosing this item will range from 0.15 to 0.65, with every utility level in this range equally likely to occur. The stochastic character of the payoffs generated reflects idiosyncratic productivity or taste factors, but we can also think of the random component of the payoffs as measurement errors of a given item's actual value.

Notice that we do not have any increasing real returns to scale of any form, no change in taste, endogenously determined utility depending on the number of adopters, nor are there complementarities or network externalities. Each individual agent's utility of a certain item i is simply an independent draw from the same uniform distribution characterized by the item's expected value EV_i. Figure 5.1 gives an example of two items with expected values $EV_1 = 0.40$, and $EV_2 = 0.55$.

The agents, then, face their choice problem sequentially, with the order of the agents being random. Although each individual agent himself has no experience with these two specific new items, he can draw six random samples from the people who have already made a decision before him. For each of the elements in his sample, he can observe the choice made, and the value actually experienced by the agent.[2] Given this sample information, an agent makes a choice himself, and then the next agent in the queue makes his decision, until the end of the queue is reached. Before the first agent in the sequence makes

[1] The ACE model will be described in Sections 5.2.1 to 5.2.3. The pseudo-code of the model can be found in the appendix.

[2] Notice that the sample size is fixed exogenously, and we do not analyze the issue of optimal sampling strategies. This follows the existing literature on information contagion. More in particular, the functional specification and the number of observations sampled, six, are based on the experimental study by Narduzzo and Warglien (1996).

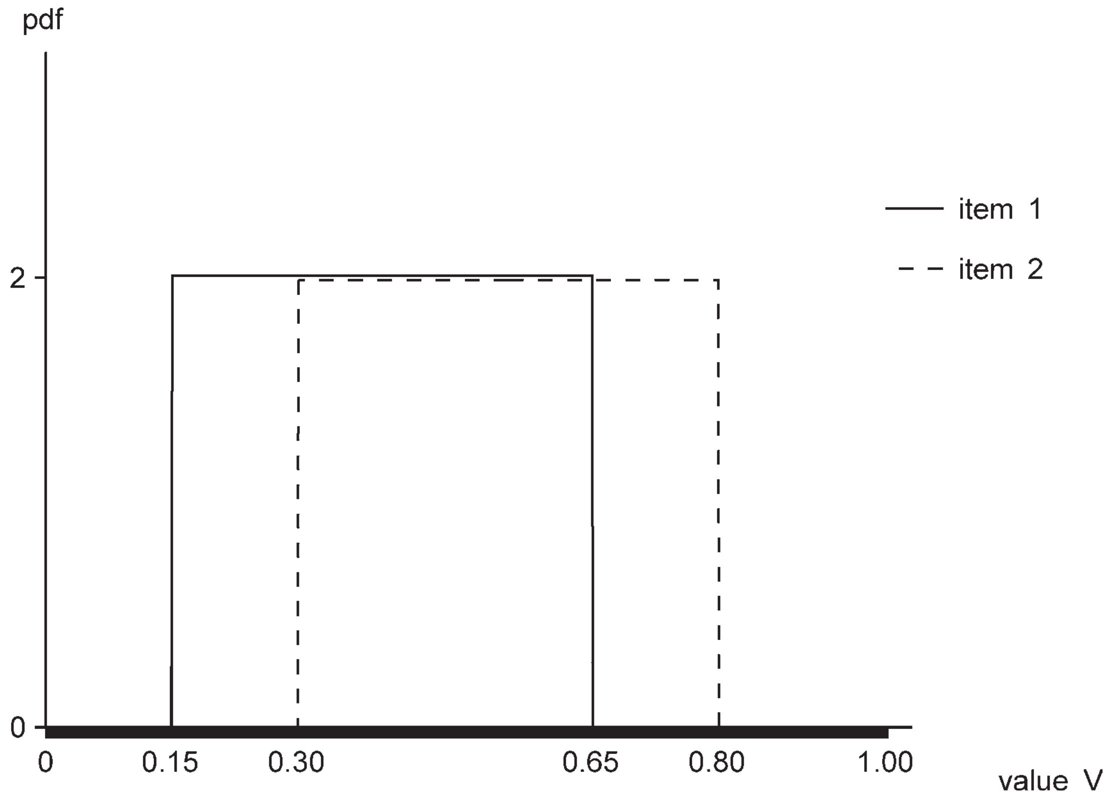

Figure 5.1: Probability density function for the values of two items, with $EV_1 = 0.40$ and $EV_2 = 0.55$.

his decision in a given period, we add six dummy agents. Three of these dummies choose one item, and the other three the other item. This 50–50 seeding prevents any bias at the start of a period. The reason to do this is that lock-in due to the choice of the very first agents would be an uninteresting artifact.

As Figure 5.1 illustrates, in general the information sampled will be far from conclusive to determine which of the two items has the greatest expected value. For example, a utility level of 0.60 experienced by a specific agent in a sample could have been generated by an item with an expected value of 0.35, but also by an item with an expected value of 0.85. Obviously, this uncertainty matters a great deal for an agent that needs to make such a decision. We assume that each agent has in mind a set of simple rules of thumb to choose an item, and that the propensity to use any of these rules may change over time as a result of an agent's experience in the use of these rules. Therefore, before we explain in detail the modeling of the decision making and learning by the individual agents, we need to clarify how the individual agents face a similar basic choice problem over and over again.

5.2.2. Choice dynamics

All individual agents face the same basic choice problem for 25,000 periods. In every period, two new items arrive that are completely independent from all earlier items, and all agents sequentially face a choice between them, with the order of the agents being determined at random in every single period. The fact that we modeled the sampling in a given period as random is a short-cut to take into account that for every day-to-day decision an individual agent may have a different relevant "neighborhood." As we want to focus on the issue of information contagion (analyzing the meaning of the information externality), we do not want to impose any given, fixed structure on these neighborhoods, nor do we want to consider the endogenous formation of neighborhood structures.

As we explained in Section 5.2.1, every item appearing is characterized by its expected value. This expected value itself, which is unknown to the agents, is also a random draw from a uniform distribution; this time with support from 0.25 to 0.75. Hence, the worst item that can ever appear has an expected value of 0.25 (generating values for individual agents between 0.00 and 0.50), and the best possible item has an expected value of 0.75 (yielding utility levels between 0.50 and 1.00). Obviously, the ranges of utility levels that can be generated by intermediate items overlap with each other, as shown in Figure 5.1. Every 500th period, the expected values of the two items are identical (0.50). These identical expected value cases will serve as useful benchmarks to see how much information contagion has emerged. While we use this benchmark every 500th period, in all other periods the expected value of the two items will not be identical, with one of the two items being superior in a statistical sense.

Although we have not said much about individual decision making and learning yet, intuition might suggest that this must be a trivial problem. If we run the model for 25,000 periods, and if in every period (apart from the benchmark periods) one of the two items is superior, then, surely, eventually every agent will easily discover which item is better. However, matters are slightly more complicated. Every period, two new, unknown items appear, and each item is up for choice only once during the entire history. Hence, the learning concerns the general rules of behavior, and not the specific, particular items as such. The fact that the agents learn the usefulness of general rules of behavior, and not the value of specific items also implies that if an agent oversees a certain sample of prior adoptions by other agents he might choose item 1, whereas he might choose 2 if he were confronted with the same two items but a different sample of prior adoptions.

5.2.3. Individual decision making and learning

The individual agent's decision making is modeled for each individual agent separately by means of a Classifier System. Figure 5.2 presents one such stylized Classifier System.

A Classifier System consists of a set of rules, each rule consisting of a condition part ("*if* …"), and an action part ("*then* …"), plus to each rule attached a measure of its strength. The Classifier System does two things. First, it decides which of the rules will be the active rule in a given period. Hence, it checks the condition part, and all rules satisfying the "*if*…" condition make a "bid" as follows: bid $=$ strength $+\, \varepsilon$, where ε is white noise. The rule with the highest bid in this "stochastic auction" wins the right to be active. Second, the Classifier System updates the strength s of a rule that has been active, and has generated a reward from the environment in a given period $t-1$, as follows: $s_t = s_{t-1} - c \cdot s_{t-1} + c \cdot \text{reward}_{t-1}$, where $0 < c < 1$. Hence, $\Delta s_t = c \cdot (\text{reward}_{t-1} - s_{t-1})$. In other words, as long as the reward generated by the rule in period $t-1$ is greater than its strength at $t-1$, its strength will increase. As a result, the strength of each rule converges to the weighted average of the rewards from the environment generated by that rule.[3] In the Classifier System implemented in our model, the strengths of all rules are equal at the start.

Classifier Systems are a form of reinforcement learning. Reinforcement learning is related to multi-armed bandit problems, and is based on two principles. First, agents try actions. Second, actions that led to better outcomes in the past are more likely to be repeated in the future. There is a family of stochastic dynamic models of such individual behavior in the scientific literature, for which different backgrounds can be distinguished. The idea was first developed in the psychological literature. See especially Hull (1943) and Bush and Mosteller (1955), on which Cross (1983) is based. Much later, reinforcement learning was independently reinvented twice as a machine learning approach in computer science. See, e.g., Sutton and Barto (1998) for a survey of an approach called reinforcement learning. The other reinforcement learning approach in computer science is known as Classifier Systems. See Holland (1975) for early ideas on this, or Holland et al. (1986) for a more elaborate treatment of the issue of induction in general. In the economics literature reinforcement learning became better known more recently through Roth and Erev (1995).

It should be stressed that the Classifier Systems are not models of agents using only simple decision rules. Although each rule for itself in a Classifier

[3] We presented this specific learning model in Kirman and Vriend (1995), see also Kirman and Vriend (2001).

condition	action	strength
if	then	....
.. ....	....	....
.. ....	....	....

Figure 5.2: Classifier System.

System is a simple rule, it is the *set* of rules that forms the link between actions and previous actions and outcomes, and it is not the individual rules that matter. As is well known, this type of representation of knowledge is not restrictive in any sense, and any program that can be written in a standard programming language can be implemented as a Classifier System. That is, these systems are "computationally complete" (Minsky, 1967). Hence, a Classifier System may be thought to model the most complex and sophisticated human decision procedures, as well as the most simple. In other words, *any* decision can be modeled *as if* made by a Classifier System.

Table 5.1 summarizes the set of rules we actually use in our model. A more detailed explanation of each "if… then…" rule can be found in the appendix. To illustrate that these rules of thumb compete with each other, and that, given the six sample observations, different rules of thumb may lead to different product choices, consider the following example. If the choices in an agent's sample are three times item 1, and three times item 2, with utility levels of 0.48, 0.71, and 0.28 for item 1, and 0.41, 0.37, and 0.44 for item 2, then rule 1 (choose highest average) would point to item 1, while rule 4 (choose highest minimum) would lead to item 2. The relative importance of each rule of thumb in a decision maker's decision process depends on the payoffs generated by these rules of thumb, such that rules that gave rise to higher payoffs are more likely to be used. As explained above, the agents continuously update their beliefs in this respect.[4] Besides

[4] A more general analysis, including also the issues of creativity and innovation, would allow for new rules of thumb to be generated (rules we perhaps could not even imagine right now). This could be modeled with a Genetic Algorithm combined with our Classifier System.

Table 5.1: Decision rules.

Rule	**Choice**
1	Highest average
2	Highest average (2)
3	Highest average (3)
4	Highest minimum
5	Highest minimum (2)
6	Highest minimum (3)
7	Highest maximum
8	Highest maximum (2)
9	Highest maximum (3)
10	Majority
11	Majority (3)
12	Majority (5)
13	Follow last
14	Follow last (2)
15	Follow last (3)
16	Random
17–31	Opposite choice of rules 1–15

through the white noise added to the "bids" of the Classifier System (see above), the agents experiment through some kind of "trembling hand," mistakenly picking the item they did not intend to with a given small probability.

5.3. ANALYSIS OF THE MODEL

In this section we will show that the ACE model described in Section 5.2 provides a possible explanation for information-contagious behavior. Moreover, we will see that information-contagion is an inherently complex dynamic phenomenon. In order to analyze the properties of our ACE model, we examine 10 runs of the model, each with 100 agents for 25,000 periods.

From an objective point of view, in almost every period one of the two items is superior, but knowledge is very much divided in our model. Each individual agent has a sample of six observations, and such a sample may overlap with the samples of some other agents. Hence, some more specific questions we want to answer are the following: Do the agents through their interaction learn to use rules of thumb that solves the division of knowledge problem? How do the market outcomes look like? And do we get path-dependence and lock-in effects?

5.3.1. Path-dependence and lock-in

We first focus on the benchmark periods in which the expected value of both items is 0.50, i.e., the periods that are a multiple of 500. We want to know how the market shares of the two items develop as we go down the sequence of 100 agents in a given period, and in particular we want know how this development changes over time as the agents learn which rules of thumb to use. Figure 5.3 shows some examples of a typical run: the development of the cumulative market share of one of the items in the periods 500, 10,000, and 20,000.[5] Each sequence starts with a market share of 0.50 because of the initial choices by the six dummies. The market share of the other item is just one minus the share of the item shown, i.e., the curve shown mirrored in the straight line at 0.50.

If there were no information externalities at all, every choice would be an independent decision, with each of the two items being equally likely to be chosen (as in these benchmark periods the two items were equally good), and the development of the market shares would more or less zigzag around a 0.50 market share. As we see in Figure 5.3(a), the cumulative market share curve for period 500 looks as if there is no information externality. This is because the agents have had only little opportunity to learn, and they basically behave like "zero intelligence" agents (Gode and Sunder, 1993), choosing behavioral rules at random. As a result, no information contagion occurs. If we showed this curve in period 500 for different runs, or other benchmark periods towards the beginning of a run, we would get a series of different zigzag curves that all stay close to the 0.50 market share line.[6] The market share curve shown in Figure 5.3(b) for period 10,000 looks very different. Just as in period 500, we see some deviations from a 0.50 share early on, but unlike in period 500, this time we see that the item that gets a smaller market share early on continues to lose ground. Eventually, its share stabilizes at a level of about 16%. The rather smooth curve for period 20,000 shows the positive feedback effect even stronger. Right from the beginning of this period, one item (the one not shown) increases its market share continually until it dominates the market completely. Although the two items are identical in this period, the information contagion leads to lock-in. Which of the two items gets to dominate is basically random, due to small historical events. That is, it is path-dependent.

Figure 5.4 looks at the same phenomenon, the emergence of information contagion, focusing on the individual choices of the 100 agents as such in the same

[5] We will see below that these examples have been carefully selected in a certain sense.

[6] The fact that the zigzag pattern appears to become smoother towards the end of the sequence is due to the fact that each additional decision maker carries less weight in the cumulative market share as we move down the line of 100 agents.

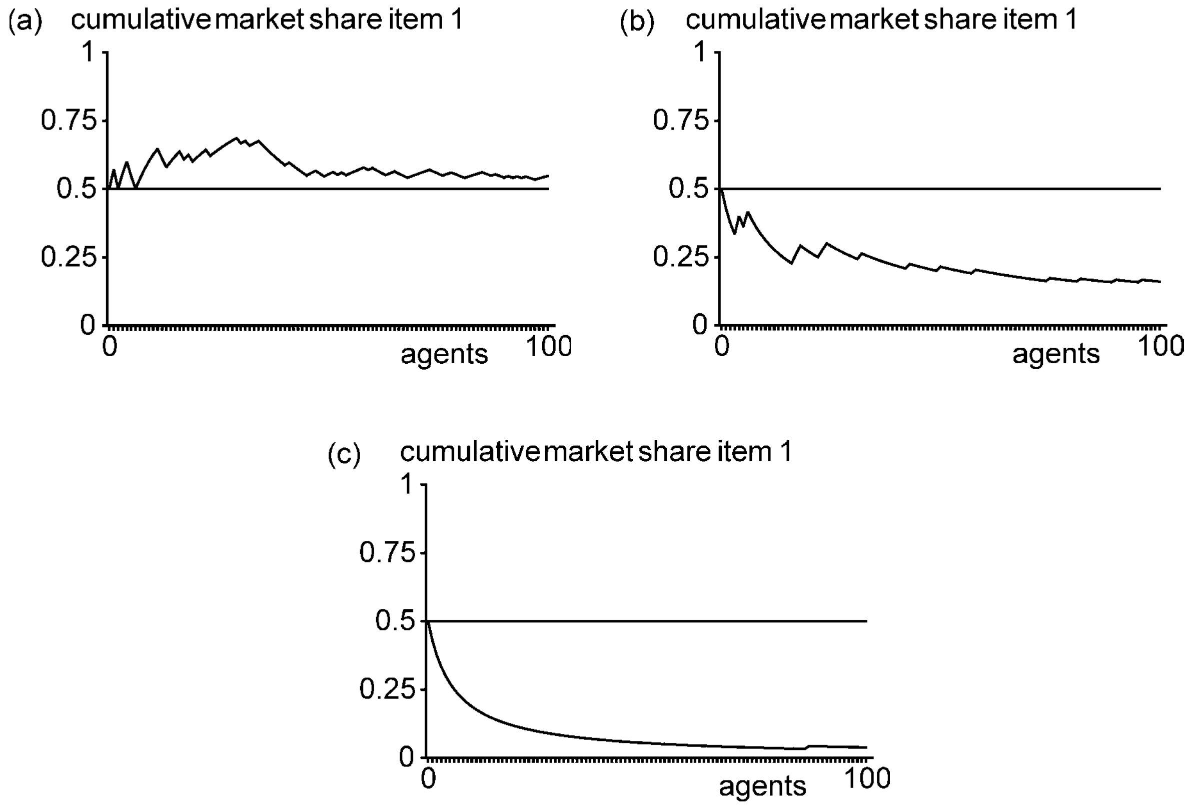

Figure 5.3: (a) Cumulative market share, period 500 (run #8). (b) Cumulative market share, period 10,000 (run #8). (c) Cumulative market share, period 20,000 (run #8).

periods as shown in Figure 5.3. In Figure 5.4(a) we see an almost random sequence of choices in period 500, shifting from one item to the other item all the time, and there is very little order, if any. In Figure 5.4(b), showing the same for period 10,000, and in Figure 5.4(c) for period 20,000, we see an increasingly orderly pattern. In Figure 5.4(c), although the two items are identical, item 1 does not seem very fashionable, with agent after agent choosing item 2, and only an occasional deviation from the norm.

As Figures 5.3 and 5.4 show, the decentralized interaction of the individual agents leading to a situation in which almost all agents choosing the same item emerged as a spontaneous order, where by "order" we mean here that knowledge of a sequence of individual choices would allow us to make a more than educated guess about the next individual choice.

The market share curves and the individual choice curves shown suggest a simple story. As time goes on, the more the choice behavior of the population becomes self-organized the more the information contagion develops, and as a result the development of market shares increasingly gets a particular pattern, with rather smooth curves concentrated in a relatively small space with either a very

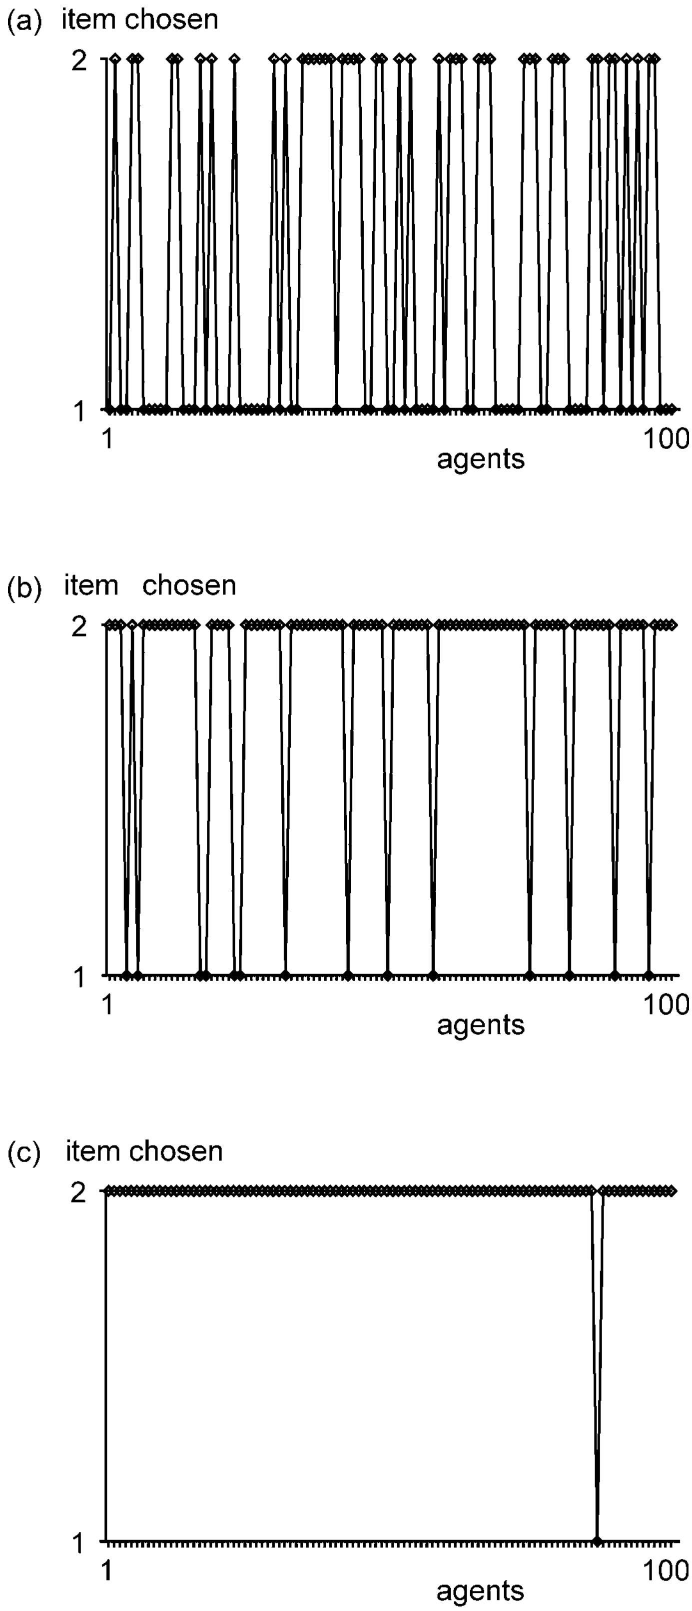

Figure 5.4: (a) Individual choices, period 500 (run #8). (b) Individual choices, period 10,000 (run #8). (c) Individual choices, period 20,000 (run #8).

high or a very low cumulative market share. However, as we will show in a moment, matters are slightly more complicated. The spontaneous order emerging turns out to be far from absolute, and the examples just shown have been carefully selected. In every run it takes some time before the information contagion emerges, giving rise to lock-in and path-dependence effects, but once the population gets self-organized this turns out not to be a monotonic process at all. Giving the curves shown in Figures 5.3 and 5.4 for different benchmark periods would show market share curves going all over the place. Sometimes one item almost completely dominates the market, other times we see the fashion switching at some point from one item to the other, and sometimes this switching occurs so frequently that we get a zigzag curve similar to the one shown for period 500. Hence, the curves seem to drift about in all directions, and the system moves all the time between almost complete order and almost complete disorder, but never stays at either of these. We will explain this phenomenon in Section 5.3.2, but first we will illustrate it by using different measures for what goes on in these markets.

Obviously, the final market share of an item is not exhaustively informative concerning the amount of lock-in generated. One change in fashion at the middle of the sequence would be sufficient to end up with 5050 shares. Therefore, we take as a measure of the path-dependence in the population's decisions the size of the area between the cumulative market share curve (as shown in Figure 5.3) and the straight line at 0.50, relative to the area of the rectangle defined by the axes and the 0.50 line. The more systematically the market stays away from a 50–50 distribution, the more lock-in we have. This measure, the lock-in rate, is a number between 0 and 1, and is shown in Figure 5.5 for the same run number 8 for all benchmark periods, i.e., those periods that are a multiple of 500. As we see, lock-in ranges from low values around 0.10 at the beginning, and tends to get higher values as time goes on, up to about 0.80, but there remains a lot of variation all the time, with lock-in regularly falling back to the low initial values. The three benchmark periods that we used in Figures 5.3 and 5.4 are indicated with a dot.

The same kind of picture results in each of the other runs. The only difference being that the exact benchmark periods in which the upward or downward shifts occur differ from run to run.[7] Table 5.2 considers the second half of the span for which we examined the model, i.e., the benchmark periods from 13,000 to 25,000. For each of the 10 runs we compute the average, the standard deviation, the minimum, and the maximum lock-in rate. The table shows for each of these four variables the run with the lowest and the run with the highest values among

[7] These graphs are available from the author upon request.

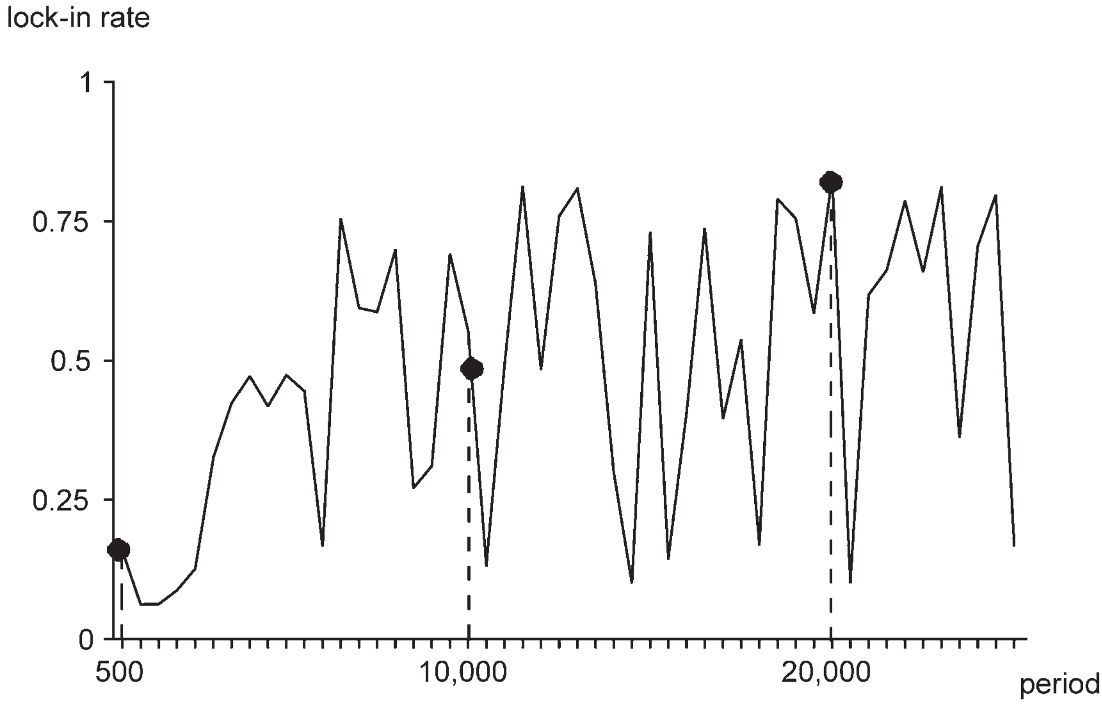

Figure 5.5: Lock-in rates in benchmark periods (run #8).

the 10 runs. As we see, for each of these four statistics the run with the lowest value and the run with the highest value are within a relatively narrow range.

Another way to measure how much lock-in into one of the two items is present is the rate at which the choices of the agents switch from one to the other item in the benchmark periods. If each individual decision is taken independently, and the items are equally good, the expected switch rate is 0.50. Figure 5.6 shows the switch rates for each of the benchmark periods in run number 8. As we see, the switch rate starts indeed around 0.50, and then comes down as time proceeds, but just as with the lock-in rates above, this goes with a lot of fluctuations. The switch rate regularly comes down to

Table 5.2: Lock-in rates across runs.

	Out of 10 runs	
	Lowest	**Highest**
Lock-in rates		
Average	0.483	0.572
Standard deviation	0.182	0.265
Minimum	0.046	0.113
Maximum	0.793	0.832

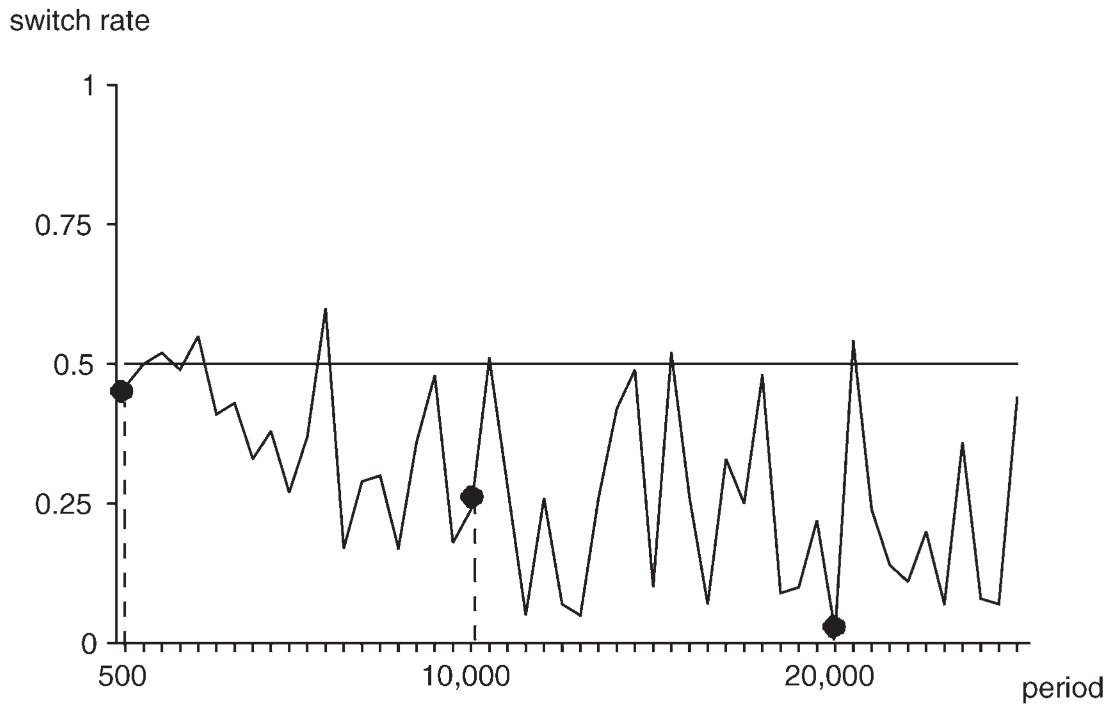

Figure 5.6: Switch rates in benchmark periods (run #8).

values close to 0, implying a very orderly state in which every agent chooses the same item, but almost equally regularly the switch rate jumps back to levels close to 0.50, the maximum disorder, as if all agents choose randomly.

Just as for the lock-in rates, qualitatively similar pictures emerge across the 10 runs.[8] Table 5.3 shows the run with the lowest and the run with the highest value for the same statistics as used in Table 5.2; the average, the standard deviation, the minimum, and the maximum rate over the benchmark periods from 13,000 to 25,000 for a given run. As we see, for each of the four statistics the differences across the runs are relatively minor.

5.3.2. Performance over time

In the benchmark periods analyzed in Section 5.3.1, the two items were always identical. In those periods, any item was as good as the other item, and hence any decision rule was as good as any other decision rule. We used those periods to see how much information contagious behavior the agents had developed during the periods in between the benchmark periods, periods in which the two items were

---

[8] These graphs are available from the author upon request.

Table 5.3: Switch rates across runs.

	Out of 10 runs	
	Lowest	**Highest**
Switch rates		
Average	0.213	0.312
Standard deviation	0.118	0.164
Minimum	0.000	0.060
Maximum	0.470	0.560

generally not identical. Before we analyze the behavior of the individual agents, we first want to see what the effects of the learning of the agents is on the overall outcomes for the society.

Figure 5.7(a) shows the 100-period moving average of the relative frequency with which the best of the two items is picked in each of the first 5000 periods, and Figure 5.7(b) does the same for the periods 20,100–25,000. The three curves drawn are based on the moving average performance curve of each of the 10 runs. They show the upper and lower contour of these 10 curves plus the average curve. As we see, all moving average performance curves are placed within a rather narrow band. The relative frequency with which the superior item is chosen, i.e., the final cumulative market share of the superior item in a given period, is a good measure of social efficiency. At the start, with people making almost random choices, about 50% of the agents pick the correct item. This average frequency increases over time, and in Figure 5.7(b) it has reached a level of about 84.2%, without any further increase suggested by a trend. That we do not reach higher efficiency levels is related to the fact that often the two items have an expected performance that is extremely close. In fact, on average the expected performance for the worst item turns out to be 0.42, and for the best item 0.58. In many periods the difference in expected performance is close to zero.

Instead of the 100-period moving average of the performance of the 10 runs, Figure 5.8 shows the performance for every single period of run number 8, and reveals that underlying these moving average performances something interesting is happening.[9] Notice that every 500 periods the items are equally good, and

[9] The graphs may seem to present multiple observations for each single period. This false impression is solely due to the fact that 5000 observation points are crammed into a small space. The graphs for the other nine runs are available from the author upon request. They show a very similar picture. In Figure 5.7, we showed already that the average performance was very similar across runs. The same applies to the spread of the period-to-period performance. If we take, for example, the standard deviation of this performance measure for the periods 20,001–25,000 in a given run, we see that this ranges from 0.230 to 0.249 across the 10 runs.

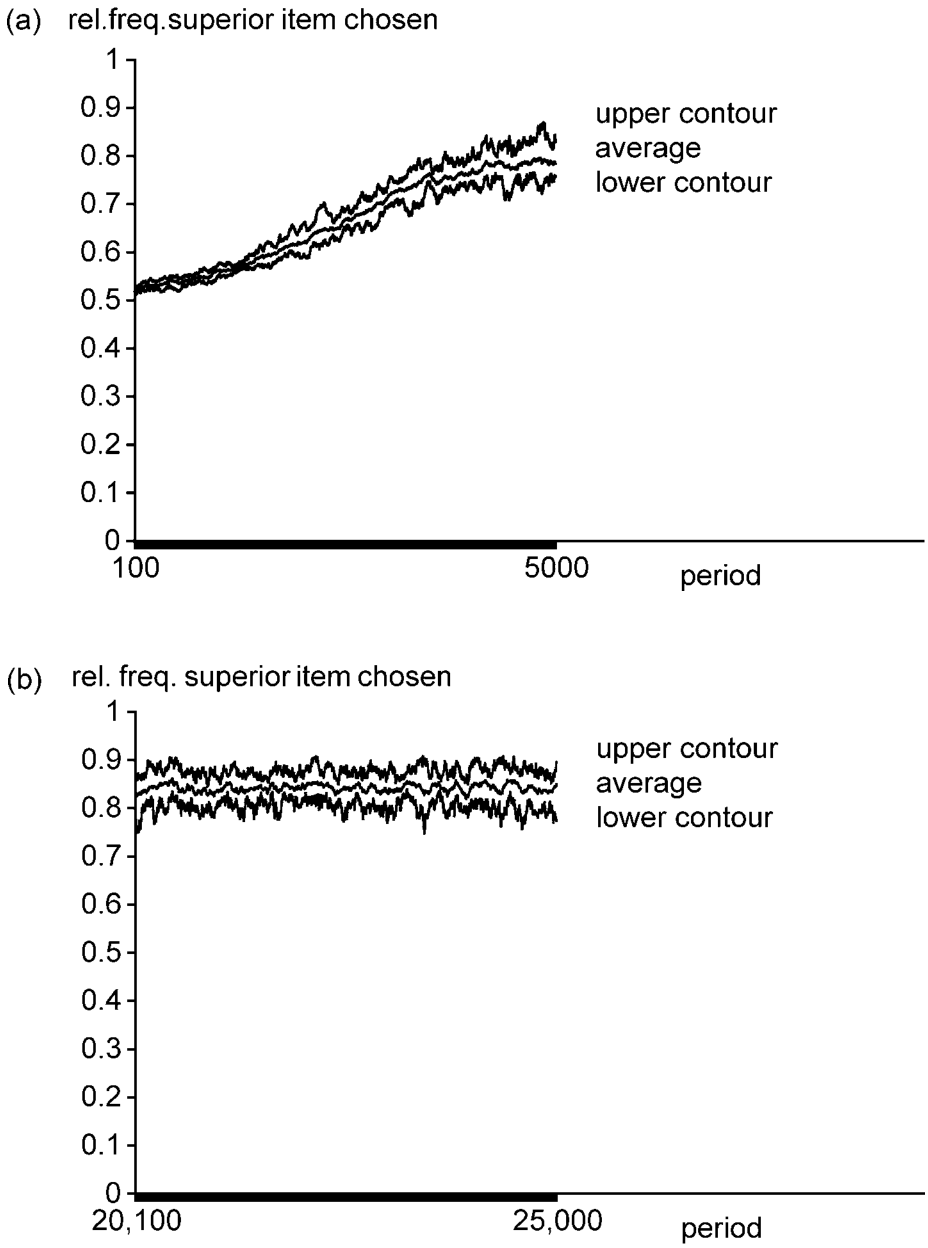

Figure 5.7: (a) Moving average performance, periods 100–5000 (all runs). (b) Moving average performance, periods 20,100–25,000 (all runs).

hence everybody makes the right choice. More interesting is the observation that while the (moving) average performance goes up, the spread increases as well. In the beginning, in every period about 50% of the agents choose the correct item. Sometimes this is a little bit lower, and sometimes a little bit higher, but never very much so. For some time performances never exceed the 35–70% band. But as times goes on, and average performance goes up, occasionally periods occur in which only 30% of the agents pick the superior item. Later on there are periods

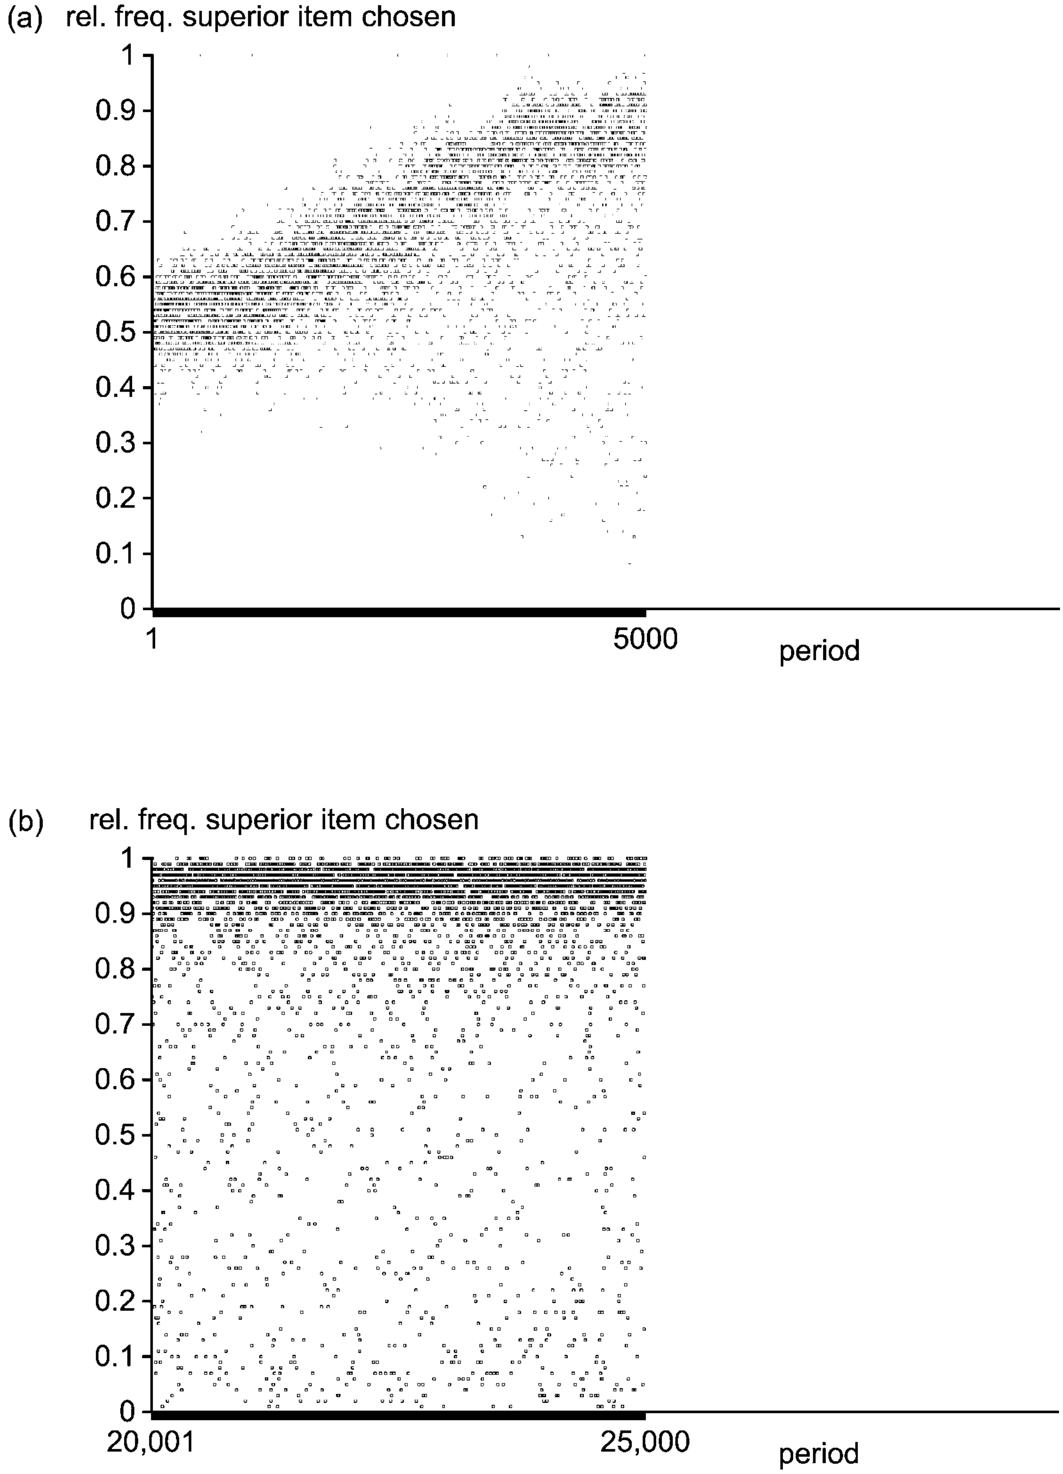

Figure 5.8: (a) Performance, periods 1–5000 (run #8). (b) Performance, periods 20,001–25,000 (run #8).

with just 15% choosing correctly, and eventually, (after about 5000 periods) it sometimes even happens that almost nobody recognizes which is the best item. All the time, though, the moving average of the performance shows an upward trend. As noticed above, in part the spread in performance is due to the fact that on average the expected performances of the two items is rather close, occasionally leading many people to the wrong choice. But the frequency with which the expected performances are close to each other (making mistakes likely) does not change over time. Hence, the change in spread over time that we observe is due to the adaptive behavior of the agents. As they learn, they improve their average performance, but occasionally this leads to disasters, with almost everybody choosing the wrong item.

The big question to be answered, then, is: what is it in the behavior of the individual agents that has adapted in such a way that information contagion emerges? And how is this related to the reported effects of increased average performance and increased intensity of social disasters? Basically, the model implies two things for the behavior of the individual agents that need to be distinguished. First, the agents learn to use better rules as such, i.e., the rules that lead to higher utility levels because they are better at recognizing the superior item on the basis of six sample observations. The dynamics are in part the result of this evolution of the rules being used. Second, the agents learn to use rules that *aggregate* information. The possibility to aggregate information is due to the presence of an information externality. As an agent chooses an item, it gives the choosing agent a certain utility, but at the same time, there is also an *externality*, as the choice of the given agent is added to the information pool on which the choices of future agents will be based. Some rules take advantage of this externality by aggregating information, while others do not.

For example, consider the "highest average" rule. This rule does not aggregate information. It bases its choice on the six observations sampled, i.e., on the items chosen and the payoffs actually generated for those six agents. It is not sensitive to how many people in the sample of six had chosen one item or the other. That is, the choice made by a non-aggregating rule is not affected by the information samples used by each of the six people in an agent's own sample. In other words, an agent using the highest average rule is not bothered by explaining why the agents in his sample had made their choices.

Now, consider the rule that tells an agent to follow the choice of the majority in his sample. This rule does not consider the actual payoffs generated for the six agents in the sample. But if each of the six agents in the sample had considered the payoffs in their samples of six (e.g., following the highest average rule), then the "majority" rule implicitly considers six times six or 36 sample payoffs instead of only six. That is, the majority rule aggregates the information available to each of the agents in the sample.

More in general, the information aggregating rules are those rules that are affected by the choices of the other agents. That is, they are sensitive to how many people in an agent's sample of six had chosen each of the two items. Obviously, because of the information externality, the two forms of learning (i.e., the learning to use better rules as such, and the learning to use rules that aggregate information) are closely related. As one agent learns and changes his behavior, other agents are learning as well, partly in response to this. The value of aggregated information depends on the quality of the choices made by the other agents. Hence, this is a coevolutionary process. The rules that an agent uses evolve in response to the evolution of other agents' rules.

To analyze the relevance of these two forms of learning we did the following experiment that excludes the information externality. The basic choice situation in this variant of the model is the same as above. But this time every agent, when making his choice, does not observe what other agents did before him, nor the payoffs they realized. Instead, when an agent's turn comes, he can six times randomly choose and try an item himself, and observe the payoffs.[10] Hence, the only difference with the standard model is that there is no interaction between the agents, hence no information externality, and thus no possibility of information aggregation.

Although there is no information externality in this variant, the agents still learn which rules are more likely to pick the superior item on the basis of a sample of six observations. Figure 5.9(a) shows the 100-period moving average of the relative frequency with which the best of the two items is picked in each of the first 5000 given periods, and Figure 5.9(b) does the same for the periods 20,100–25,000. We show again the upper and lower contour of the moving average performance curves plus the average of the 10 runs. We observe that the performance, starting from a level of 0.50, which even random choice would achieve, increases to a level of about 0.775. That is, the agents do learn to improve their performance by the use of the better rules, but they stay below the average performance in the standard version, when it reached a level of 0.842. In other words, taking advantage of the information externality by aggregating

---

[10] These payoffs are generated using exactly the same underlying distributions as in the base model, including the noise term added to each observation. This implies that the stochastic element of the payoffs can no longer be interpreted as idiosyncratic taste or skill factors, but should be seen as measurement errors in this variant. Notice also that to follow the previous setup closely, we do not consider the issue of what the optimal sampling strategy would be.

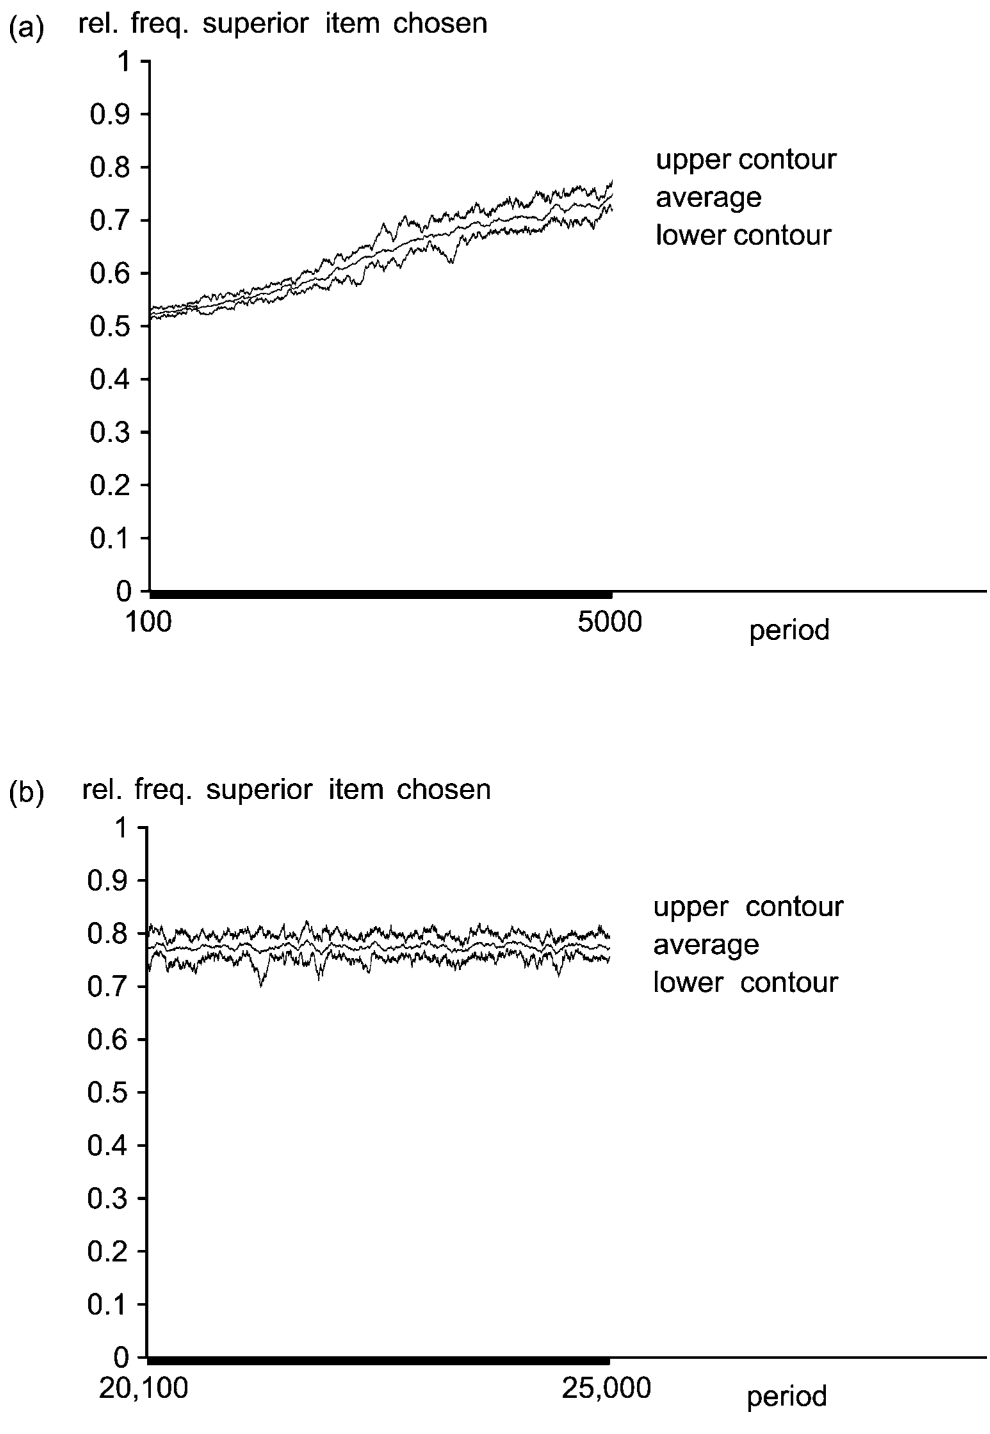

Figure 5.9: (a) Moving average performance, periods 100–5000 (variant; all runs). (b) Moving average performance, periods 20,100–25,000 (variant; all runs).

knowledge, the agents succeeded in winning another 8.6% in performance in the standard version.[11]

Figure 5.10 shows the performance in every single period of a given run of the variant. As we see, the performance tends to rise, but, apart from the benchmark periods, it almost never gets close to 1, and there are also no disasters. In the most unfortunate periods, it is still about 40% of the agents that succeed in choosing the superior item.[12]

This analysis illustrates at the same time the advantage and the limits of information aggregation, as occurring in the standard version. By aggregation, the agents succeed in reaching very high performance levels in many periods, higher than they could ever achieve on their own. But when agents aggregate information (e.g., following the majority rule), an agent wastes some information as well, since he does not use the information concerning the actual payoffs realized by the six people in his sample. As we explained above, if a single agent aggregates information he implicitly uses six times six, i.e., 36 observations instead of only the six in his own sample. But if each of the six agents in his sample would also be aggregating information, they would each implicitly use 36 observations, and hence our single agent would be using six times 36, i.e., 216 observations. Hence, the more agents use aggregating rules, the more aggregation of knowledge occurs. But, when too many agents aggregate information, too many agents waste their own information. At some point a tiny little bit of knowledge starts getting aggregated ad absurdum. In some sense, the agents start aggregating ignorance instead of knowledge.

Figure 5.11 shows the upper and lower contour of the cumulative market shares for all benchmark periods of all 10 runs. As we see, the cumulative market shares stay around 0.50. At the end of each given period, each item has a cumulative market share between 0.358 and 0.679. There is no lock-in or path-dependence. This was to be expected, because in the benchmark periods the two items are identical, and all agents make their choice independently. Since there is no information externality, we cannot get path-dependent lock-in.

[11] If we compute for each single run the average performance over the periods 20,001–25,000, we see that this ranges from 0.835 to 0.850 in the standard version of the model, and from 0.770 to 0.779 in the variant. In other words, even in the single worst run of the standard model average performance is 7.2% higher than in the best run of the variant.

[12] The graphs for the other nine runs are available from the author upon request. They show a very similar picture. Besides a similar average performance (see Figure 5.9), the standard deviation of this performance measure for the periods 20,001–25,000 in a given run is also very similar across the 10 runs. It ranges from 0.139 to 0.142.

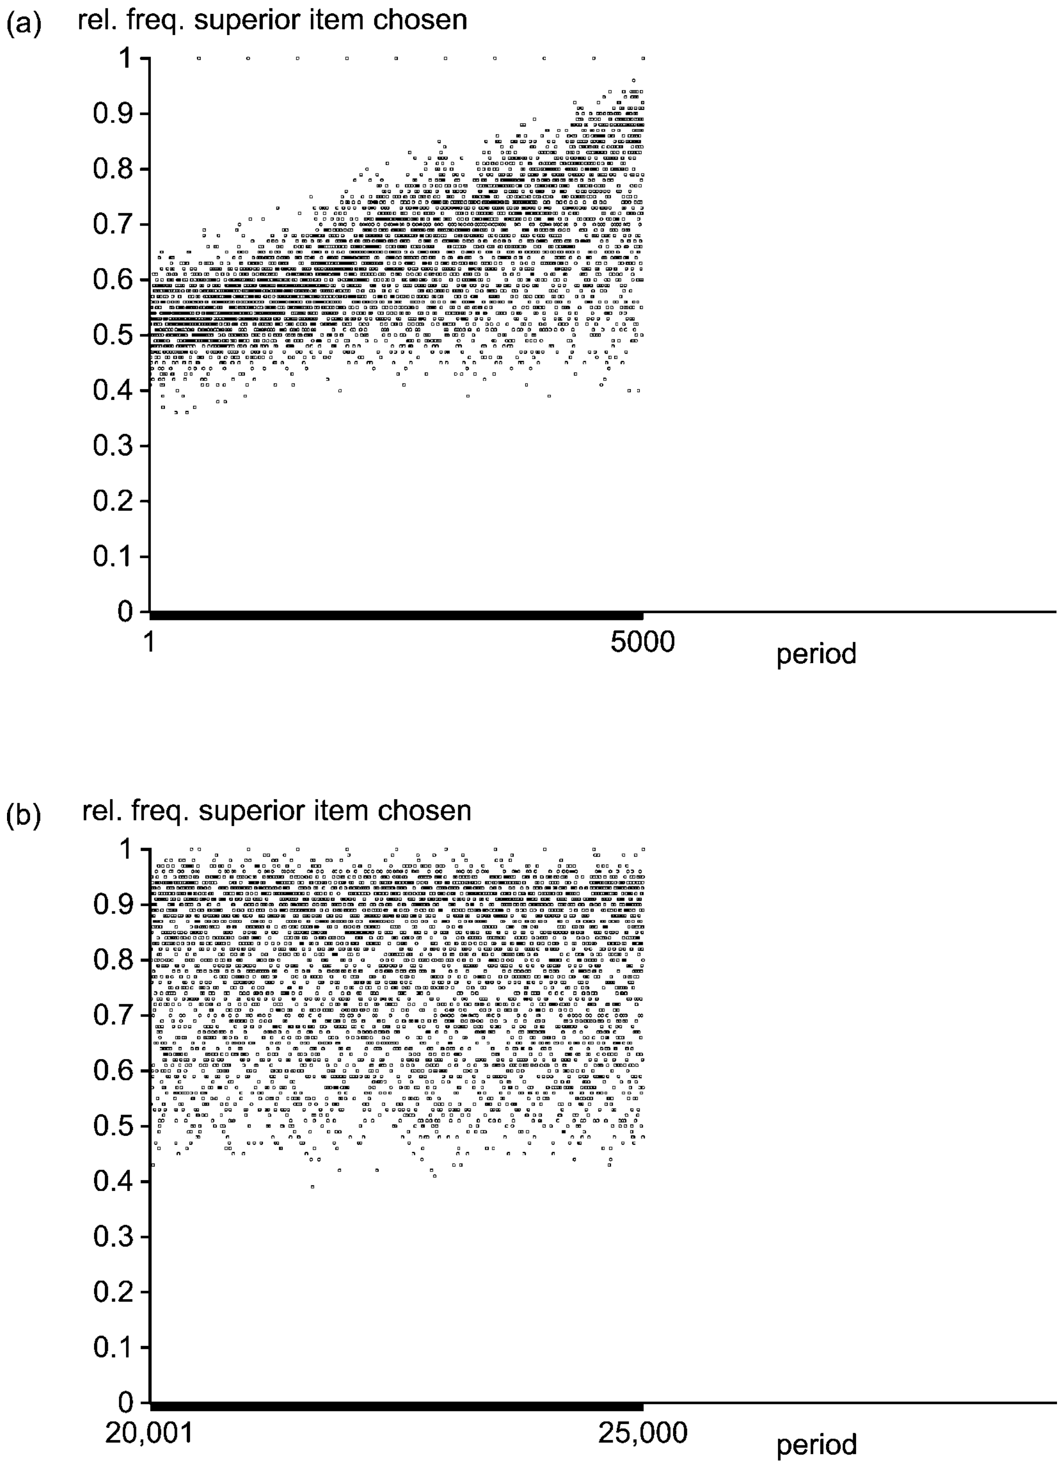

Figure 5.10: (a) Performance, periods 1–5000 (variant; run #1). (b) Performance, periods 20,001–25,000 (variant; run #1).

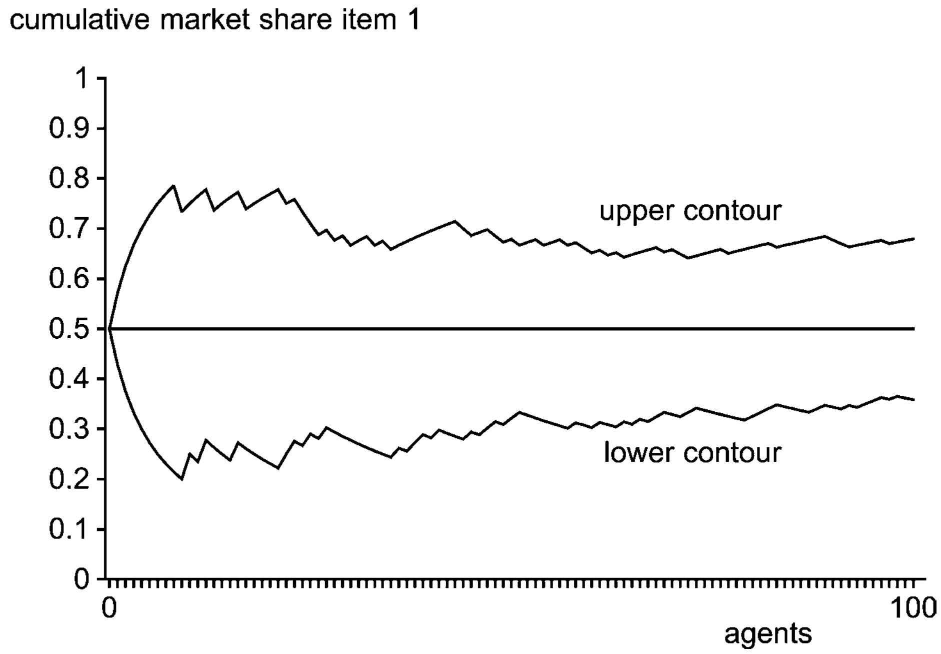

Figure 5.11: Cumulative market shares, all benchmark periods (variant; all runs).

5.3.3. The individual decision rules

One of the advantages of an ACE approach is that we, as modelers, know for each single period which of the two items is superior. Hence, for each single decision to be made by any of the agents, given his sample of six observations, we can check for each of the 31 rules whether it would have picked the superior item. Obviously, the individual agents do not obtain this information. They only try one rule of behavior in every period, and observe the payoffs they generate doing so. Figure 5.12 shows the time series of the relative frequencies that a given rule would have picked the superior item, averaged over the 10 runs.[13] That is, the graph shows the relative frequency that a given rule belongs to an agent's best-reply correspondence.

As we see, picking an item at random (rule "random") leads to the superior item in about 50% of the cases, and this remains constant over time. Looking at just one other agent, and imitating whatever he picked (rule "last") starts close to 50%, but as other agents learn to make better choices, the performance of this rule increases

[13] These frequencies are normalized for eligibility, since, as we explained above, in some cases the "*if* ..." part of a rule is not satisfied. Each observation concerns one cycle of 500 periods (from one benchmark period to the next). For presentational reasons we only show the rules 1, 4, 7, 10, 13, and 16 (see Table 5.1).

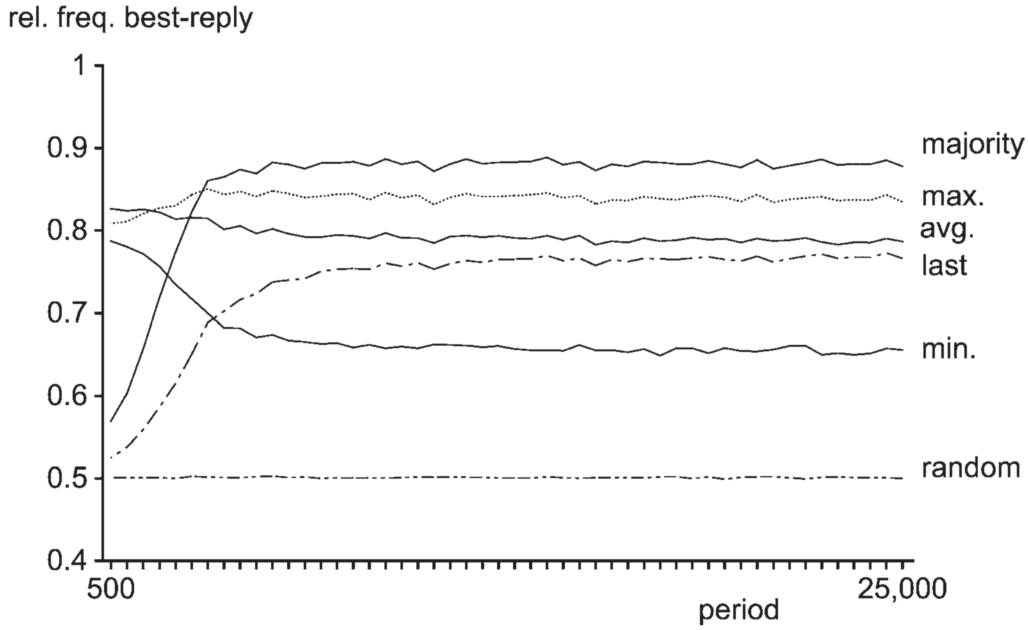

Figure 5.12: Specific rules as best-replies (all runs).

considerably and gets close to the rule that chooses the highest average in the sample (rule "average"). This increase in performance applies even much more to the rule that says to follow the majority of the six observations sampled (rule majority). This rule, which does not use any of the available information concerning the utility levels obtained by the six agents sampled, at some point starts beating all other rules.[14] Two other rules stand out. The rule that chooses the highest minimum (rule "minimum") deteriorates over time. The explanation for this is that, implicitly, it does the opposite of information aggregation. It favors the item that is the least often chosen, because the more an item is chosen, the more likely it is that some observation will be in the lower part of the distribution, and hence be the lowest minimum in the sample. Exactly the opposite applies to the rule that chooses the highest maximum (rule "maximum"). The more an item is chosen, the more likely it is it will provide the highest maximum in the sample.

The important thing to notice here is that the degree to which a given rule is objectively good changes over time as a result of the other agents changing the rules they use. To show how the effect of the information externality makes it a coevolutionary process, i.e., the agents adapting to each others' adaptation to each other…, Figure 5.13 presents the frequencies (averaged over the 10 runs) with

[14] Arthur and Lane (1991) argue that lock-in resulting from the simple imitation of other people is not interesting, but what makes it interesting here is that we contribute to an explanation of the phenomenon of imitative behavior itself.

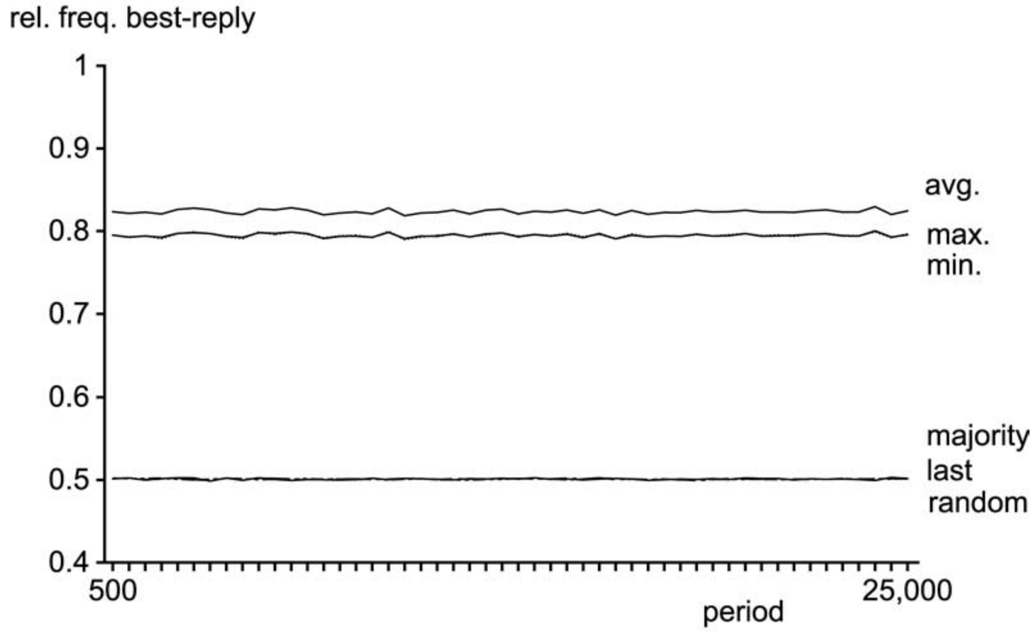

Figure 5.13: Specific rules as best-replies (variant; all runs).

which the individual rules form part of an agent's best-response correspondence in the variant in which there are no information externalities. As we see, these frequencies remain constant over time, apart from some random noise. The only thing the agents need to learn is to figure out which of these rules are most often good in a given situation. Obviously, for different situations different rules might be best. But which rule is good for a given sample configuration does not change over time. This is very much unlike Figure 5.12, where the learning of the agents influences in turn what the other agents have to learn.

To conclude our analysis of the model, could it be that the famous QWERTY lock-in has less to do with network externalities and other real payoff matters than with information contagion? After all, with the current technology, and most people using a personal computer, switching a keyboard layout is relatively easy. It is true that it requires a little bit of personal investment (time and effort to change the layout itself, plus some re-training), but if individual agents knew an alternative keyboard were superior, that would be no obstacle. The only problem seems that individual agents do not know whether it is worth choosing an alternative keyboard layout, and generating own sample observations by trying various different keyboard layouts is rather costly. Hence, an individual agent needs to base his decision on the choices made by other people, and as our ACE model demonstrates, it might be that it is the emergence of information-contagious behavior that leads to a QWERTY lock-in.

5.4. SOME RELATED LITERATURE ON INFORMATION CONTAGION AND SOCIAL LEARNING

Alongside increasing returns (Arthur, 1989), network externalities (Katz and Shapiro, 1985, 1980), information cascades (Bikhchandani et al., 1992), and herding behavior (Banerjee, 1992), information contagion (Arthur and Lane, 1991) has been presented in the literature as an explanation for particular patterns of macrobehavior (for example, path-dependence and lock-in effects) that may seem at odds with the underlying micromotives. But whereas these other explanations have been shown to have a proper microfoundation (either related to changing productivity or changing preferences, or to Bayesian updating in the face of uncertainty), information contagion has remained a phenomenon that occurs only when certain ad hoc rules of thumb for individual behavior are assumed.[15]

Making use of the theory of generalized Polya urn schemes (Hill et al., 1980), Arthur and Lane (1991) show that in a population of naive Bayesian optimizers information contagion may drive the market to stable (but not complete) domination by a single product. This is caused by the fact that the more a product has been chosen already by others, the more likely it is to be in an individual agent's sample on which he has to base his choice, and hence the more precise his information concerning the true quality of the product. Dosi et al. (1994) show similar results for a given rule of thumb that is more directly imitative.

Narduzzo and Warglien (1996) carried out two experiments with human subjects to test the empirical relevance of this theoretical possibility of information contagion. Between 50 and 170 experimental subjects were instructed that they faced the choice between two products of which the objective value is uncertain, while the same product can generate a different value for different subjects. The players made their choice sequentially, and each player received information concerning the choice and value generated of a random sample of six subjects that had already made their choice at that point. In fact, the values generated by the two products had the same uniform distribution on [0.25, 0.75] and in another treatment these values were drawn for both products from [0.60, 1.00]. They observed that path-dependent dynamics emerged and that lock-in of market shares occurred, with early accidental choices giving rise to seemingly stable cumulative market shares with the prevalence of one product over the other. Narduzzo and Warglien, then, tried to find out which choice heuristics the players used. Therefore, they interviewed some players after the experiment and they did some "thinking-aloud"

[15] A difference between the information contagion literature on the one hand, and the literature on information cascades and herding behavior on the other hand, is that in the latter an agent does not observe the payoffs generated by other agents, but only their choices as such.

protocol analysis. They found four basic choice heuristics: the mean rule (highest average), the min rule (highest minimum), the max rule (highest maximum), and the popularity rule (follow majority). They observe that these rules are not necessarily used in isolation, and that subjects may have changed their rules based on their experience from one run to another. Also, there might be context-sensitivity, with different samples inducing the use of different rules.

Lane and Vescovini (1996) analyzed the contagious effects of these four rules, with the assumption that all subjects follow the same rule. They find that the mean and the min rule never produce path-dependent behavior. The popularity rule always generates path-dependent market shares, and the max rule only when the products are exactly identical. Hence, three of the four rules reported by the subjects do not generate path-dependence. So the question is, where did the observed path-dependence come from?[16] Lane and Vescovini note that the emergence of path-dependence and information contagion is related to the mix of rules actually used in the population, and that an important question is how people change rules after they experience outcomes. Since learning is a coevolutionary process (while one agent is learning all other agents are learning simultaneously), these two points should be considered combined. That is what we do in our model. We want to understand the process through which information contagion emerges. Narduzzo and Warglien's experiments are one-shot games, but the players must have faced very many analogous decision problems outside the laboratory. How does information contagion emerge, what role does it play, and what are the effects (both with respect to individual players and the society as a whole)? Where does a configuration of rules used in a population that leads to path-dependence come from? Is it based on arbitrary, extremely bounded rational behavior? Or is it reasonable to learn rules of behavior that give rise to information contagion?

Our ACE model addresses these questions. ACE modeling has two advantages relative to the experimental method followed by Narduzzo and Warglien. First, organizing a laboratory experiment with a large number of players making their choice sequentially is not easy, and organizing an experiment with a large number of such periods in order to study the learning dynamics seems a very arduous task. Second, using experimental data to characterize individual behavior faces the problem that the rules of thumb used by the players are not directly observable, while subjects' reports are not very reliable. One reason for this might be that each

[16] Of course, it could be the case that the subjects in the Narduzzo and Warglien (1996) experiments were actually naive Bayesian optimizers as outlined in Arthur and Lane (1991), but this seems unlikely. For one thing the Arthur and Lane model assumes that the subjects know the parameters of the underlying distributions, which was not the case in the experiments.

rule he exactly uses at a certain moment. Hence, questionnaires tend to be rather inconclusive. Therefore, this seems an excellent case to use an ACE model. Not only does this allow us to analyze very long run dynamics, but we can also do an explicit analysis of the rules of thumb actually used.

The basic choice situation in our ACE model as described in Section 5.2.1 follows closely the one used in the information contagion papers discussed in this section. The main difference is the following. In the information contagion experiments by Narduzzo and Warglien (1996), the expected values of the two items are actually identical, although this was not known to the subjects. In our ACE model, the two expected values are generally different in every period (apart from the benchmark periods). The reason to use these different expected values in all periods not being a multiple of 500 is that if the two products would have an identical performance distribution in every period, then there would be no relation between the rules of thumb used and the payoffs generated. That is, since on average any of the two items is equally good, any rule of thumb is as good as any other rule, and hence there would be no selection pressure at all. Therefore, in our model in each cycle of 500 periods we have 499 periods in which the average performance of the two products is different, in order to give selection some bite, and then we check in the 500th period what the selection process has achieved by using the two products with identical performance as a benchmark case.

Ellison and Fudenberg (1993) consider a closely related problem, i.e., the use of rules of thumb in a situation where agents need to choose between two items with unknown value, and where social learning takes place. They assume that players use exogenously specified, simple rules of thumb. One justification they give for this is that they do not consider fully Bayesian learning a realistic assumption, because it requires calculations that may be too complicated. In each period, some fraction of the players have the opportunity to revise their choices. They only observe last period's payoffs and choices of all agents. They present some simple and some more complicated rules which are all some form of popularity weighted choice of the highest average, and they derive the optimal weight of the popularity for some of these rules. However, since they assume rules of thumb that are exogenously given, an important question is whether it is likely that these optimal weights will actually be used by the agents. In support of a positive answer, Ellison and Fudenberg note that they showed in a working paper that all agents using the optimal weights constitutes an equilibrium, but they stress that important extensions of their analysis are needed. The reason for this is that the precise specification of these rules supposes more sophistication of the agents than they find themselves compelling. Therefore, conjecturing that the optimal popularity weighing might emerge from an adaptive process, Ellison and Fudenberg

Table 5.4: Overview of models.

	(Series of) once-and-for-all choice(s)	**(Series of) choice(s)-with-revision**
Static choice behavior (no learning)	Arthur and Lane (1991)	Ellison and Fudenberg (1993)
Dynamic choice behavior (learning)	Our ACE model	………

explicitly ask for a complementary study. As they put it, "it would be interesting to complement these results with an analysis of a dynamic process by which players adjust their rules of thumb along with their choice of technology" (p. 638).

Such an analysis would imply an application of our approach to the Arthur and Lane (1991) model to the Ellison and Fudenberg (1993) model. Table 5.4 illustrates this. Starting point is the Arthur and Lane (1991) model, in which the agents face a single binary once-and-for-all choice, while their choice behavior is static, i.e., there is no learning. We, then, distinguish two types of dynamics. Our ACE model introduces dynamic choice behavior in a series of binary once-and-for-all choices. Ellison and Fudenberg (1993), on the other hand, study a single binary choice problem with static choice behavior while allowing for period-to-period revision dynamics. The model that Ellison and Fudenberg solicit would be one in which our dynamic choice behavior is applied to a series of binary choices-with-revision.

5.5. DISCUSSION

We presented a model to explain the phenomenon of information contagion in a study in which individuals repeatedly have to make a choice between two previously unknown items while they can rely only on some information from previous adopters. We showed how one could provide a microfoundation for information contagion, based on a simple model of adaptive behavior with agents trying to do the best they can, and without needing to assume ad hoc rules of thumb. We also showed that information contagion, unlike increasing returns to scale, network externalities, information cascades, and herding behavior, is an inherently complex phenomenon.

Our ACE model exhibits self-organization, and the emergence of spontaneous orders in which typically most agents choose the same, superior item. This results from the emergence of information contagious behavior of the individual agents. Information contagion is a way to aggregate distributed knowledge in society, allowing the individuals and the society to achieve higher performance levels. This comes together with path-dependent lock-in effects, but in some sense, the remarkable thing is not so much the emergence of these effects, but the fact that this is related to a high average performance. The model explains this through a coevolutionary process in which the rules of behavior used by the individual agents evolve simultaneously. Notice that models based on fixed rules of thumb would not work. For example, in our model the information aggregation, and in particular the rule to follow the majority *emerge*. If we specify a priori that the individual agents follow the majority rule then we would stay at a performance level of 0.50. In addition, when the majority rule emerges as a good rule, this does not imply that everybody should follow it. If they did, then the performance would fall back again to 0.50. Hence, what matters is also the precise configuration of rules used in the population. The continuously changing configurations that emerge, lead to both a high performance level and information contagion with path-dependent lock-in. But it turns out that this is not a simple monotonic process from disorder to order until the solution has been reached, with a happy ending. Instead, the system continually moves back and forth between order and disorder. That is, the self-organization is a continuing, ongoing story, in which the emerging order unravels time and again.

The emerging spontaneous order is beneficial, that is, on average. But along with the improved average performance we also see an increase in both the number and degree of disasters. This is related to the tension between generating knowledge and aggregating knowledge. If enough knowledge is generated by the individual agents, aggregation leads to good outcomes, but if everybody would merely aggregate repeatedly a little bit of knowledge, this might lead occasionally to very bad outcomes for the society. In fact, this keeps the self-organizing process from being a monotonic one. If it were monotonic, we would get stuck with only disasters.

Our ACE model shows that it is not simple to argue whether information contagion as such is beneficial or not. Pointing to occasional disasters (QWERTY?, VHS?) is not sufficient to argue that it is damaging. It is not even the case that occasional disasters are just unlucky draws of a given stochastic mechanism. In some sense, the disasters and successes are flip-sides of the same *dynamic* process, as the bad outcomes are essential to generate the successes. It might be that, although the outcomes are not optimal in each and every single period from a static point of view, this is the best that is dynamically achievable in a decentralized setting (see also Bak (1997) for a similar argument).

APPENDIX A

Table A.1 explains each of the 31 decision rules of the Classifier System listed in Table 5.1. Table A.2 presents the pseudo-code of our ACE model.

Table A.1: Classifier System.

1	Highest average *If* both items are present in the sample of six observations *then* choose the item that has the highest average performance in the sample. *Otherwise*, if the condition is not satisfied, the rule is not eligible and will be neglected.
2	Highest average (2) *If* both items are present in the sample of six observations *and* the item with the highest average performance occurs at least twice in the sample *then* choose the item that has the highest average performance. *Otherwise*, neglect this rule.
3	Highest average (3) *If* both items are present in the sample of six observations *and* the item with the highest average performance occurs at least three times in the sample *then* choose the item that has the highest average performance. *Otherwise*, neglect this rule.
4	Highest minimum *If* both items are present in the sample of six observations *then* choose the item that has the highest minimum performance in the sample. *Otherwise*, neglect this rule.
5	Highest minimum (2) *If* both items are present in the sample of six observations *and* the item with the highest minimum performance occurs at least twice in the sample *then* choose the item that has the highest minimum performance. *Otherwise*, neglect this rule.
6	Highest minimum (3) *If* both items are present in the sample of six observations *and* the item with the highest minimum performance occurs at least three times in the sample *then* choose the item that has the highest minimum performance. *Otherwise*, neglect this rule.
7	Highest maximum *If* both items are present in the sample of six observations *then* choose the item that has the highest maximum performance in the sample. *Otherwise*, neglect this rule.
8	Highest maximum (2) *If* both items are present in the sample of six observations *and* the item with the highest maximum performance occurs at least twice in the sample *then* choose the item that has the highest maximum performance. *Otherwise*, neglect this rule.
9	Highest maximum (3) *If* both items are present in the sample of six observations *and* the item with the highest maximum performance occurs at least three times in the sample *then* choose the item that has the highest maximum performance. *Otherwise*, neglect this rule.
10	Majority *If* there is a strict majority in the sample choosing one item, *then* this rule follows that majority.
11	Majority (3) *If* there is a strict majority in the sample choosing one item *and* this majority is at least three elements greater than the minority, *then* this rule follows that majority.
12	Majority (5) *If* there is a strict majority in the sample choosing one item *and* this majority is at least five elements greater than the minority, *then* this rule follows that majority.
13	Follow last This rule chooses the same item as the one in the last observation sampled.
14	Follow last (2) *If* the last two observations sampled concerned the same item, *then* this rule chooses that item as well.

(*Continued*)

Table A.1: Continued.

15	Follow last (3) *If* the last three observations sampled concerned the same item, *then* this rule chooses that item as well.
16	Random This rule randomly selects one of the items, each with equal probability.
17–31	opposite choice of 1–15 These rules operate just as the rules 1–15. However, when any of the corresponding rules 1–15 determines a choice of item 1, then the current rule selects item 2, and the other way round.

Table A.2: Pseudo-code of the model.

```
program CONTAGION;
begin
 for all 100 players do for all 31 rules do fitness := 1.00;
 for all 25,000 periods do
 begin
  draw expected_value_item_1 from uniform distr. with support [0.25, 0.75];
  draw expected_value_item_2 from uniform distr. with support [0.25, 0.75];
  if period is multiple of 500 then
  begin
   expected_value_item_1 := 0.50;
   expected_value_item_2 := 0.50;
  end;
  put all 100 players in random order;
  create six dummy observations (either 121212 or 212121 with corresponding values);
  for all 100 players do
  begin
   sample six observations;
   for all 31 rules do
   begin
    check conditional part;
    if condition satisfied then bid := fitness + ε, where ε ≃ N(0, 0.025);
   end;
   determine highest bidding rule;
   pick item implied by that rule;
   with probability 0.025 pick instead item not intended;
   draw actual value of item chosen from uniform distr. with support
    [expected_value - 0.25, expected_value + 0.25];
   with winning rule do fitness := 0.975 * fitness + 0.025 * value_item;
  end;
 end;
end.
```

Arthur, W. B. (1989). Competing technologies, increasing returns and lock-in by historical events. Economic Journal, 99, 1160Ò31.

Arthur, W. B. and Lane, D. A. (1991). *Information Constriction and Information Contagion.* Working Paper 91-05-026, Santa Fe Institute.

Bak, P. (1997). *How Nature Works. The Science of Self-organized Criticality*, Oxford: Oxford University Press.

Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, *107*, 797–817.

Bikhchandani, S., Hirshleifer, D. and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100*, 992–1026.

Bush, R. R. and Mosteller, F. (1955). *Stochastic Models for Learning*, New York: Wiley.

Cross, J. G. (1983). *A Theory of Adaptive Economic Behavior*, Cambridge: Cambridge University Press.

Dosi, G., Ermoliev, Y. and Kaniovski, Y. (1994). Generalized urn schemes and technological dynamics. *Journal of Mathematical Economics*, *23*, 1–19.

Ellison, G. and Fudenberg, D. (1993). Rules of thumb for social learning. *Journal of Political Economy*, *101*, 612–643.

Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: markets as a partial substitute for individual rationality. *Journal of Political Economy*, *101*, 119–137.

Hayek, F. A. (1948). “Economics and Knowledge (reprint),” in *Individualism and Economic Order*, F. A. Hayek (ed.), Chicago: University of Chicago Press, pp. 33–56.

Hill, B. M., Lane, D. A. and Sudderth, W. D. (1980). A strong law for some generalized urn processes. *Annals of Probability*, *8*, 214–226.

Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*, Ann Arbor: University of Michigan Press.

Holland, J. H., Holyoak, K. J., Nisbett, R. E. and Thagard, P. R. (1986). *Induction: Processes of Inference, Learning, and Discovery*, Cambridge, MA: MIT Press.

Hull, C. L. (1943). *Principles of Behavior*, New York: Appleton-Century-Crofts.

Katz, M. and Shapiro, C. (1985). Network externalities, competition and compatibility. *American Economic Review*, *75*, 424–440.

Katz, M. and Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy*, *94*, 822–841.

Kirman, A. P. and Vriend, N. J. (1995). *Evolving Market Structure: A Model of Price Dispersion and Loyalty* (mimeo). Paper presented at the Workshop Economic ALife, Santa Fe Institute.

Kirman, A. P. and Vriend, N. J. (2001). Evolving market structure: an ACE model of price dispersion and loyalty. *Journal of Economic Dynamics and Control*, *25*(3/4), 459–502.

Lane, D. and Vescovini, R. (1996). Decision rules and market share: aggregation in an information contagion model. *Industrial and Corporate Change*, *5*, 127–146.

Minsky, M. L. (1967). *Computation: Finite and Infinite Machines*, Englewood Cliffs, NJ: Prentice-Hall.

Narduzzo, A. and Warglien, M. (1996). Learning from the experience of others: an experiment on information contagion. *Industrial and Corporate Change*, *5*, 113–126.

Roth, A. E. and Erev, I. (1995). Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, *8*, 164–212.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*, Cambridge, MA: MIT Press.

Vriend, N. J. (2002). Was Hayek an ace? *Southern Economic Journal*, *68*(4), 811–840.