

Social learning in a changing world[★]

Giuseppe Moscarini¹, Marco Ottaviani², and Lones Smith³

¹ Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven, CT 06510, USA

² Department of Economics and ELSE, University College London, Gower St., London WC1E 6BT, UK

³ Department of Economics, Massachusetts Institute of Technology, 50 Memorial Dr., Cambridge, MA 02139, USA

Received: December 5, 1996; revised version: February 11, 1997

Summary. In the social learning model of Banerjee [1] and Bikhchandani, Hirshleifer and Welch [2] individuals take actions sequentially after observing the history of actions taken by the predecessors and an informative private signal. If the state of the world is changing stochastically over time during the learning process, only temporary informational cascades – situations where socially valuable information is wasted – can arise. Furthermore, no cascade ever arises when the environment changes in a sufficiently unpredictable way.

JEL Classification Number: D83.

1 Introduction

Models of learning and experimentation typically predict that information accumulation eventually stops. This prediction contrasts with the observation that information is repeatedly, if not continuously, collected and used. In this paper we study the problem of information aggregation in a model of observational learning by a society of individuals acting sequentially in a changing environment. In this setting, new information is worth more than old information and learning either never stops, or is at least sure to resume after a long enough period of time.

The social learning model of Banerjee [1] and Bikhchandani, Hirshleifer and Welch [2] (BHW) describes the decision problem faced by a sequence of

[★]An earlier version of this paper is contained in the first part of Chapter 3 of Ottaviani's doctoral dissertation at MIT. We thank Massimiliano Amarante, Abhijit Banerjee, Glenn Ellison, Bengt Holmström, Robert Shimer, and Peter Sørensen for helpful suggestions. We also received useful comments from seminar participants at Harvard, MIT, SEDC 1995 annual meeting in Barcelona, and Università Bocconi. The financial support of Banca Nazionale del Lavoro (Moscarini) and Mediocredito Centrale (Ottaviani) is gratefully acknowledged.

exogenously ordered individuals each acting under uncertainty about the state of the world. Everyone conditions his decision in a Bayes-rational fashion on both a privately observed informative signal, and the ordered history of all predecessors' decisions; he observes neither predecessors' realized private signals, nor their realized payoffs. Private signals are assumed to be drawn from an identically and independently distributed random variable correlated with the state of the world. Thus, were all signals publicly observable, their aggregation would eventually reveal the state almost surely, by the Strong Law of Large Numbers. Instead, signals are private, and individuals can only imperfectly infer information from actions.

The striking result in this setting is that with positive probability everyone eventually settles on a common inefficient action, and as BHW note, on a single action regardless. Either outcome is occasioned by the lumpy way in which information is filtered through actions. Suppose that the evidence from the publicly observed action history sufficiently favors one state that it swamps the private signals, as can initially happen by chance. Then private information is perforce ignored, and one's action is dictated by history. Successors then learn nothing from the action, and the system has reached a steady-state where everyone rationally "herds" on the same (either good or bad) action. In this "informational cascade," socially valuable private signals are lost.¹

In this paper we consider a modified but natural model where it is common knowledge that the state of the world changes stochastically over time. We then ask how robust are the above findings to such an evolving environment. We find that because of the resulting information depreciation, only temporary informational cascades can arise. BHW discuss at length the fragility of cascades to the release of small amount of public information. Here, we note that any cascade will eventually come to an end without new informational input, but instead simply because of the fading relevance of old information. Cascades on a single action arise only if the state of the world is sufficiently persistent. When state changes are sufficiently unpredictable no cascade ever arises, because past information depreciates so fast that the belief can never be too extreme. Finally, for completeness more than for economic relevance, cascades on alternating actions arise when the state of the world is changing frequently enough, because here too information depreciates slowly. We conclude that temporary cascades on single actions arise when the environment changes slowly. In this sense, herding survives as a temporary phenomenon only.

This model naturally encompasses many interesting economic situations, like decision making in organizations and consumer choice, where social learning induces inertia in the behavior given a discrete action space. This is the natural extension of the never-ending herd idea of BHW and Banerjee to a stochastic environment. Provided information does not depreciate too quickly, such inertia still arises. BHW themselves briefly discuss the possi-

¹ Smith and Sørensen [7] show that this result holds iff the quality of private signals is bounded.

bility of the state changing in a nonstationary fashion. They provide a specific numeric example where the state changes with 5% chance after 100 periods; they find that cascade reversals are more likely than 5%. In this paper, we feel that a stochastically evolving world is worthy of serious investigation on a par with other recent research on learning.

The problem of optimal experimentation in a changing environment has already been addressed by Kiefer [4]. Recently Keller and Rady [3] have built a tractable continuous-time version of the model. The complexity of the optimization problem precludes a characterization of the long-run learning outcome. In our simple model of social learning the optimization problem is straightforward and the aggregation is simple. See Rustichini and Wolinsky [6] for a manageable framework for long-run analysis in an individual experimentation problem. Intuitively, information depreciation in these contexts increases the immediate value of new information but diminishes its long-term value. In a social learning context, since agents are myopic, only the former effect is present, and the depreciation unambiguously discourages the noninformative cascade stage.

The paper proceeds as follows. Section 2 describes our model of social learning in a changing world. In Section 3 it is shown that only temporary informational cascades can arise, and that cascades arise only if the change in the state of the world is sufficiently predictable. Section 4 concludes.

2 Model

A countable number of individuals take sequentially one of two possible actions, a_0 and a_1 . Payoffs to actions are contingent on an unknown state of the world, ω_0 or ω_1 . Let q^1 be the common prior belief that the state is *initially* ω_1 . Action a_1 is more rewarding than action a_0 in state ω_1 , while the opposite is true in state ω_0 : the payoff of action a_i in state ω_j is 1 if $i = j$ and 0 if $i \neq j$, with $i, j \in \{0, 1\}$. For example, action a_i may be “buy good i ,” and the state of the world ω_i : “good i is better than the alternative good $j \neq i$.”

After an individual’s decision, the state of the world changes with chance ε , assumed for simplicity to be Markovian and independent of the current state of the world:

$$\Pr(\omega^n = \omega_i \mid \omega^{n-1} = \omega_j) = \Pr(\omega^n = \omega_j \mid \omega^{n-1} = \omega_i) = \varepsilon$$

for $i, j \in \{0, 1\}$, $i \neq j$, and any n . Our results can be easily extended to a general two-state Markov transition matrix.

Before choosing an action, individual n both observes a private signal $\sigma^n \in \{\sigma_0, \sigma_1\}$ and the public history of action decisions of all preceding individuals $1, 2, \dots, n-1$. He cannot see predecessors’ signals. Private signals are drawn from a state-dependent Bernoulli distribution, and are independent conditional on the current state. The probability that the signal σ_i is realized in state ω_j is $\alpha > 1/2$ if $i = j$ and $1 - \alpha < 1/2$ if $i \neq j$, with $i, j \in \{0, 1\}$. The quality of the private signal is assumed bounded, i.e. $\alpha < 1$.

For $n \geq 2$ let $H^n \equiv \{a_0, a_1\}^{n-1}$ be the space of all possible period n histories of actions chosen by the $n - 1$ predecessors of individual n . Let h^n denote an element of H^n . Let $q^n \equiv \Pr(\omega_1|h^n)$ be the public probability belief that the state is ω_1 in period n conditional on the publicly observed history of actions chosen by the predecessors of individual n .

Similarly let $r_i^n \equiv \Pr(\omega_1|h^n, \sigma_i)$ be the posterior belief that the state is ω_1 conditional on both the public action history h^n and the realization σ_i of the private signal observed by individual n . A simple application of Bayes' rule yields

$$r_i^n = \frac{\Pr(\omega_1 \cap \sigma_i|h^n)}{\Pr(\sigma_i|h^n)} = \frac{\Pr(\sigma_i|h^n, \omega_1)\Pr(\omega_1|h^n)}{\Pr(\sigma_i|h^n)},$$

so that

$$r_0^n = r_0(q^n) = \frac{(1 - \alpha)q^n}{\alpha(1 - q^n) + (1 - \alpha)q^n} \tag{2.1}$$

$$r_1^n = r_1(q^n) = \frac{\alpha q^n}{\alpha q^n + (1 - \alpha)(1 - q^n)} \tag{2.2}$$

These posterior probabilities are used to compute the expected payoffs from taking the two different actions in the two states. The decision rule of the agent n is to choose the action a^n which gives her the highest expected payoff. Given our simple payoffs, if individual n receives the private signal $\sigma^n = \sigma_1$, then it is optimal to take action $a^n = a_1$ if and only if $r_1^n \geq 1/2$. After substituting from (2.2), this becomes $q^n \geq 1 - \alpha$. The decision rule can be summarized as:

- if $\sigma^n = 0$, then $a^n = a_0 \Leftrightarrow q^n \leq \alpha$, and $a^n = a_1 \Leftrightarrow q^n > \alpha$
- if $\sigma^n = 1$, then $a^n = a_0 \Leftrightarrow q^n < 1 - \alpha$, and $a^n = a_1 \Leftrightarrow q^n \geq 1 - \alpha$

where the action choice when indifferent WLOG minimizes the possibility of herding.

3 Informational cascades with a changing world

A. Belief dynamics

Were ε equal to 0, as in the standard models referenced earlier, the public prior belief of individual $n + 1$ equals the posterior belief that leads Mr. n to act according to her signal $\sigma^n = \sigma_i$, i.e. $q^{n+1} = r_i^n$. There is an *informational cascade* (or *cascade*) on action a_i at time n whenever action a_i is taken by individual n regardless of the individual's private signal σ^n . Thus, a cascade on a_1 (respectively a_0) arises as soon as $q^k > \alpha$ (respectively $q^k < 1 - \alpha$), for then the public belief swamps either private signal. A cascade once started would never end, because public belief would remain unchanged.

With $\varepsilon > 0$, the dynamics change drastically; however, the cascade region is unaffected by the state switching, since that event occurs only after the

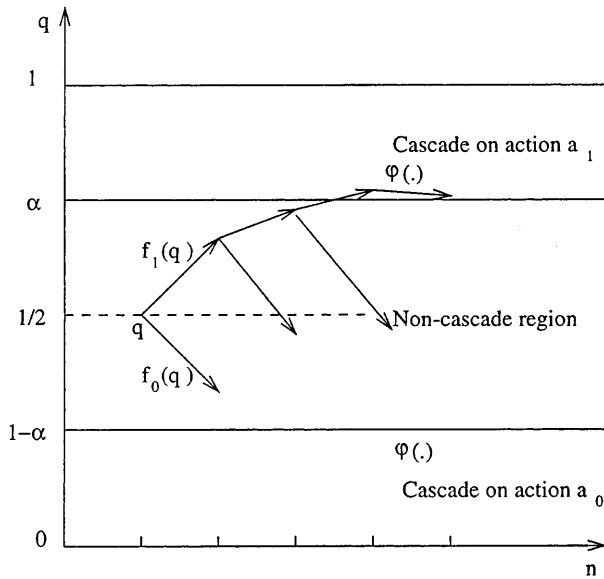


Figure 1

decision is made. When the possibility that the state of the world has changed in the meantime is accounted for, the public prior belief of individual $n + 1$, coming after individual n who chose $a^n = a_i$ according to her signal $\sigma^n = \sigma_i$, satisfies $q^{n+1} = (1 - \varepsilon)r_i^n + \varepsilon(1 - r_i^n)$, which can be rewritten by (2.1) and (2.2) as

$$q^{n+1} = \begin{cases} f_0(q^n) \equiv \frac{(1-\varepsilon)(1-\alpha)q^n + \varepsilon\alpha(1-q^n)}{(1-\alpha)q^n + \alpha(1-q^n)} & \text{if } a^n = a_0 \\ f_1(q^n) \equiv \frac{(1-\varepsilon)\alpha q^n + \varepsilon(1-\alpha)(1-q^n)}{\alpha q^n + (1-\alpha)(1-q^n)} & \text{if } a^n = a_1 \end{cases} \quad (3.1)$$

Note that anyone can compute these probabilities. We will consider the case $q^k > \alpha$, since the other case can be treated symmetrically. The action chosen will be $a^k = a_1$, regardless of the signal σ^k . The next individual $k + 1$ knows that $a^k = a_1$ is uninformative, and computes the public prior belief $q^{k+1} = (1 - \varepsilon)q^k + \varepsilon(1 - q^k)$. In general the following individual $n + 1$, as long as $q^n > \alpha$ or $q^n < 1 - \alpha$, will update her prior belief during the cascade in the same fashion, according to the (uninformative) *cascade dynamics* below:

$$q^{n+1} = \varphi(q^n) \equiv (1 - \varepsilon)q^n + \varepsilon(1 - q^n) \quad (3.2)$$

The public belief dynamics are stochastic and determined by (3.1) as long as $1 - \alpha \leq q^n \leq \alpha$ (when not in a cascade) and are deterministic and follow (3.2) when either $q^n > \alpha$ or $q^n < 1 - \alpha$ (during the cascade). These dynamics are depicted in Figure 1.

B. Cascades are temporary

Now we argue that any cascade will eventually stop.

Proposition 1. *For any $\varepsilon \in (0, 1)$, if a cascade exists, then for some $k = k(\varepsilon) < \infty$, the cascade must end in $k(\varepsilon)$ periods.*

Proof. When in a cascade, the dynamics of the system are described by (3.2). The unique fixed point of this first order linear difference equation is $\bar{q} = 1/2$. Global stability of \bar{q} exactly requires $\left| \frac{dq^{k+1}}{dq^k} \right| = |1 - 2\varepsilon| < 1$, i.e. $0 < \varepsilon < 1$. For $\varepsilon \in (0, 1/2)$ the convergence to the fixed point is monotonic, and oscillatory for $\varepsilon \in (1/2, 1)$. For $\varepsilon = 1/2$ the convergence to \bar{q} is immediate. Since $\alpha > 1/2$ then $[1 - \alpha, \alpha]$ has positive Lebesgue measure and contains $\bar{q} = 1/2$. Therefore there exists $k = k(\varepsilon)$ such that $q^k \in [1 - \alpha, \alpha]$. □

This result is robust to many states. For during a cascade, public beliefs are always drawn back to the fixed point of the cascade dynamics, i.e. the Markov steady state which exists if one can get from one state to any other (formally if the Markov matrix M is regular, i.e. $M^n > 0$ for some n with positive probability), as it is reasonable to assume in a truly changing environment. Cascades are temporary provided that the Markov steady state lies outside the (extremal) cascade set. Finally, the introduction of insurance actions can expand the cascade set (see [7]); since cascades on insurance actions may arise for strictly interior public beliefs, the fixed point of the cascade dynamics may well lie inside the cascade set.

One can also compute the maximum length of a cascade. The longest possible cascade on a_1 starts with a belief $f_1(\alpha)$. After h periods in a cascade the belief is

$$\varphi^h(f_1(\alpha)) = \varepsilon \left[\sum_{i=0}^{h-1} (1 - 2\varepsilon)^i \right] + (1 - 2\varepsilon)^h f_1(\alpha) .$$

This cascade terminates as soon as $\varphi^h(f_1(\alpha)) \leq \alpha$, or equivalently $(1 - 2\varepsilon)^{h+1} \leq [1 - 2\alpha(1 - \alpha)]$, so that $K(\alpha, \varepsilon) \equiv \log[1 - 2\alpha(1 - \alpha)] / \log|1 - 2\varepsilon|$ is a tight upper bound on the length of a cascade. Since $\partial K(\alpha, \varepsilon) / \partial \alpha < 0$ and $\text{sign } \partial K(\alpha, \varepsilon) / \partial \varepsilon = \text{sign}(2\varepsilon - 1)$ for $\varepsilon \neq 1/2$, the lower the quality of private information and the more predictable the evolution in the state of the world, the longer a cascade can possibly last.

C. Cascades arise

To show the possibility of temporary cascades, it suffices to use (3.1) to prove that there exists a fixed point of $f_1(\cdot)$ and one of $f_0(\cdot)$ in the cascade region (i.e. outside $[1 - \alpha, \alpha]$), and that these fixed points are global attractors.

Consider first the case of cascades on a single action. They arise only if the state of the world is sufficiently persistent. If instead state changes are rather unpredictable, i.e. for middling $\varepsilon \approx 1/2$, information depreciates the most. Here, the public belief q^n will never venture far from $1/2$, and so no cascade will ever arise.

Proposition 2 (cascades on a single action). *For any $q^1 \in (1 - \alpha, \alpha)$, with probability one, a cascade on some action arises in finite time if and only if $\varepsilon < \underline{\varepsilon}(\alpha) \equiv \alpha(1 - \alpha)$.*

Proof. It suffices to show that with positive probability, a cascade arises on some action in finitely many periods. Consider first the possibility of a cascade on action a_1 , requiring that $q^k > \alpha$ for some $k > 1$. The most favorable case for the occurrence of such a cascade is when *all* consecutive signal realizations are σ_1 . Then the dynamics follow

$$q^{n+1} = f_1(q^n) = \frac{(1 - \varepsilon)\alpha q^n + \varepsilon(1 - \alpha)(1 - q^n)}{\alpha q^n + (1 - \alpha)(1 - q^n)}. \tag{3.3}$$

We claim that for small enough ε , $\langle q^n \rangle$ converges globally to some attractor strictly above α .

The fixed points of the difference equation (3.3) are the roots of the quadratic equation

$$q^2(2\alpha - 1) + q(1 + \varepsilon - 2\alpha) - \varepsilon(1 - \alpha) = 0. \tag{3.4}$$

Note that $f_1(1) = 1 - \varepsilon$ and $f_1(0) = \varepsilon$, and one can calculate $\text{sgn}f_1'(q^n) = \text{sgn}(1 - 2\varepsilon)$ and $f_1''(q^n)^2 < 0$. For any $\varepsilon \in (0, 1)$ there is only one positive root $\bar{q}(\alpha, \varepsilon)$ of the quadratic equation (3.4). Note that $\bar{q}(\alpha, 1/2) = 1/2$ and $\bar{q}(\alpha, \varepsilon) \in (\varepsilon, 1 - \varepsilon)$ for any α .

To show that the fixed point $\bar{q}(\alpha, \varepsilon)$ is globally stable, let $L(q) \equiv [f_1(q) - \bar{q}](q - \bar{q})$. Then L is a Liapunov function if: (i) $L(\cdot)$ is continuous in q ; (ii) $L(f_1(q)) > L(q) > 0$ for any $q \neq \bar{q}$; (iii) $L(f_1(\bar{q})) = L(\bar{q}) = 0$. It is easy to show that these conditions are always satisfied for $\varepsilon \in (0, 1/2)$. Similarly, for $\varepsilon \in (1/2, 1)$ these conditions are satisfied if $f_1'(\bar{q}) > -1$, which always holds; therefore, the fixed point \bar{q} is globally stable for $\varepsilon \in (0, 1)$.

Finally, there is a positive chance of cascade on a_1 iff ε satisfies

$$\alpha < \bar{q}(\alpha, \varepsilon) = \frac{-(1 + \varepsilon - 2\alpha) + \sqrt{(1 + \varepsilon - 2\alpha)^2 + 4(2\alpha - 1)\varepsilon(1 - \alpha)}}{2(2\alpha - 1)} \tag{3.5}$$

For then, and only then, is it possible that in finite time the public beliefs surpass α , and reach any left neighborhood of this fixed point. Finally, (3.5) reduces to $\varepsilon < \underline{\varepsilon}(\alpha) \equiv \alpha(1 - \alpha)$, as required. The argument for cascades on a_0 is symmetric: If $\underline{q}(\alpha, \varepsilon)$ is the positive fixed point of f_0 , then $\varepsilon < \underline{\varepsilon}(\alpha)$ is equivalent to $\underline{q}(\alpha, \varepsilon) < 1 - \alpha$. \square

Consider a belief on the border of the non-cascade region, e.g. $q = \alpha$. If $\varepsilon = \underline{\varepsilon}(\alpha)$ then $f_1(\alpha) = \alpha$, so that the posterior belief is fixed, but for $\varepsilon < \underline{\varepsilon}(\alpha)$ the state changes less rapidly, so that the cascade region is entered after a signal σ_1 departing from $q = \alpha$. If instead the state changes slightly more rapidly, then $f_1(\alpha) < \alpha$, so that the posterior after a signal σ_1 would be interior to the non-cascade region.

Temporary cascades arise when information depreciates the least, i.e. when the next state is most predictable based on current beliefs. This happens not only when ε is small, but also when it is near 1. When the state of the world changes rapidly enough and individuals alternate between the two actions, the system may enter an *alternating cascade* in which individuals alternate between the two actions regardless of private signals.

Corollary 1 (cascades on alternating actions). *For $q^1 \in (1 - \alpha, \alpha)$, with probability one, cascades on alternating actions arise in finite time if and only if $\varepsilon > \bar{\varepsilon}(\alpha) \equiv 1 - \alpha(1 - \alpha)$.*

The appendicized proof argues that the public belief oscillates until it enters either cascade region (in finite time). Cascade dynamics (3.2) are nonmonotone given $\bar{\varepsilon}(\alpha) > 1/2$.

From Proposition 2 and Corollary 1, and $\alpha > 1/2$ it follows immediately that if the environment is changing in an unpredictable way there will never be a cascade.

Corollary 2. *No cascade ever arises for $\varepsilon \in [1/4, 3/4]$.*

4 Conclusion

Social learning decouples the behavior of the agents from the economy's fundamentals. During an informational cascade on a single action, the same action persists predictably while the environment changes with positive probability. Then the action eventually switches and learning resumes. This simple model of observational learning can explain why common practice can persist more than it should; agents stick to the practice even if they have contrary private information, without possibly knowing in an informational cascade whether others have similar contrary information.

This paper has dwelt on the short-run aspects of learning in a changing world. In a fixed world a never-ending cascade starts eventually with probability one, while in our setting the long-run analysis becomes more interesting. For that we refer to the companion paper Moscarini and Ottaviani [5].

This model is amenable to applications across the economics spectrum, where an invariant state of nature does not ring true. For a salient microeconomics example, suppose that actions are the sequential and observable consumer purchase decisions. Then a changing environment might correspond to the stochastic but unobserved quality innovations of the products purchased.

Appendix: Proof of Corollary 1

Proof. The most favorable case for the occurrence of such a cascade is when all signal realizations exactly alternate up to n . WLOG suppose first that $\sigma^{n-1} = \sigma_0$ and $\sigma^n = \sigma_1$ for n even. Therefore the public belief dynamics are

$q^n = f_0(q^{n-1}) = f_0(f_1(q^{n-2}))$. The function $f_0(f_1(\cdot))$ will be denoted by $g_0(\cdot)$. Substituting from (3.1) one gets

$$g_0(q) = \frac{[\alpha(1-\alpha)(1-2\varepsilon) - \varepsilon(1-\alpha-\varepsilon)]q + (1-\alpha)\varepsilon(1-\varepsilon)}{\varepsilon(2\alpha-1)q + (1-\alpha)(\alpha-\varepsilon+2\alpha\varepsilon)}.$$

Similarly denote $q^{n+1} = f_1(q^n) = f_1(f_0(q^{n-1}))$ by $g_1(q^{n-1})$, where

$$g_1(q) = \frac{[\alpha(1-\alpha)(1-2\varepsilon) - \varepsilon(\alpha-\varepsilon)]q + \alpha\varepsilon(1-\varepsilon)}{-\varepsilon(2\alpha-1)q + \alpha(1-\alpha-\varepsilon+2\alpha\varepsilon)}.$$

For the selected sequence of realizations of signals and for $\varepsilon > 1/2$ the public belief follows oscillating dynamics, while the odd and even subsequences of beliefs follow monotonic dynamics. Following for g_0 and g_1 the same steps as in Proposition 2 for f_1 , we find that both g_0 and g_1 have exactly one fixed point in $[0, 1]$, denoted by $q_0(\alpha, \varepsilon)$ and $q_1(\alpha, \varepsilon)$. Both points are global attractors for the two dynamics and are in the cascade region (i.e. $q_0(\alpha, \varepsilon) > \alpha$ and $q_1(\alpha, \varepsilon) < 1 - \alpha$) if and only if $\varepsilon > \bar{\varepsilon}(\alpha)$. \square

References

1. Banerjee, A. V.: A Simple Model of Herd Behavior. *Quarterly Journal of Economics* **107**, 797–817 (1992)
2. Bikhchandani, S., Hirshleifer, D., Welch, I.: A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* **100**, 992–1026 (1992)
3. Keller, G., Rady, S.: Optimal Experimentation in a Changing Environment. The University of Edinburgh Discussion Paper (1995)
4. Kiefer, N.: A Dynamic Model of Optimal Learning with Obsolescence of Information. Cornell University mimeo (1991)
5. Moscarini, G., Ottaviani, M.: Long-Run Outcomes of Social Learning in a Changing World. Yale University and University College London mimeo (1996)
6. Rustichini, A., Wolinsky, A.: Learning about Variable Demand in the Long Run. *Journal of Economic Dynamics and Control* **19**, 1283–1292 (1995)
7. Smith, L., Sørensen, P.: Pathological Outcomes of Observational Learning. Nuffield College Discussion Paper 115 (1996)