What's Cool? - Modelling Fashion-like Collective Behavior Emergence from Individual Neuro-psychological Conditioning

Jorge Simão[1] and Peter M. Todd[2] and Luís Moniz Pereira[3]
[1,3]Centro de Inteligência Artificial – CENTRIA
Faculdade de Ciências e Tecnologia – Universidade Nova de Lisboa
2829 - 516 Caparica, Portugal
[2] Center for Adaptive Behavior and Cognition
Max Planck Institute for Human Development
Lentzeallee 94 – 14195 Berlin, Germany
E-mail:[1] **jsimao@di.fct.unl.pt**; [2] **ptodd@mpib-berlin.mpg.de**; [3] **lmp@di.fct.unl.pt**

May 3, 2003

KEYWORDS: Fashion, conditioning, Neuro-psychology, Human social behavior, preference dynamics, culture change, agent-based simulation

Abstract

We present an agent-based model of the emergence of fashion like collective behavior, based on a simple abstraction of individuals neuro-psychological conditioning. Our results show that fashion like collective behavior can emerge from social interaction and the working of individuals' neuro-psychological mechanisms, within a wide range of plausible assumptions about the levels of social stratification within a population.

1 Introduction

What do miniskirts, afro haircuts, and body tattoos have in common? They are all forms of body accessories that have had a characteristic fashion like career. They emerge out of obscurity and spread through a population very fast, only to shortly after they reached their maximum popularity to vanish again from the cultural landscape, sometimes to resurgence again long after. For example, in the USA, turtleneck shirts where first popularized in the 1920s, enjoyed a resurgence in the 1960s, and became popular again in the 1990s (Long, 2002).

Although contemporary history of popular culture is riddled which such fashion-like phenomena, current scientific proposals to explain such forms of collective behavior are in many way unsatisfactory or limited in scope. One classical explanation was proposed by Georg Simmel's (Simmel, 1957). According to Simmel, humans are endowed with two opposing drives or instincts: on the one hand they tend to *imitate* those individuals they admire, and on another they tend to *distinguish* themselves from those to whom they are indifferent or who they despise (Benvenuto, 2000). From this it follows that, in a stratified society, characteristic traits of upper social classes should spread down to lower classes by imitation. As this occurs, individuals in upper-classes will want to distinguish themselves from those in lower social classes. This makes them drop previously desired traits and replace them by others.

While Simmel's effect is logically sufficient to make fashion-like collective behavior to emerge (see (Pedone & Conte, 2000, 2001), for a computer simulation), its assumptions present some problems. First, postulating the existence of *instincts* raises an important ontological question (Bateson, 2000): What do we mean, exactly, by instinct? Can't it be used to label any behavioral regularity or pattern? If so, we should then further identify the neuro-psychological mechanisms that might generate the behavior. For example, while one can rationalize the adaptive advantages of imitating somebody else (Boyd & Richerson, 1985), the advantages involved in distinguish oneself are much less clear. Second, how do the imitation and distinguishing behavior relate to the set of values and norms shared by a population? Could the upper-classes embrace any characteristic trait, or are these individuals also somewhat constrained by what is acceptable? While, as observers, one can retrospectively find some cultural traits strange or even bizarre (given our own cultural background), it seems that even upper class individuals have been constrained in their choices by their peers and members of their social network (e.g. by virtue of negative social reinforcement that would discourage too extravagant a behavior (Bandura, 1985)). Moreover, lower class individuals would certainty not imitate everything the upper-classes might envision. Clearly, some cultural traits are more likely to be found more offen than others. That is, the existing values and norms of individuals in the lower classes also constrain what they imitate and what spreads through the population. Finally, how do individuals learn what cultural specific traits are characteristic, and how can they identify others has being members of higher class? Without having an appropriate theory of the way in which preferences and meaning are associated with symbols and traits, one ends up with a circular explanation for the emergence of fashion.

Economists too have hurdled with the problem of explaining the strange career of products that have their value change by fashion-like behavior (Bergman, 2002; Banerjee, 1992; Bikhchandani, Hirshleifer, & Welch, 1992). According to standard micro-economic theory, the demand for a product should mostly depend on the supply and an individual's own evaluation of the product (Frank, 2003). In general, however, such theoretical framework is insufficient to explain the rapid spread of product consumption and the rapid change in individual preference characteristic of fashions. Thus, to explain this kind of phenomena, additional factors (often designated as *externalities* (Schelling, 1978)) are required. One such proposal is that the value of some products is proportional to the total number of its users (e.g. as in telephone or other communication systems). A high number generates a positive feedback loop that leads to rapid spread of products once an initial slow inception phase is surpassed (Katz & Shapiro, 1985). A related proposal, employed to explain the introduction and spread of innovation in markets, is to assume that the likelihood that individuals buy a product is dependent not only on its price (and therefore, its intrinsic value), but also on the number of individuals that have bought it already (Holanda, Bazzan, Gerolamo, H. A. Franco, & Martins, 2003). These explanations are however limited, because they do not explain why many products suddenly lose their value. As put by Margo Bergman: "[...] Most products are still useful to people even if they are the only ones who own them. [...], miniskirts are still clothing, and therefore useful, even if they aren't popular". To fix this, this author as proposed a model where in addition to considering imitation effects in product acquisition, she considers that an individual's interest in a valued product decays with time (Bergman, 2002). While this generates a fashion-like characteristic behavior, this is no more than an ad-hoc fix in the basic model assumptions, as it does not explain why individuals' valuation decays in the first place, neither why it could increase later in time.

An alternative type of hypothesis, proposed by some economists, is based on the assumption that individuals often imitate others blindly, possibly ignoring information available to them. This can create *information cascades* (Banerjee, 1992; Bikhchandani et al., 1992), that lead to the spreading of bad products if the initial choice is poor or ill-advised. Thus, when the product's "real" value is later discovered and the attending information spreads, this creates a sudden drop in product use. The latter kind of explanation presents two shortcomings though. First, similarly to the previous explanations, it does not apply to cases where the product is still functional and its "real" value is high. Most importantly, is assumes that there is such thing as a "real" value that can be independently assessed by the isolated individual (at least given enough information, resources, and time), without really explaining where this value comes from.

The above discussion suggests that one possible avenue to explain the emergence of fashion-like collective behavior is to identify a plausible (neuro-psychological) mechanism that would explain how individuals attribute value to traits or products, and how they might change over time. With this mechanism one could then model a population of agents with some such psychology, and examine what kind of macro patterns it would generate. Fortunately, we know of (at least) one such mechanism: classical (or Pavlovian) conditioning (LeDoux, 1998).

In classical conditioning, previously neutral stimuli (here traits) are made to trigger emotional or behavioral responses, by recurrent association in space and/or time between them and other stimuli that, taken alone, already produced such responses (LeDoux, 1998). A neutral stimulus is said become *conditioned* by other (unconditioned) stimuli. Although this kind of neuro-psychological process is often setup in controlled laboratory environments for experimental scientific studies, one can find evidence that the same or similar processes are also operating in more natural settings. For example, studies have shown that when physically attractive people are rated in other attributes, such as linguistic or speech skills, they receive from observers higher scores than below average attractiveness of individuals (Rosenthal & Simmerman, 1978; Kunda, 1999). And an all too well known example of "application" of conditioning mechanisms, is the marketing strategy of associating value charged stimuli (e.g. sex or status symbols) with the products or brands producers want to sell (Quart, 2003). In fact, it is reasonably well understood that learning to trigger similar responses when presented with similar or spatiotemporal correlated stimuli is a basic organizational principle of cognitive systems — both in artificial AI systems (Holland, Holyoak, Nisbett, & Thagard, 1986), and in natural brains (Arbib, Érdi, & Szentágothai, 1998).

In this paper, we present an agent-based model where a population of agents observe others' traits, and change their preferences for those traits by a rule akin to that of classic conditioning. This allows us to study the macro-level dynamics of preferences change and trait usage in the population. Our results show that fashion-like collective behavior can emerge out of this individual neuropsychological mechanism, within a wide range of plausible assumptions about the level of social stratification within a population. While agents themselves change their preferences in a very natural and reasonable way — making the plausible assumption that conditioning is in general of adaptive value (Miller & Todd, 1990; Todd & Miller, 1991) — at the population level the pattern of change is found to be quite capricious and seemingly irrational.

The rest of this paper is organized as follows: In section 2, we describe our model design. In section 3, we present and discuss the model's results. And in section 4, we set forth our conclusions.

2 A Model of Fashion Emergence by Conditioning

We assume a population of P agents, with two attributes and two preferences. One attribute is real-valued and remains fixed during the entire simulation runs. It represents some quality or cluster of qualities of the agent that is positively valued by all members of the population. For example, physical attractiveness, speech ability, or some aggregated (holistic) perception of several of such traits[1]. We dub this attribute the *quality attribute* or *quality* for short. Its value q_i is assumed to be randomly generated from a normal distribution with mean 0 and standard deviation σ^2. A second attribute t_i is a binary trait, and represents some agent trait or item that the agent either carries with him (when $t_i = 1$), or not (when $t_i = 0$). For example, t_i may represent a type or brand of clothes, or a body ornament. This second attribute can be changed by the agent at will, according to what is perceived by the agent as being the most valued option (from its subjective aesthetical point of view). To distinguish it from the quality attribute, we will call it the *trait attribute* or *trait* for short. Two additional preference attributes represent the agent's subjective perception of what is the value of not having the trait v_i^0, or of having the trait v_i^1. The values of v_i^0 and v_i^1 are updated according to a process of conditioning (or association) between observed traits and qualities, described next. Overall, an agent a_i can be represented formally as a four element vector: $< q_i, t_i, v_i^0, v_i^1 >$.

At each time step, every agent observes a set of N different agent role models. This is the set of agents which influence an agent's current perception of the value of having or not the *trait*. As a very simple abstraction of the process of conditioning, we assume that the value (v^1) of having the trait is changed by an amount proportional to the average of the qualities of the role models that have the trait. Additionally, since we want to explore the effects of memory within agents, we let an agent's new valuation to be influenced by the previous valuation. For this purpose, we define a learning rate parameter α, that specifies how insensitive agents are to new observations. That is, it specifies how slow or how fast agents are in forgetting previous valuations and changing to new ones.

[1]On the specific topic of mating preference this as been abstracted by Donald Symons as *mate value* or *mate quality* (Symons, 1979).

Specifically:

$$v_i^1(t) = v_i^1(t-1) \cdot \alpha + \sum_{a_j : a_j \in M_i \wedge t_j = 1} q_j \cdot (1 - \alpha)$$

In the equation above, M_i is the set of models for agent a_i ($\#(M_i) = N$), $\alpha \in [0, 1]$ is the learning parameter, and $v_i^1(t)$ and $v_i^1(t-1)$, are respectively the new and previous valuations of trait usage for agent a_i. In a dual manner, we define the value of not having the trait to be the average of the qualities of the models that do not have the trait, weighted by the learning parameter α. Formally:

$$v_i^0(t) = v_i^0(t-1) \cdot \alpha + \sum_{a_j : a_j \in M_i \wedge t_j = 0} q_j \cdot (1 - \alpha)$$

The set of role models for an agent is selected stochastically at each time step, such that the probability of an agent being a model to another agent is proportional to how close their qualities are. This is intended to model a scenario of a socially stratified society, where agents are more likely to interact (or observe) agents that are closer to them in quality than they are to interact with average quality individuals. We implement this by using a roulette-wheel selection mechanism (Goldberg, 1989), and making the probability of agent q_j to be a model of q_i proportional to: $max\{(3\sigma^2 - |q_i - q_j|)^{-E}, 0.001\}$, where the exponent parameter E specifies the degree of social stratification (assortment) in the population. This setting implies that, for positive values of E, individuals with extremely high and low qualities have little chance of interacting amongst them, while individuals with medium quality have maximum probability of interaction. When $E = 0$, the population is not stratified, and individuals have an equal probability of interacting with other individuals with regard to qualities. Thus, higher E values imply higher correlation coefficients between the quality of agents and their models (numerically illustrated in the next section).

Each time step a simulation run goes through two phases. In the first phase, each agent is assigned a model set of size N, and each agent updates its subjective value of having and not having the trait as described above. In the second phase, each agent a_i decides whether or not to switch its trait attribute value. If $v_i^1 > v_i^0$, the agent will start using the trait by setting $t_i = 1$ (if not set already). Conversely, if $v_i^0 > v_i^1$ the agent will set $t_i = 0$, indicating that it does not have the trait. Additionally, to allow the model to escape stationary situations where all or none of the agents have the trait (see next section), we define a parameter θ that specifies the probability that an agent will choose a random value for t_i in reference to its own valuation (e.g. due to some exogenous factor). A simulation run is said to have T of such two-phase time steps. This two-phase procedure can be anecdotally interpreted as follows: individuals go out every day sporting or not the trait (e.g. wearing a certain type of clothes). During the day they interact and/or observe other people and, based on their observations, they update their mental subjective perception of the value of having or not having the trait (possibly sub-consciously (LeDoux, 1998)). At night, back home, they sleep on the subject. In the morning, with their new sense of aesthetics properly in place, they make a fresh new decision of whether or not to wear the trait during the day. (In this version of the model, we assume that there is no intrinsic cost in using the trait or cost in switching from the situation where the agent starts using or stops using the trait.)

To study the model dynamics and to identify fashion-like behavior, we are interested in looking for situations where trait usage changes from being very popular to being very unpopular or, conversely, from being very unpopular to very popular. Specifically, we want to see how the proportion of individuals carrying the trait changes with time, why and when does it happen, and how often does this occur during a certain time period. For this purpose, we first define p to be the proportion of agents having the trait. That is: $p = \frac{\#\{a_i \in P : t_i = 1\}}{P}$. When referring to a specific time step t, we specify this quantity as p_t. Next, we define $p^{\uparrow}$, to be the fraction of time p is above .8(80%), and $p^{\downarrow}$ to be the fraction of time p is below .2(20%). We measure this simply by counting the number of time steps in which p is above .8, and below .2, normalizing this value with the duration of the time period considered (here, the number of time steps of a simulation run T). Formally: $p^{\uparrow} = \frac{\# p_t \geq .8}{T}$ and $p^{\downarrow} = \frac{\# p_t \leq .2}{T}$. We now define two somewhat arbitrary, but appropriate measures of fashion-like behavior $\mathcal{F}^{+}$and $\mathcal{F}^{*}$, as follows:

$$\mathcal{F}^{+} = p^{\uparrow} + p^{\downarrow}$$

$$\mathcal{F}^{*} = \sqrt{p^{\uparrow} \times p^{\downarrow}}$$

$\mathcal{F}^{+}$measures a value proportional to the time p has *any* of the extreme values (above .8, and below .2), while $\mathcal{F}^{*}$measures a value proportional to the time p takes on extreme values and simultaneously oscillates between extreme values. Thus, if complete trait usage or non usage stabilizes in the population early on then $\mathcal{F}^{+}$will be high, because p is taking extreme values. On the other hand, $\mathcal{F}^{*}$will be low because one of the extreme values dominates. Intuitively, a high $\mathcal{F}^{*}$value characterizes a scenario where trait use (or non-use) becomes fashionable for a while, but then somehow stops being fashionable, with this repeating over and over again. In addition to this, we also want to measure how fast the overall rate of change in trait use is. For this purpose, we define $\mathcal{F}^{'}$to be the mean change in the proportion of trait usage from one time step to the next over the duration of a simulation run. And, to measure particularly abrupt changes in p, we define the measure $\mathcal{F}^{t}$to be the fraction of time p is above or below .8 to .2 in a certain time step, and changes to .2 and .8, respectively, in the immediate time step afterward. Formally:

$$\mathcal{F}^{'} = \frac{\sum_{t \in [1,T]} |p_t - p_{t-1}|}{T}$$

$$\mathcal{F}^{t} = \frac{\#(p_t \geq .8 \wedge p_{t-1} \leq .2) + \#(p_t \leq .2 \wedge p_{t-1} \geq .8)}{T}$$

3 Results

To study the model dynamics, we set the population size parameter $P = 50$, and the model set size $N = 10$. The standard deviation of agents' quality σ^2 is set to 2. We make each simulation run to have $T = 500$ time steps. And then

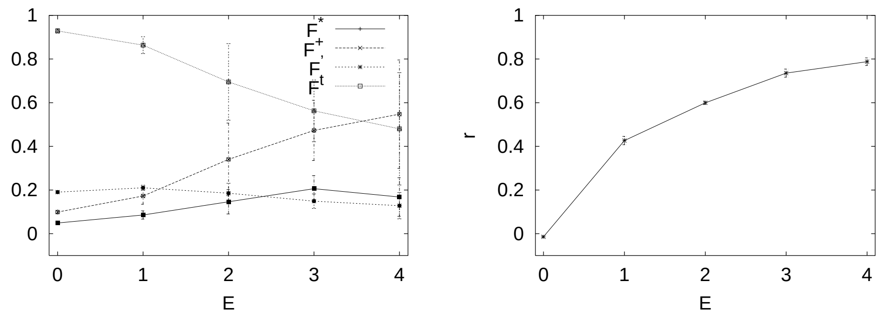

Figure 1: a) Measures of fashion-like behavior as a function of the assortment parameter E, with $\alpha = 0$ and $\theta = 0$; b) Correlation coefficient between quality of agents and their models also a function of E, and the same values for α and θ.

we compute averages and standard deviation of results across 10 such runs. We start all simulation runs with 50% of the agents having the trait ($t_i = 1$), and other 50% not having the trait ($t_i = 0$). And we make the initial value that agents attribute to having or not having the trait equal to 0. This corresponds to a situation where trait usage is initially neutral for all agents, and therefore it's indifferent for them whether or not they have it.

We begin by considering the case where the learning rate parameter $\alpha = 0$, and the random change parameter $\theta = 0$. In figure 1a, we plot the measures of fashion behavior defined in the previous section as a function of the assortment parameter E, and in figure 1b, we plot the value of the average correlation coefficient r between quality of agents and their role models. As expected, the higher the value of E the higher that of r. We choose the interval range to be from $E = 0$, where there is no assortment, to $E = 4$, where r raises to .75. Intermediate values of E render intermediate values of r.

Note that the high values of r considered here are quite plausible for modern societies, given the data collected in a wide range of studies about levels of assortment in individual interaction on many different traits. For example, when considering assortment among individuals of the opposite sex in mating relationships, correlation values between .3 and .8 have been found in attributes such as IQ, physical attractiveness, and others (Kalick & Hamilton, 1986; Marcus W. Feldman, 2000). In previous work, we have shown, using an agent-based model, how such patterns of correlation can arise even when *all* individuals have preference for high quality individuals (Simão & Todd, 2002, 2003). Here, we just assume that the pattern of interaction between individuals leads to assortment without trying to explain how such a pattern emerges. Similarly, it has been found in a series of empirical studies that the same pattern of assortment does occur amongst friends of the same-sex (Grusky, 2000).

From figure 1a, we can see that as E (and r) increases the measures of fashion-like behavior $\mathcal{F}^+$and $\mathcal{F}^*$also increase. On the other hand, the measures of rate of change $\mathcal{F}^t$and $\mathcal{F}^{'}$decrease as E and r increase. To help understand why this occurs, we plot in figure 2a a typical simulation run depicting the proportion p of individuals having the trait ($t_i = 1$), when there is a high degree

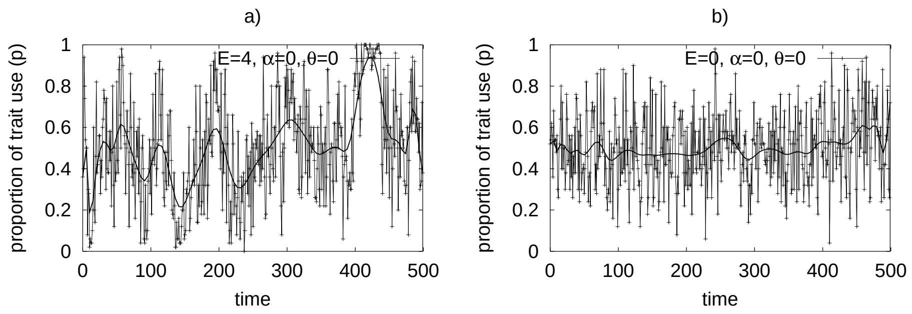

Figure 2: Proportion of trait use in a sample simulation run, with $\alpha = 0$ and $\theta = 0$: a) with $E = 4$ (and $r \approx .74$); b) $E = 0$ (and $r \approx 0$) .

of assortment ($E = 4$ and $r \approx 0.75$). And to contrast, in figure 2b we plot the same proportion p when there is no assortment ($E = 0$ and $r \approx 0$). (To better depict the general trend of change in p, we also plot, super-imposed in the same figure, a smoothed version of the same values using a bezier spline.) As can be seen in figure 2b, in the no assortment case the value of p oscillates around $.5 (= 50\%)$ at a high rate. This occurs because when the trait starts to be common it loses value due to the high proportion of use by low quality individuals, and this prompts individuals to drop the trait. Conversely, when the trait starts to be uncommon it gains in value because not too many low quality individuals are using it, and this prompts more individuals to start using the trait. In general, when the trait is used by more than 50% of the population then the value of not using the trait always has the *potential* to become more valued than that of using the trait; and the reverse occurs when the trait is used by less than 50%. Thus, in a population of omniscient agents, which could observe all other population members at every time step, this would prompt a collective frenzy behavior of continuous preference reversals. While in the no assortment case agents can only observe a fraction $\frac{P}{M}$ of the population at each time step, since there are no structural constraints to the set of members that they are most likely to observe, the role model set tends to be representative of the whole population. This makes agents' computed values v^0 and v^1 not to differ significantly from those that would be perceived by a omniscient agent, thus making collective behavior also to approximate the hypothetical case of a population with fully omniscient agents. This makes $\mathcal{F}^t$and $\mathcal{F}^{'}$ take large values. On the other hand, although in the no assortment case p changes fast it never stabilizes for long within the extreme values $p \geq .8$ and $p \leq .2$. This happens because as soon as p takes extreme value it is pulled almost invariably in the opposite direction. Thus, near full trait usage or near full trait avoidance, never becomes a characteristic population trend. This cause the values of $\mathcal{F}^+$and $\mathcal{F}^*$to be small.

In contrast with the above, we glean from figure 2a that when assortment is very high ($E = 4$ and $r \approx .74$), p not only reaches extreme values of high trait usage and avoidance, but it also remains within extreme values for a substantial amount of time — as measured by $\mathcal{F}^+$and $\mathcal{F}^*$. This occurs because at the same time trait usage becomes fashionable among high quality individuals, the structural constraints on interaction created by a high assortment rate makes lower quality individuals to take some time to observe that. This creates a delay in the devaluation of the trait, during which trait use (or avoidance) becomes a distinguishing feature of the "elite" individuals. However, as more and more lower quality individuals start using the trait its devaluation becomes unavoidable, and this makes the opposite fashion to take over eventually. An important point to notice here is that for this first set of parameter settings, $\alpha = 0$ and $\theta = 0$, this effect is not very robust to changes in the degree of assortment. As figure 1 shows, for an intermediate level of assortment, $1 \leq E \leq 3$, $\mathcal{F}^+$and $\mathcal{F}^*$take values increasingly comparable with the no assortment case very fast. This in spite of the fact that for $E = 1$ the assortment level is still about .4 (figure 1b).

In figure 3a, we plot the same measures of fashion behavior, but now for the case where the learning parameter α is set to .5 (with θ still set to 0). This corresponds to a situation where agents have some internal memory of their trait valuation, possibly creating some inertia in their preference reversal. As can be seen from figure 3a, this modification in the value of parameter α causes $\mathcal{F}^+$to increase very fast as soon as some degree of assortment is introduced. Thus the values of $\mathcal{F}^+$for $E = 1$ ($r \approx .4$), are closer to the values of high assortment $E = 4$, than in the no assortment case $E = 0$. On the other hand, the value of $\mathcal{F}^*$remains very small, and virtually non changing for all values of E. To see better why this happens, we plot in figure 4a the value of p across time in an example simulation run when $E = 4$, and in figure 4b for the no assortment case ($E = 0$). In figure 4a we see that although p stabilizes fast in an extreme value, it never reverses its value to the other extreme value. This process of rapid convergence to extreme p values can be explained as follows: First, when high quality individuals establish a certain trait preference, they induce lower quality individuals to match their preference. As before, this makes higher quality individuals devaluate their current trait preference, thereby causing preference reversals. However, due to the inertia in preference change caused by memory, *on average* there are more lower quality individuals taking the preference of higher quality individuals, than higher quality individuals reversing their preference[2]. This has the effect of making it difficult for changes in the value of p to reverse direction, and has the effect of reducing the difference between the mean quality of the individuals having or not having the trait. This implies that when an agent takes a sample set of role models in a succeeding time step, the changes in the preference values v^0 and v^1 become increasingly smaller. This reinforces the processes of unbalancing the number of preference reversals in a particular direction. Over time, the difference v^0 - v^1 becomes so small that individuals always keep their preference.

In contrast to the above scenario, we can check in figure 4b that when there is no assortment($E = 0$), p never stabilizes with an extreme value. This occurs because high quality individuals (tend to) observe a representative sample of the population, making the less common trait preference in a time step always the more valued one in the next, causing continuous oscillations between extreme p values. However, as indicated by the small value of $\mathcal{F}^*$(figure 3a), trait preferences never remain stable in extreme values.

This far, we have seen that fashion-like collective behavior as indicated by

[2] If the standard deviation of the quality distribution is increased this effect is reduced.

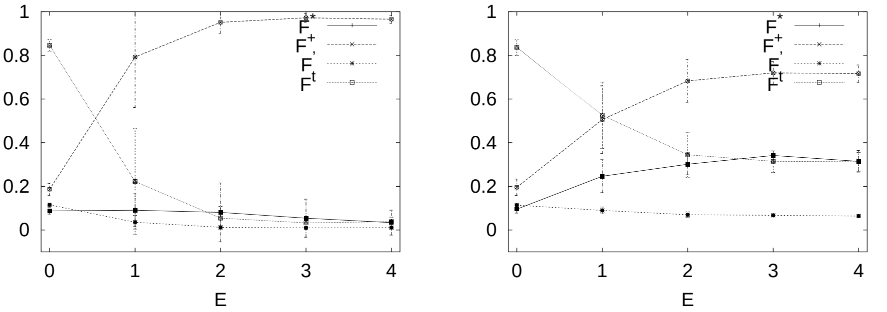

Figure 3: Measures of fashion-like behavior as a function of the assortment parameter E: a) with $\alpha = .5$ and $\theta = 0$; b) with $\alpha = .5$ and $\theta = .01$.

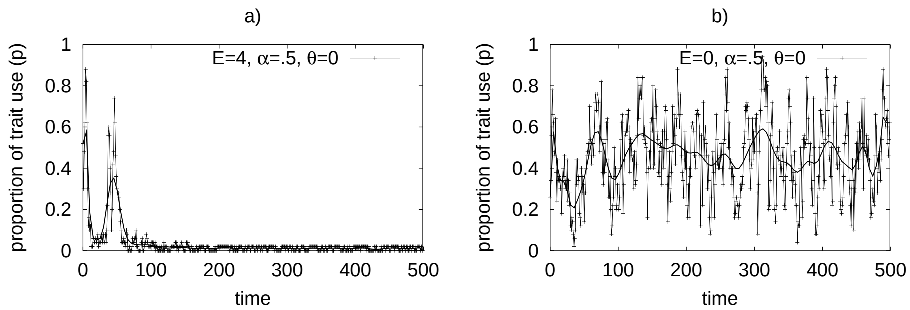

Figure 4: Proportion of trait use in a sample simulation run, with $\alpha = .5$ and $\theta = 0$: a) with $E = 4$ ($r \approx .75$); b) $E = 0$ ($r \approx 0$) .

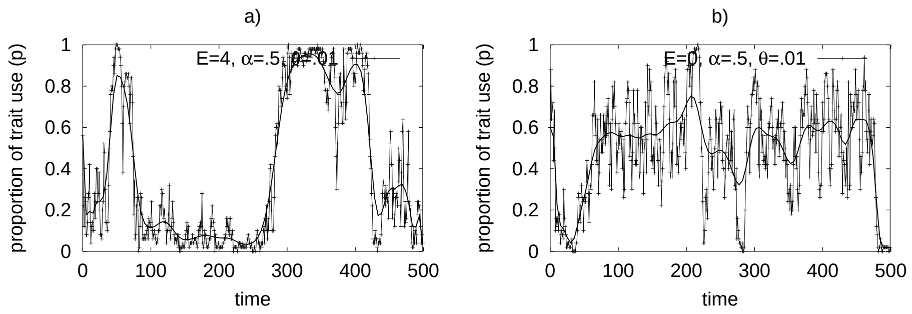

Figure 5: Proportion of trait use in a sample simulation run, with $\alpha = .5$ and $\theta = .01$: a) with $E = 4$ ($r \approx .75$); b) $E = 0$ ($r \approx 0$) .

a high value of $\mathcal{F}^*$, emerges only for a very high level of assortment, and that the introduction of memory causes complete stabilization of trait preferences in some direction (i.e. either full trait use or full trait avoidance). Next, we show that the introduction of a minimal level of randomness in the system, is enough to generate a robust emergence of fashion-like behavior (and consequently of high values of $\mathcal{F}^*$), even for an intermediate level of assortment.

In figure 3b, we plot the different measures of fashion behavior for the case where the learning parameter α is set to .5 as before, but now with θ set to 0.01 (i.e. 1 in every 100 agent decisions are made randomly in what regards preference). With these parameter settings, we see that even for a small value of assortment, $E = 1(r \approx .4)$, a high value of $\mathcal{F}^*$is obtained. To better understand this change in model behavior we depict in figure 5a a typical simulation run depicting the proportion p of individuals having the trait when $E = 4$, and in figure 5b the case when $E = 0$. As can be gathered from figure 5a, and in contrast with figure 5b, while a certain trait preference can fully invade a population for some period of time, it is still possible for the reverse preference to take over the population. This may occur because even if a situation of full trait usage (or non usage) is reached, with a small difference in perceived valuation of traits (i.e. $v^0 \approx v^1$), there is some probability that high quality individuals may change trait usage. This can create an increase in the perceived value of that trait preference, which is conducive to making many agents reverse trait preferences, by following the high quality agent. Sometimes the new trait preference can be devalued very fast due to low quality individuals taking it on, but often enough it occurs that a small cluster of high quality agents induce a complete preference reversal in the population. This makes global population behavior to oscillate between somewhat long periods of near dominance of one trait preference, with periods of near dominance of the other trait preference. In figure 5b, we can see again that, if there is no assortment, oscillations between extreme values still find their way, but trait preference never stabilizes long in extreme values (as indicated by a low value of $\mathcal{F}^*$in figure 3b for $E = 0$).

4 Conclusions

In this article we have presented a simple model that shows how mechanisms of neuro-psychological conditioning at the individual level can generate the emergence of fashion-like collective behavior. The model shows that for even moderate levels of social assortment, trait usage can oscillate continuously between stable periods of near full trait use to near total trait avoidance.

Acknowledgments

This work was partially supported by a PRAXIS XXI, FCT/FSE, Ph.D. scholarship, and project FLUX, funded by FCT/MCES, Portugal. We would like to acknowledge the scientific and intellectual support and/or comments to earlier draft versions of this paper to Paulo Gama Mota, Nuno Preguiça, and João Sousa.

REFERENCES

Arbib, M. A., Érdi, P., & Szentágothai, J. (1998). *Neural organization: Structure, function, and dynamics.* Cambridge, MA: MIT Press/Bradford Books.

Bandura, A. (1985). *Social foundations of thought and action: A social cognitive theory.* N.J.: Prentice-Hall, Inc.

Banerjee, A. V. (1992). Simple model of herd behavior. *The Quarterly Journal of Economics*, *cvii*(3), 797–817.

Bateson, G. (2000). *Steps to an ecology of mind.* Chicago: University of Chicago Press.

Benvenuto, S. (2000). Fashion: Georg simmel. *Journal of Artificial Societies and Social Simulation*, *3*(2). (http://www.soc.surrey.ac.uk/JASSS/3/2/forum/2.html)

Bergman, M. (2002). When a fad ends: An agent-based model of imitative behavior. In *Preproceeding of the workshop on self-organization, evolution and social behavior.* In Press.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as information cascades. *The Quarterly Journal of Economics*, *100*(5), 992–1025.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process.* Chicago: University of Chicago Press.

Frank, R. H. (2003). *Micro-economics and behavior (fifth edition).* New York: McGraw-Hill.

Goldberg, D. E. (1989). *Genetic algorithms: in search, optimization & machine learning.* Addison Wesley.

Grusky, D. B. (Ed.). (2000). *Social stratification: Class, race, and gender in sociological perspective.* Westview Press (2nd edition).

Holanda, G. M., Bazzan, A. L. C., Gerolamo, G. P. B., H. A. Franco, J. ao, & Martins, R. B. (2003). Modeling the bass diffusion process using an agent-based approach. In *Proceeding of the 4th workshop on agent-based simulation* (pp. 147–152). SCS European Publishing House.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: Processes of inference, learning, and discovery.* Cambridge, MA: MIT Press.

Kalick, S. M., & Hamilton, T. E. (1986). The matching hypothesis reexamined. *Journal of Personality and Social Psychology, 51*(4), 673–682.

Katz, M., & Shapiro, C. (1985). Network externalities, competition and compatibility. *American Economic Review, 75*(3), ?–?

Kunda, Z. (1999). *Social cognition: Making sense of people.* Cambridge, MA: MIT Press/Bradford Books.

LeDoux, J. (1998). *The emotional brain: The mysterious underpinnings of emotional life.* Touchstone Books.

Long, M. A. (2002). *Bad fads.* ECW Press.

Marcus W. Feldman, F. B. C., Sarah R. Otto. (2000). Genes, culture, and inequality. In K. Arrow, S. Bowles, & S. Durlauf (Eds.), *Meritocracy and economic inequality* (pp. ?–?). Princeton University Press.

Miller, G. F., & Todd, P. M. (1990). Exploring adaptive agency i: Theory and methods for simulating the evolution of learning. In D. Touretzky, J. Elman, T. Sejnowski, & G. Hinton (Eds.), *Proceedings of the 1990 connectionist models summer school* (pp. 65–80). San Mateo, CA: Morgan Kaufmann.

Pedone, R., & Conte, R. (2000). The simmel effect: Imitation and avoidance in social hierarchies. In S. Moss & P. Davidsson (Eds.), *Proceeding of the 2th workshop on agent-based simulation* (pp. ?–?). Berlin: Springer.

Pedone, R., & Conte, R. (2001). Dynamics of status symbols and social complexity. *Social Science Computer Review: Special Issue on the Simulation of Social Agents, 19*(3), ?–?

Quart, A. (2003). *Branded: The buying and selling of teenagers.* Perseus Publishing.

Rosenthal, T. L., & Simmerman, B. (1978). *Social learning and cognition.* Academic Press.

Schelling, T. C. (1978). *Micromotives and macrobehavior.* New York: W.W. Norton & Company, Inc.

Simão, J. P., & Todd, P. M. (2002). Modeling mate choice in monogamous mating systems with courtship. *Journal of Adaptive Behavior, 10*(2).

Simão, J. P., & Todd, P. M. (2003). Emergent patterns of mate choice in human populations. *Journal of Artificial Life (special issue), (to appear).*

Simmel, G. (1957). Fashion. *American Journal of Sociology*, *62*(6), 541–558.

Symons, D. (1979). *The evolution of human sexuality.* New York: Oxford University Press.

Todd, P. M., & Miller, G. F. (1991). Exploring adaptive agency ii: Simulating the evolution of associative learning. In J.-A. Meyer & S. W. Wilson (Eds.), *From animals to animats: Proceedings of the first international conference on simulation of adaptive behaviour* (pp. 306–315). Cambridge, MA: MIT Press/Bradford Books.