A MODEL FOR PREJUDICED LEARNING IN NOISY ENVIRONMENTS

ANDREAS U. SCHMIDT

ABSTRACT. Based on the heuristics that maintaining presumptions can be beneficial in uncertain environments, we propose a set of basic requirements for learning systems to incorporate the concept of prejudice. The simplest, memoryless model of a deterministic learning rule obeying the axioms is constructed, and shown to be equivalent to the logistic map. The system's performance is analysed in an environment in which it is subject to external randomness, weighing learning defectiveness against stability gained. The corresponding random dynamical system with inhomogeneous, additive noise is studied, and shown to exhibit the phenomena of noise induced stability and stochastic bifurcations. The overall results allow for the interpretation that prejudice in uncertain environments can entail a considerable portion of stubbornness as a secondary phenomenon.

1. INTRODUCTION

As almost all terms denoting affects, the term 'prejudice' is as ubiquitous as ill-defined, and resists naïve attempts to provide it with meaning in any more epistemologically rigorous sense. Yet the last century saw, with the advent of game theory as a means of formalisation and modelling, the paving of a scientific access path to such notions, partially as a side effect of the growing interest in the behaviour of intelligent beings (usually called agents) in social environments [1], which in turn was spurred by the mathematisation of economics. With the emergence of the Internet as a mass medium, this kind of research has obtained a new test bed, a source of statistical data, and an independent study object, and has thus gained further impetus [2, 3, 4]. Consequently, hard science is lead to occupy itself with formerly foreign concepts from the psychological and sociological domains. Recent efforts in this direction combine game theory and logic with nonlinear systems theory and stochastics, and some of them are quite bold, as for instance [5]. The present study parallels these lines of thought, and is, to the best of our knowledge, the first one to focus on the concept of prejudice. It is based on the report [6], but is essentially enlarged and refined.

Section 2 provides the theoretical framework on which the subsequent construction of a model for prejudice rests. We set up a simple, general model for an environment in which an agent needs to learn a fluctuating *risk*, and argue that prejudice might be beneficial in such circumstances.

Date: 6th August 2004.

2000 *Mathematics Subject Classification*. 91A26, 34F05.

Key words and phrases. Learning, prejudice, uncertainty, noise, random dynamical system, noise induced stability, stochastic bifurcation.

This research was partially supported by a project within the fifth framework programme's priority 'Information Society Technologies' of the EU. The provision of parallel computing resources by the GMD institute SCAI (now Fraunhofer SCAI) is gratefully acknowledged.

The special instance of a prejudiced learning rule we consider is constructed in Section 3. In particular, we define what we mean by a *prejudiced* deterministic learning rule, in the language of reinforcement learning. We then set out three axioms we regard as heuristically sound for prejudiced learning in environments polluted with noise. We then look for the most simple prejudiced learning rule fulfilling them, a rule which in particular does not require the agents to have memory. In the noiseless case, the ensuing rule turns out to be equivalent to the logistic map. In correspondence to the well known dynamical features of this map we classify the general behaviour of prejudiced learners, and obtain restrictions on their internal parameters. When the input to the prejudiced learning rule is subject to random fluctuations, it still retains the form of the logistic map, but with additively coupled, inhomogeneous noise.

To corroborate the argument that the limitation of rationality presented by prejudiced learning can be beneficial in certain circumstances, a numerical analysis of the performance of a subclass of prejudiced learners in an uncertain environment is carried out in Section 4, for various levels of noise. The result is twofold. On the one hand, the prejudice causes the agents to make a small error in their belief about the true risk, generally overestimating it. On the other hand, prejudiced learning can efficiently stabilise an agent's behaviour, in particular for higher values of the logistic map's single dynamical parameter. Therefore, if both factors are taken into account, prejudiced learning has the potential to be advantageous in noisy environments.

Our special prejudiced learning rule with noise is an example for a *noisy dynamical system*, a class of systems which has attracted a lot of interest from physicists and mathematicians in recent years. Reference [7] is one earlier, seminal work. Noisy, or random dynamical systems show a host of additional phenomenology over ordinary ones, including phenomena that are to be expected from natural systems. Furthermore, they present a combination of nonlinearity and stochastics on which now a whole branch of mathematics thrives, see [8] and its vast bibliography. The system defined by the prejudiced learning rule is therefore interesting in its own right, and we devote Section 5 to its study. In it, we rediscover the phenomenon of *stochastic resonance* or *noise induced stability* [9, 10], for a large range of parameter values and noise levels, as well as the phenomenon of *stochastic bifurcations* or *noise induced transitions* [11, 12] as these parameters vary. Both phenomena have been empirically confirmed in a vast variety of natural systems and models thereof, ranging from biophysics [13] and chemistry [14], over financial markets [15] and signal processing [16], to quantum information theory [17], without any claim to completeness. We determine the stability domain of our prejudiced learning system analytically using the Lyapunov exponent, and study the bifurcating transition from that domain, which is evoked by lowering the noise level, using numerics. In particular, we determine the *critical exponent* of this transition, in analogy to concepts of statistical mechanics.

Finally, Section 6 contains a comprising assessment of our model for prejudiced learning. It concludes with some suggestions for further research to be based on this and similar models.

2. Background and Heuristics

To place our work into theoretical context and delineate its scope we briefly recapitulate some background. The framework for the construction of our model for prejudiced learning is that of basic game theory and reinforcement learning [18, 19]. We briefly sketch the necessary background.

The classical, game-theoretical subject of *decision making under uncertainty* considers single- or multi player games of intelligent agents against an environment (nature). Each agent has an *utility function* depending on the action he takes and the (unknown) state of nature. This function describes the payoff of the execution of the corresponding action given that nature is in the corresponding state. The utility function is taken to be the input for the agent's *decision rule* which uniquely determines the action to be taken. In the more realistic case when a probability distribution over the states of nature is known to the agent, one is in the realm of *statistical decision theory*. There, this knowledge is fed into a selection scheme which determines one from a set of decision rules accordingly. For example, the well known *Bayesian decision rule* is the combination of a selection scheme and a decision rule which assigns to each action the average sum of utilities weighted with the known probability distribution, and then chooses the action maximising this value. The knowledge can be *a priori* or learnt by experimentation using some *statistical learning rule*, which can be as simple as taking means over a finite set of experimental outcomes (*Bayesian learning*). When the result of taking a specific action is fed back as the outcome of an experiment into the learning rule and the whole cycle is repeated many times, the agent becomes a *learning automaton*. This is the basic object we consider.

The Bayesian rule is obviously not the only possible, there is a multitude of different learning rules and selection schemes. Numerous statistical learning rules have been considered in pursuit of the optimal one in a given environment, see, e.g., [20, 21], and references therein. In realistic cases it can be sensible to choose selection schemes which differ significantly from the apparently most rational Bayesian rule. Here, we propose a model which can justifiably be said to represent agents who are *prejudiced* by construction and in behaviour. We derive the heuristics for our construction from a key example:

Example. At each time step, the agent takes an action, selected from a finite set A, with a certain payoff whose maximal value we normalise to 1, for simplicity. Choosing action k, there is a — generically small — probability p_k for the occurrence of a damage d_k which diminishes the maximal payoff to $1 - d_k$. Assume that the actual damage is symmetrically distributed with small variance around the mean value $\overline{d}_k \in [0, 1]$. The expected payoff then is the (true empirical) **weight** $\overline{w}_k \equiv 1-\overline{r}_k \equiv 1-\overline{d}_k \cdot p_k$ of action k and $\overline{r}_k$ is its true **risk**. The weights are the quantities rational decision makers would base their decisions on. Approximations for them are learnt by the agents from the frequency and amount of previous damages they actually incurred. These are fed into a selection scheme which in turn determines the decision rule for action selection.

The general feature of this example rendering the application of Bayes' rule less attractive is that the damage rate $1/p_k$ and therefore the learning speed of any statistical inference rule for the set $\{\overline{r}_k\}$ can be very low compared to the frequency with which an action has to be taken. Thus, initial probabilistic fluctuations could result in the (costly) selection of a non-optimal action for many steps. On the other hand, the expected damage d_k can be close to 1, resulting in relatively high risk. Thus, the agent has high interest in using a reliable *a priori* risk estimate, to keep it stable against fluctuations, and still learn as quickly as possible.

Numerous studies in algorithmic learning theory are concerned with optimal learning performance in adverse environments polluted with noise, see [22]

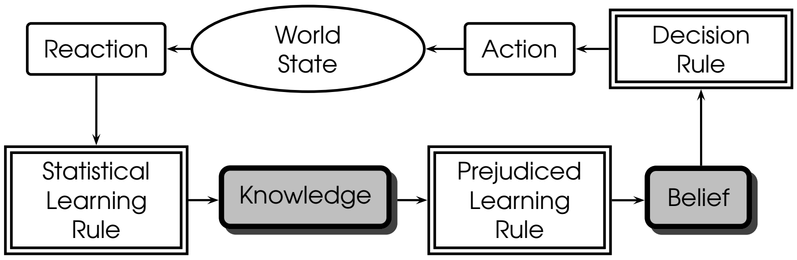

FIGURE 1. Prejudiced learning in the context of decision making under uncertainty. Shaded rectangles are internal states of the agent, unshaded ones stand for actual events. Framed rectangles are algorithms.

and references therein. There, the strategies of the learner, and of its adversaries creating the noise, are rather elaborate. Our present approach differs from these by focusing on constructive simplicity of the model and its behavioural features rather than optimality. Furthermore, our adversary will be a single, very simple, noise model.

3. MODEL BUILDING

Given these conflicting goals above, it can be sensible to use what we would like to call **prejudiced learning system** for weights $\overline{w}_k$, by which we mean the following. The agent is given a start value $w_k(0)$ for its **belief** about the value of $\overline{w}_k$. At time t it infers its (empirical) **knowledge** $\eta_k(t)$ from the $t{+}1$ observations made up to this time using some statistical learning rule not further specified, with the single requirement that $\eta_k(t)$ asymptotically approaches the real value $\overline{w}_k$ as $t \to \infty$. The agents then updates its **belief** w_k about $\overline{w}_k$ at time $t+1$ by **prejudiced learning rule**

$$w_k(t+1) = L\big((w_k(s))_{s \le t}, \eta_k(t)\big), \quad t = 0, 1, \dots$$

The essence of this constructive definition of prejudice is the heuristically sound distinction between knowledge and belief, close in spirit to [23].

Figure 1 shows a schematic view of the internal structure of the prejudiced agent and places it into the general context of decision making under uncertainty. The scope of our model below is merely the transition from knowledge to belief through a learning rule of the above kind, i.e., the three components in the lower right quadrant of the diagram. We will assume in particular that the reaction consists merely in informing the agent of its error, i.e., the numerical difference between its belief and the actual state of the world. Furthermore, since we consider the parameters $\overline{w}_k$ to be independent from and w_k to be unrelated to each other, we focus on learning of a single exterior parameter, and thus drop the index k from now on.

Before proceeding, a caveat might be in order with respect to the terminology we introduced. The Webster's dictionary definition of the term "prejudice" is a "preconceived opinion, usually unfavourable; the holding of such an opinion; an unjustified and unreasonable bias". But then, "prejudiced learning" seems self-contradictory from the outset, and is certainly a controversial expression. We nevertheless stick to this slight provocation, though the bare learning rule above and the three axioms set out below would seem to only justify more cautious terms like "prudence". We shall see that at least one very simple model for "prejudiced learning" will show characteristics, like staying away from the true value of the parameter to be learnt, usually attributed to prejudice proper.

3.1. **Axioms.** The construction of the special model of a prejudiced learning rule we propose is based on three basic axioms and guided by the principle of simplicity. The latter means in particular that we consider only *memoryless* prejudiced learning, i.e., $L = L(w(t), \eta(t))$. Furthermore, we seek to render L in the simplest (functional) form possible. The axioms for the model are derived from the heuristic meaning of three special cases of learning situations, cases which can be viewed as constituting ancillary conditions under which a prejudiced rule has to function. In particular, b) and c) are meant to justify the attribution "learning rule", while a) exhibits an extreme case of "prejudice":

a) **Inability is preserved:** The value $w(0) = 0$ is taken to express the initial inability of the agent to perform the action. This must remain constant, i.e., $w(0) = 0$ implies $L(0, \eta(t)) = 0$.

b) **Importance spurs learning:** The higher an agents ranks an action, i.e., the higher $w(t)$, the faster it shall adapt its belief to its knowledge. That is, $|L(w(t), \eta(t)) - w(t)|$ is monotonously increasing in $w(t)$ for all $\eta(t)$.

c) **Truth is preserved:** If the agent's belief equals its knowledge, then it is kept constant, i.e., $L(\eta(t), \eta(t)) = \eta(t)$.

These axioms imply a fundamental asymmetry between high and low risk, which will reemerge in the behavioural patterns of any model fulfilling them, as will be seen in the particular case below. Heuristically, the axioms incorporate a certain cautiousness, in that they tend to preserve a belief of high risk (low weight). Regarding axiom a), note that we formulated it in the least restrictive way, since it does not exclude a value $w(t) = 0$ at later times. The first axiom also entails a straightforward, implicit assumption on the decision rule, namely that actions with weight zero are not taken.

3.2. **Noiseless Case: The Logistic Learning Map.** The combination of the axioms with the definition of prejudiced learning is certainly satisfied by numerous functional forms of learning rules. We now concentrate on a single special model which, because of its simplicity allows to derive some phenomenological consequences which seem to be of general importance and might pertain also to more complex variants of prejudiced learning rules. It should however be noted that any conclusion drawn from this special model holds properly speaking only for this model and not for general prejudiced learning rules, even if they satisfy the axioms above.

For the construction of a prejudiced learning rule satisfying the above axioms, we specialise to the noiseless case, i.e., we assume the input $\eta(t) = w = 1 - r$ to be identical to the true weight for all times. Expressed in the believed risk $r(t)$ and its **error** $\Delta(t) \stackrel{\text{def}}{=} r - r(t)$, the simplest prejudiced learning rule is linear in either of these variables:

$$\Delta(t+1) = \alpha \cdot r(t) \cdot \Delta(t),$$

with a parameter $\alpha > 0$. Re-expressed in the original variables it reads

$$\begin{aligned} r(t+1) &= r - \alpha \cdot r(t) \cdot (r - r(t)) \quad \text{or} \\ w(t+1) &= (1-\alpha) \cdot w + \alpha \cdot w(t) \cdot (1 - w(t) + w). \end{aligned}$$

This rule reflects the cautiousness inherent in the axioms, in tending to maintain a prejudice of high risk (low weight). It satisfies conditions b) and c) and can be augmented by a simple condition on α to also satisfy a), see Section 3.4 below. Two extremal cases for the parameter α are $\alpha = 0$, in which case the agent immediately learns the input value r, i.e., no prejudice is present, and $\alpha = 1/r(0)$ prohibiting any learning.

Introducing the **relative error** $\delta(t) \equiv \Delta(t)/r$ the learning map reduces to

$$\delta(t+1) = \alpha \cdot \frac{\Delta(t)}{r}\Big(r - \Delta(t)\Big) = \alpha r \cdot \frac{\Delta(t)}{r}\left(1 - \frac{\Delta(t)}{r}\right) = \\ = f_\rho(\delta(t)) \equiv \rho \cdot \delta(t)(1-\delta(t)), \quad \text{with } \rho \equiv \alpha r.$$

This is the well known logistic map [24, Section 7–4], whose characteristics are entirely determined by the bifurcation parameter ρ. It is usually considered as a self-mapping of the unit interval, but note that the natural domain for δ is here $[(r-1)/r, 1]$.

3.3. **Classes of Behaviour.** We assume the intrinsic parameters $(\alpha, r(0))$ to be invariable characteristics of a given agent. Then, in the noiseless case, the behaviour of an agent is determined by these intrinsic parameters, and the environmental parameter r via $\rho = \alpha r$. We tentatively distinguish between three classes of agents, and denote them heuristically as follows.

Adaptive (A): For $\rho < 1$, the logistic map has 0 as a single, attractive fixed point. These agents are therefore bound to adapt to the external input r at an exponential rate.

Stubborn (S): In the domain $1 < \rho < 3$, the fixed point 0 becomes unstable, and the unique stable fixed point is $\delta^* = (\rho - 1)/\rho$, which is again approached at an exponential rate. Agents of this class are bound to persistently underestimate the true risk to a certain degree.

Uncertain (U): For $4 > \rho > 3$, the logistic map exhibits a bifurcating transition to deterministic chaos. Those agents exhibit an increasingly erratic behaviour as ρ rises, which we subsume under the label 'uncertainty'.

For $\rho > 4$, the logistic map is no longer a self-mapping of the unit interval. We simply ignore this case.

We will see in Section 3.4 that behaviours of the last two classes S and U can, under reasonable assumptions, only occur if the agent initially underestimates the risk $r(0) < r$. On the other hand, an initial value $r(0) > r$ will lead to agents of class A, corroborating the heuristics that cautiousness, that is overestimation of the risk, entails a rather safe behaviour.

3.4. **Viability Conditions.** In many senses, the logistic learning map is too simple to work properly. In particular, it does not satisfy axiom a), but this can be accomplished by adapting the intrinsic parameters $(\alpha, r(0))$. The value $r(0) = 1$ implies $\alpha = 1$. This condition, which we impose from now on, is a paradigm for what we call a **viability condition**. These are conditions on the intrinsic parameters that guarantee a proper functioning of the agents or improve their performance in a given environment. A biological heuristics for their prevalence is that agents not fulfilling them are naturally deselected.

We use two other viability conditions. First, we restrict the range of admissible α-values for all agent classes to $\alpha < 1/r(0)$, to make the learning map contractive at least in the first time step, and to omit the extremal case of non-learning $\alpha = 1/r(0)$ mentioned above. Though we could restrict the range of α further to avoid the occurrence of agents with $\rho > 4$, we refrain from posing the pertinent intrinsic condition $\alpha \leq 4$, since this demand seems too restrictive for all classes. As said above, we ignore agents with resulting $\rho > 4$.

The last condition concerns only agents of class A and is somewhat more severe in its consequences: For these agents, the fixed point δ^* is negative and moreover repulsive. If it lies within the range $[1-1/r, 1]$ of admissible δ-values, then the learning map would diverge to $\delta \to -\infty$ whenever $r(t)$ becomes $< \delta^*$. Although this could be avoided in the noiseless case by admitting only start values $r(0) \geq \delta^*$, it would almost surely happen in noisy environments, e.g., the additive noise model of the following subsection. To prevent this disastrous effect we require $\alpha < 1$ for A-agents as the third viability condition. This leads to $\delta^* < 1 - 1/r$, pushing the repulsive fixed point out of the range of δ.

The viability conditions already have direct consequences for S- and I-agents. Due to the condition $\alpha < 1/r(0)$, the case $\rho > 1$ can only occur if $r/r(0) > \rho > 1$, i.e., if the agent initially underestimates the risk. Furthermore, The higher values of ρ, and therefore the more complex behaviour patterns, emerge with increasing discrepancy between initial belief $r(0)$ and knowledge $\eta(t) = r$. We have given a heuristic interpretation of these features in Section 3.3.

Although it would be desirable to dispense with the viability conditions by refining the learning rule, this does not seem easy at the given level of simplicity, without giving away other desired features. For instance, consider the straightforward attempt to adapt the parameter α with time to keep it $< r(t)$, e.g., using the additional rule $\alpha(t+1) = \alpha(t)r(t)/r(t+1)$, and thus force the learning map to be contractive at all times. This would certainly enable us to lift the viability condition $\alpha < 1$ for A-agents. Yet, apart from necessitating at least a one-step memory for the prejudiced learner, it would also let the modified agents become subject to fluctuations in $r(t)$, something the heuristics for the construction of prejudiced learning rules suggests to avoid in the first place. We will propose a measure of the pertinent quality in Section 4, and see that the logistic learning map amended by the viability conditions performs well with respect to it.

As opposed to these linear boundary conditions in parameter space, we will employ boundary conditions in real space in the next subsection, when submitting the system to a noisy environment.

We can now obtain a very coarse picture on the relative proportions in which an observer would expect the three behavioural classes to occur in a population of prejudiced learners. We calculate the *a priori* probabilities for an agent to belong to one of the classes, assuming that $(r, r(0), \rho)$ are uniformly distributed in the domain $(0,1)^2 \times \{0 \leq \rho < r/r(0)\}$, ignoring the last viability condition that affects only A-agents. Integrating over the pertinent ranges and discarding agents with a resulting $\rho > 4$ yields the ratios $4/5$, $8/45$, and $1/45$ for class A, S, and I, respectively. Thus, this naïve estimation renders adaptive behaviour prevalent, while stubborn and uncertain behaviour occur with small but non-negligible probability.

3.5. **Adding Noise.** We now assume that the input $\eta(t)$ of the prejudiced learning rule underlies additive, statistical fluctuations around the true risk value r, i.e.,

$$\eta(t) = r + \Xi(t),$$

with a random variable Ξ. As the simplest possible noise model we choose Ξ to be symmetrically distributed with spread $\Sigma \geq 0$ around 0. That is, Ξ is i.i.d. in $[-\Sigma, \Sigma]$. To keep η in $[0,1]$, this limits the range of admissible values for Σ to $0 \leq \Sigma \leq \min(r, 1-r)$. Using relative coordinates $\delta(t)$ and $\xi(t) \equiv \Xi(t) \cdot \rho/r$, we can separate the fluctuations from the learning map:

$$\begin{aligned}\delta(t+1) = f_{\rho,\xi}(\delta(t)) &= f_\rho(\delta(t)) + \frac{\xi(t)}{\rho} \cdot (\rho(1-\delta(t)) - 1) \\ &= f_\rho(\delta(t)) + \xi(t)(\delta(t) - \delta^*).\end{aligned}$$

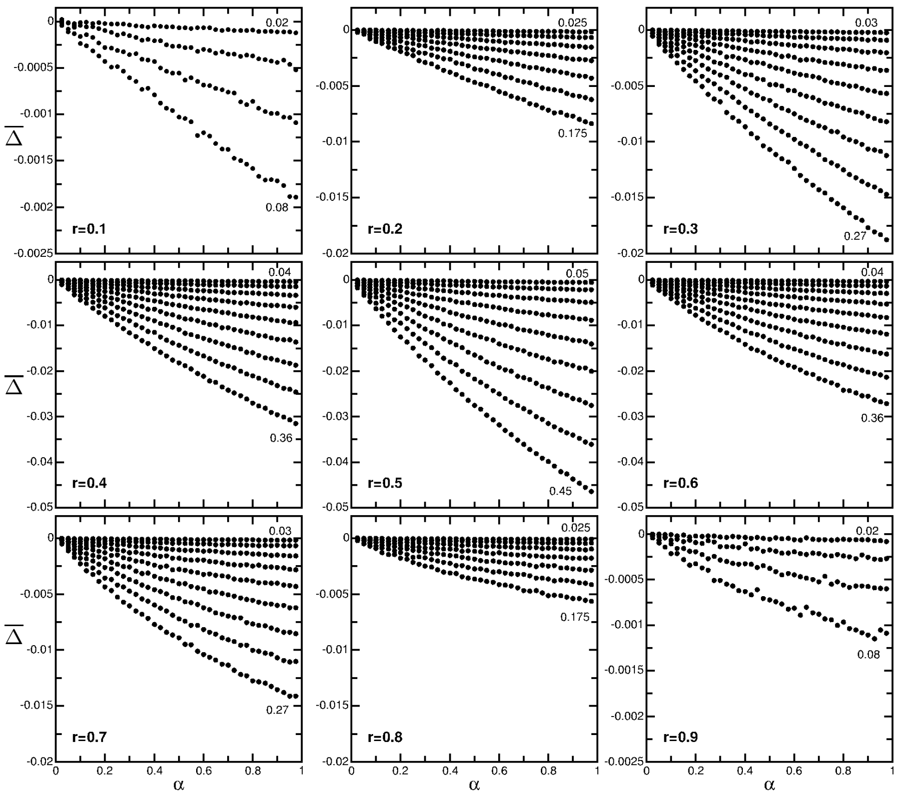

FIGURE 2. Average error of A-agents. Numbers at top/bottom curves denote minimum/maximum values of Σ, between which the remaining curves interpolate in equidistant steps.

This is an example of a *random dynamical system* (RDS). In this special case, it is a *dynamical system with inhomogeneous noise*, for which the inhomogeneity $\delta(t) - \delta^*$ vanishes at the fix point. While systems with homogeneous, additive noise have been intensively studied, the inhomogeneous case is scarce in the literature. The additive separation of the noise from the dynamical mapping means, in particular, that we will still be able to make use of the tentative classification of Section 3.3 for the noiseless case, to classify, at least partially, the behaviour of the agents under noise.

The change of variable from Ξ to ξ renders ξ an i.i.d. random variable with range $[-\sigma/2, \sigma/2]$, where $\sigma = 2\rho/r \cdot \Sigma$. Therefore, the general bounds $\alpha < 1/r(0)$ and $\Sigma \leq \min(r, 1-r)$ imply $\sigma < 2/r(0) \cdot \min(r, 1-r)$ in general, and $\sigma < 2\rho$ in dependence of ρ. On the other hand, the **noise level** σ is limited to $\sigma < 2\min(r, 1-r)$ for A-agents by the viability condition $\alpha < 1$.

Since the noise is in general non-vanishing at both boundaries of the domain $[(r-1)/r, 1]$ of δ, the functions $f_{\rho,\xi}$ are in general not self-mappings of this domain, and thus need to be augmented by boundary conditions. For the performance analysis of A-agents in the next section we use boundary conditions of von Neumann type

$$f_{\rho,\xi}((r-1)/r) = (r-1)/r, \quad f_{\rho,\xi}(1) = 1,$$

for all $\rho < 1$, ξ. That means that when the system hits the boundary, it remains there for one step and has, due to noise, a probability $\geq 1/2$ to leave it in the next step.

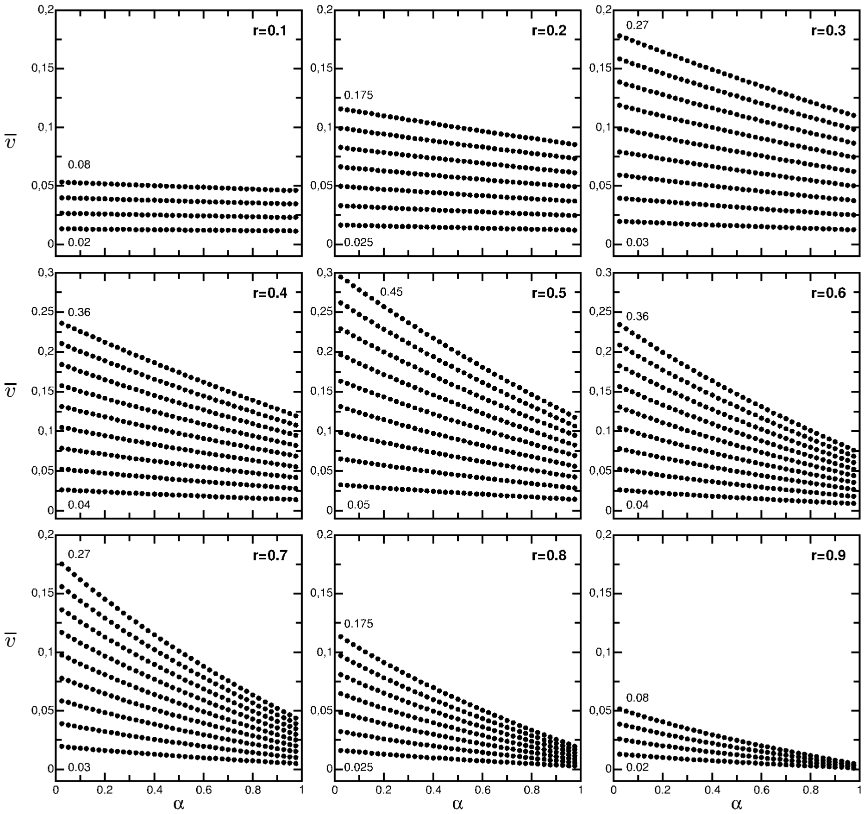

FIGURE 3. Average volatility of A-agents.

4. PERFORMANCE UNDER NOISE

Our prime heuristic for the introduction of the limitation of rationality presented by prejudice into learning was that agents can benefit from not following stochastic fluctuations, in addition to minimising the errors they make. Therefore, the performance of prejudiced agents in a noisy environment should be assessed by considering at least the following *two* natural quantities. The first one is simply the average error $\overline{\Delta}$ which is a measure for the deterioration of learning performance due to prejudice. However, we expect the learners to benefit from prejudice by capitalising on a reduction of the average **volatility** $\overline{v} \stackrel{\text{def}}{=} \overline{|r(t+1) - r(t)|}$ of their belief, a variable which might for example be associated with an energy cost, perhaps arising from an energetic price the agents would have to pay for changing their selected action. We use these two variables to analyse the performance of adaptive agents by numerical simulations.

The benign neglect of S- and I-agents is motivated by their fundamentally different behaviour that hampers a quantitative comparison with A-agents. In particular, stubborn agents have a constant error $\Delta > 0$ and volatility 0. This persists even when they are subjected to noise, as will be seen, together with further qualitative features of the other agent classes, in Section 5.

As a further restriction, following the heuristics that statistical learning is slow in comparison to the frequency of prejudiced learning and decision making, we consider the limiting case in which the fluctuation level Σ is constant, and thus ignore statistical learning altogether. Generically, we would suppose it to exhibit a slow decay, depending on the efficiency of the statistical learning rule and the characteristics of the environment.

Simulations took place in the parameter ranges $0 < r < 1$, $0 < \alpha < 1$, restricted by the viability conditions of Section 3.4, and $0 < \Sigma < \min(r, 1 - r)$, determined by the noise model of Section 3.5. The results for error and volatility are shown in Figures 2 and 3, respectively, in which every data point represents an average over 2.5×10^6 time steps in 25 independent runs with random starting values $r(0)$.

It springs to the eye in Figure 2 that the error induced by noise is always negative, i.e., leads the agents to overestimate the risk. Furthermore, the error is always small, hardly ever reaching 10 percent, and tends to be somewhat smaller above $r = 0.5$ than below. Thus, the heuristics on cautiousness inherent in the construction of the model is confirmed in this experimental setting.

On the other hand, as the last two rows in Figure 3 show, a stabilising mechanism of prejudiced learning becomes effective with increasing α, as $\overline{v}$ decreases from its unadulterated value $2/3\Sigma$ at $\alpha = 0$. This is particularly true for higher r entailing a higher dynamical parameter ρ. In fact, as the graph for $r = 0.9$ exhibits, the learners become completely stable when ρ approaches 1, even in the presence of noise. This remarkable feature will be studied further in Section 5.

Altogether, prejudice in learning has the potential to improve the performance of agents by reducing their volatility significantly, while not putting them far out in their risk estimation, or impeding the efficacy of learning too much. In particular, the learning rate, or rather the rate at which a stable equilibrium is approached, remains exponential.

5. A Noisy Dynamical System

5.1. **Noise Induced Stability of Prejudiced Learning.** In Section 4 we have seen that the prejudiced learning map becomes increasingly stable as ρ approaches 1, despite the presence of noise. In fact, the inhomogeneity of the noise, vanishing at the fixed point ρ^*, leads us to the suspicion that this point plays a special role for the dynamics of the map at higher ρ. In this section we want to pursue this trail further, and consider the features of the map for ρ between 1 and 4, the upper limit at which the noiseless logistic map reaches full deterministic chaos.

Here and in the subsequent numerical analysis we employ periodic boundary conditions, i.e., we consider the random RDS $f_{\rho,\sigma}$ defined by

$$x_{t+1} = f_{\rho,\xi_t}(x_t) \stackrel{\text{def}}{=} \rho x_t(1 - x_t) + \xi_t(x_t - x^*) \mod 1,$$

with ξ an i.i.d. random variable in $[-\sigma/2, \sigma/2]$, and $x^* = (\rho - 1)/\rho$. These boundary conditions make the presentation and the numerics somewhat simpler, while it turns out that they do not affect the behaviour of the system significantly.

Simulations of this system show that, for many combinations of the dynamical parameter ρ and a nonzero noise level σ, it rapidly approaches, and then stably remains at, the fixed point x^*. In fact, this holds for *all* admissible noise levels $\sigma < 2\rho$ in the domain $1 < \rho < 3$ corresponding to agents of class S. Yet the effect prevails even for a continuum of agents of class U, that is, in the parameter range $3 < \rho < 4$, where the noiseless system undergoes Hopf bifurcations into 2^n cycles on the route to deterministic chaos. An example for this *stochastic resonance* is shown in Figure 4b), while Figures 4c) and d) show how bifurcating behaviour is restored through a *stochastic bifurcation*, when the noise level is lowered to leave the domain of the resonance (the stability domain). These two phenomena are examined in the following subsections.

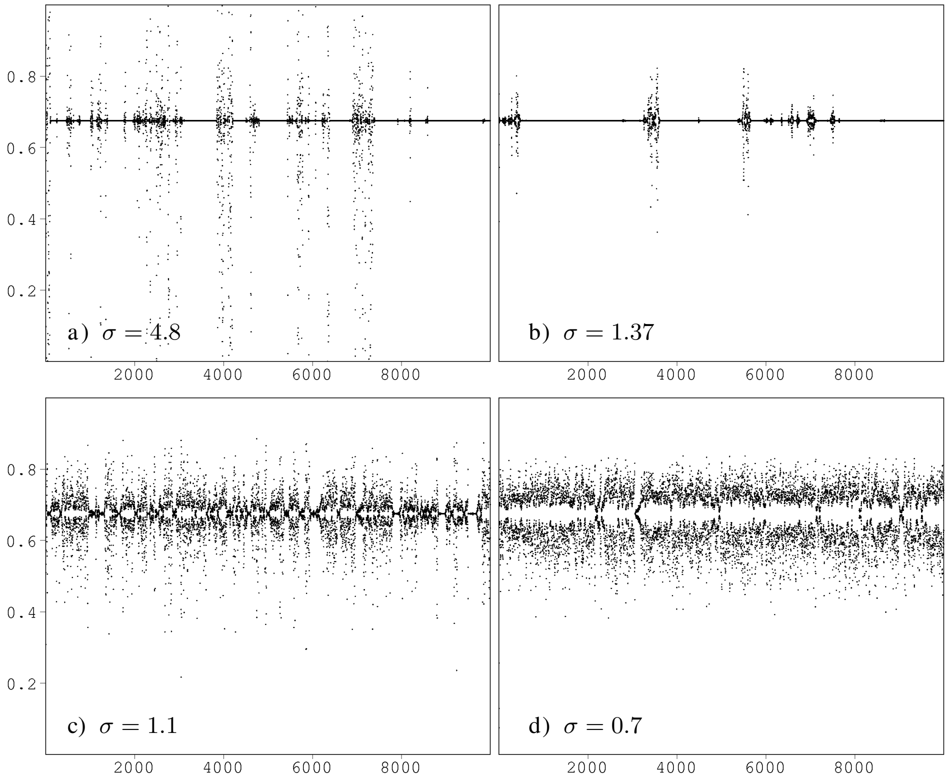

FIGURE 4. Examples for evolutions at $\rho = 3.08$.

5.2. **Basics of Noisy Dynamics.** Let us first introduce the necessary prerequisites on noisy, respectively, random dynamical systems. We rely on [25] as a primary source, and develop the material using the logistic learning map above as an example. Note that more refined theoretical tools for the treatment of such systems abound by now, see for instance [26], but are not needed in the analysis of the simple system at hand.

Every RDS, including the present, is completely characterised by its **transition density** $\mathcal{P}_{\rho,\sigma}(x, y)$, which yields the probability under $f_{\rho,\sigma}$ for ending up in an interval $J \subset [0, 1]$ upon starting in $I \subset [0, 1]$, by the formula

$$\mathrm{Pr}_{\rho,\sigma}(I \to J) = \int_I \int_J \mathcal{P}_{\rho,\sigma}(x, y)\mathrm{d}y\mathrm{d}x.$$

Figure 5 shows an example of this density. Explicitly we find

$$\mathcal{P}_{\rho,\sigma}(x, y) = (\sigma|x - x^*|)^{-1} \cdot \chi_{[x-\frac{\sigma}{2}|x-x^*|,x+\frac{\sigma}{2}|x-x^*|]}(y),$$

where χ is the characteristic function of an interval, and periodic boundary conditions are implicitly assumed on all variables. The transition density serves to define the **Perron–Frobenius (P–F) operator** of an RDS, a central tool in the system's analysis [7]. This operator acts on functions $u \in L^1([0, 1])$, i.e., probability densities by

$$\mathbf{PF}(u)(y) \stackrel{\mathrm{def}}{=} \int_0^1 \mathcal{P}_{\rho,\sigma}(f_\rho(x), y)u(x)\mathrm{d}x.$$

The eigenvalues and eigenvectors of **PF** are of key importance. In particular, the positive and normalised eigenvectors to the highest eigenvalue 1, i.e., probability densities u with $\mathbf{PF}(u) = u$, are called **invariant densities** of the

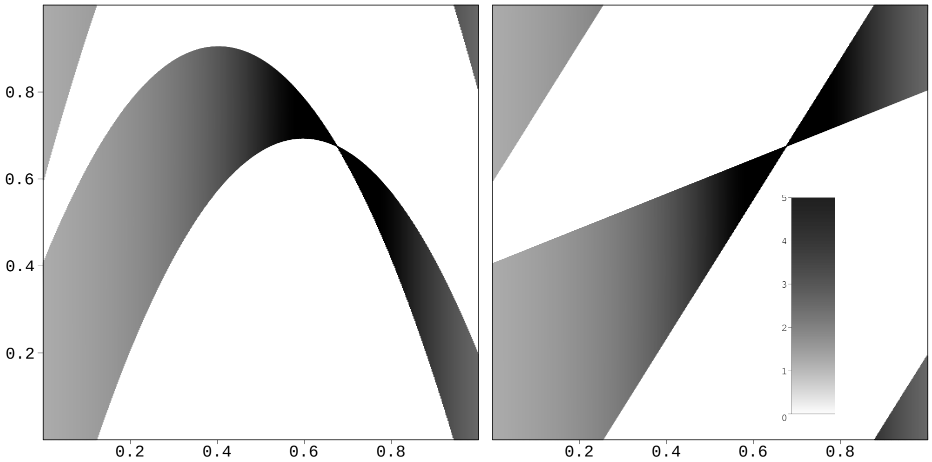

FIGURE 5. $\mathcal{P}_{3.08,1.2}$ (left) and associated noise alone (right).

system. An invariant density u defines an associated **invariant measure** μ_u by $\mu_u(A) \stackrel{\text{def}}{=} \int_A u(x)\mathrm{d}x$, where A is any Lebesgue-measurable set. An important example of an invariant measure for the present system is $\mu_{\delta_{x^*}}$ generated by the the point measure δ_{x^*} at x^*, as an easy calculation shows, using the fact that $\mathcal{P}_{\rho,\sigma}(x_n,\cdot)$ is a δ_{x^*}-sequence if $x_n \to x^*$.

Yet more interest lies in the so-called **physically significant** or **Bowen–Ruelle–Sinai (BRS) measures**. A BRS-measure μ_{BRS} is defined for an ordinary, i.e., non-random dynamical system, defined by a deterministic mapping f, by the following property. There exists a subset U of the configuration space considered and with positive Lebesgue measure, such that for every continuous function ψ, the following holds

$$\lim_{N\to\infty} \frac{1}{N} \sum_{j=0}^{N-1} \Psi(f^j(x)) = \int \Psi \mathrm{d}\mu_{\text{BRS}},$$

for all starting points $x \in U$, where f^j denotes the jth iterate of f, see [27]. By the Birkhoff individual ergodic theorem, this property is always fulfilled for μ_{BRS}-*almost all* x. The crucial strengthening of the hypothesis lies in the assumption that the ergodic hypothesis can be safely applied for *all* starting points x in a set of positive Lebesgue measure. For a RDS, the above property must be formulated in the mean with respect to the stochastic perturbation, i.e., the noise. Since most physically interesting quantities of a dynamical system are time averages, an ergodic hypothesis is regularly invoked by physicists to calculate them by space averages. This makes the existence and uniqueness of BRS-measures an important theoretical issue in the study of random and ordinary dynamical systems.

5.3. **Stable Phase and Lyapunov Exponent.** We are now in a position to bolster the heuristic conjecture of noise induced stability in the parameter range $3 < \rho < 4$ with theoretical and numerical evidence. Furthermore, we want to determine the shape of the latter **stable phase**, i.e., the set of (ρ, σ) values for which the point mass concentrated at x^* generates the unique BRS-measure, that is $\mu_{\text{BRS}} = \mu_{\delta_{x^*}}$.

The **Lyapunov exponent** λ is central to the study of dynamical systems. It is the measure for the exponential rate at which distant starting points converge to an attractor under the system's evolution, respectively, the rate at

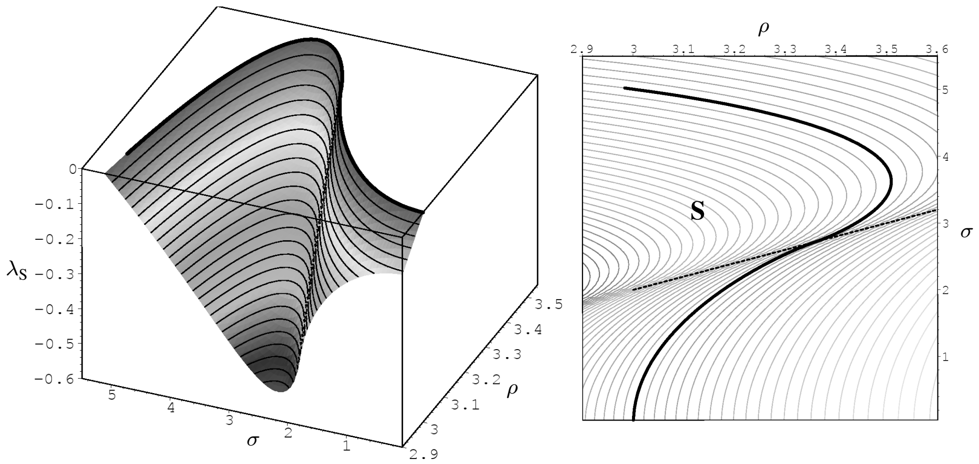

FIGURE 6. Stable Lyapunov exponent λ_S.

which nearby starting points become separated. In the former case, λ is negative, and in the latter positive, while zeros of λ mark transitions in the system's behaviour. Lyapunov exponents can be used to detect stability, bifurcations, and the onset of chaos, and one of the simplest examples for their use is, again, the logistic map [24, Section 7-4]. It has been questioned whether Lyapunov exponents play the same role for RDS, since there are examples in which a positive λ does not indicate ordinary chaotic behaviour [28, 29]. Yet, a *negative* Lyapunov exponent always corresponds to stability, and therefore we chose to tentatively characterise the stable phase of our system through the property $\lambda < 0$. The quantity itself is defined by the time average

$$\lambda \stackrel{\text{def}}{=} \lim_{N\to\infty} \frac{1}{N} \sum_{t=0}^{N-1} \ln |f'_{\rho,\xi_t}(x_t)|.$$

If a BRS-measure is known, λ can be calculated, using an ergodic theorem, as a space average

$$\lambda = \left\langle \int_0^1 \ln |f'_{\rho,\xi}(x)| \mathrm{d}\mu_{\mathrm{BRS}}(x) \right\rangle_\xi ,$$

taking into account that in the noisy case we also have to average over the random variable ξ.

The **stable Lyapunov exponent** λ_S is now calculated under the assumption $\mu_{\mathrm{BRS}} = \mu_{\delta_{x^*}}$. For the given noise model we obtain

$$\begin{aligned}
\lambda_S(\rho,\sigma) &= \frac{1}{\sigma} \int_{-\sigma/2}^{\sigma/2} \int_0^1 \ln |f'_{\rho,\xi}(x)| \mathrm{d}\mu_{\mathrm{BRS}}(x) \mathrm{d}\xi \\
&= \frac{1}{\sigma} \int_{-\sigma/2}^{\sigma/2} \ln |f'_{\rho,\xi}(x^*)| \mathrm{d}\xi \\
&= \frac{1}{\sigma} \int_{-\sigma/2}^{\sigma/2} \ln |\rho(1-2x^*) + \xi| \mathrm{d}\xi \\
&= \frac{1}{\sigma} \int_{-\sigma/2}^{\sigma/2} \ln(\rho - 2 + \xi) \mathrm{d}\xi,
\end{aligned}$$

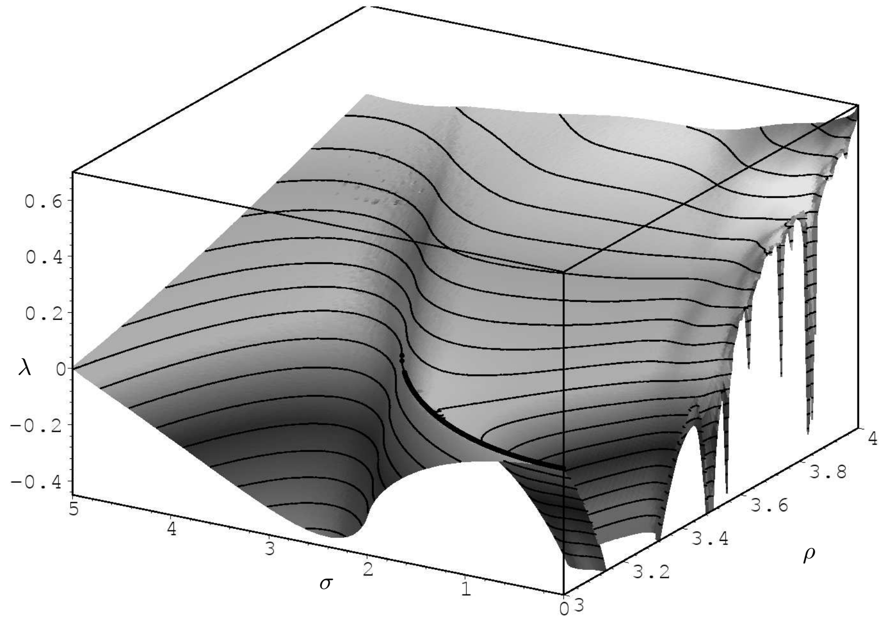

FIGURE 7. Numerical evaluation of $\lambda(\rho,\sigma)$. On a regular grid with resolution $(0.0025, 0.05)$, each point represents 25 independent runs of length 10^6,after omitting 10^4 initial iterations

for $\rho - 2 > \sigma/2$, and by a change of variable $\zeta = \rho - 2 + \xi$ we finally find, using the abbreviations $\Delta_\pm = \rho - 2 \pm \sigma/2$

$$\lambda_{\mathrm{S}} = \int_{\ln\Delta_-}^{\ln\Delta_+} \frac{\zeta \mathrm{e}^\zeta}{\sigma}\mathrm{d}\zeta = \left[\frac{\zeta-1}{\sigma}\mathrm{e}^\zeta\right]_{\ln\Delta_-}^{\ln\Delta_+}.$$

Similar calculations in the two other cases $\rho - 2 < \sigma/2$ and $\rho - 2 = \sigma/2$ yield the net result

$$\lambda_{\mathrm{S}}(\rho,\sigma) = \frac{1}{\sigma}\begin{cases} (\ln\Delta_+ - 1)\ln\Delta_+ - (\ln\Delta_- - 1)\ln\Delta_-, & \text{if } \Delta_- > 0; \\ (\ln\sigma - 1)\ln\sigma, & \text{if } \Delta_- = 0; \\ (\ln\Delta_+ - 1)\ln\Delta_+ + (\ln|\Delta_-| - 1)\ln|\Delta_-|, & \text{if } \Delta_- < 0. \end{cases} \quad (*)$$

Figure 6 displays this result. The solid curve in both pictures is the nodeline $\lambda_{\mathrm{S}} = 0$ at which the hypothesis $\mu_{\mathrm{BRS}} = \mu_{\delta_{x^*}}$ breaks down, while the dotted line in the right picture is $\rho - 2 = \sigma/2$, where the two solutions in $(*)$ connect. The conjectured stability domain is marked by $\mathbf{S}$.

We confirm the existence of a stable phase empirically by collecting numerical data for the Lyapunov exponent using time averages, obtaining the picture shown in Figure 7. We find that λ is identical to λ_{S} within the statistical error bounds for the most part of the region $\mathbf{S}$. Yet for higher $\sigma \gtrsim 3$, the empirical values of λ are significantly larger than λ_S along the inner boundary of $\mathbf{S}$. It seems plausible that the strongly intermittent behaviour of the system in that region, of which Figure 4a) shows an example, together with the periodic boundary conditions, in effect, prevents the convergence of λ to λ_S.

5.4. **Stochastic Bifurcation.** The sample evolutions in Figure 4b)–d) exhibit the stochastic bifurcation the system undergoes when the noise level is lowered to leave the stable phase. We now examine this noise induced transition closer at $\rho = 3.08$. It occurs at the transition point $\sigma_0(\rho)$, which is obtained by numerically solving for the lower solution of $\lambda_S(\rho, \sigma_0(\rho)) = 0$ in equation $(*)$, yielding

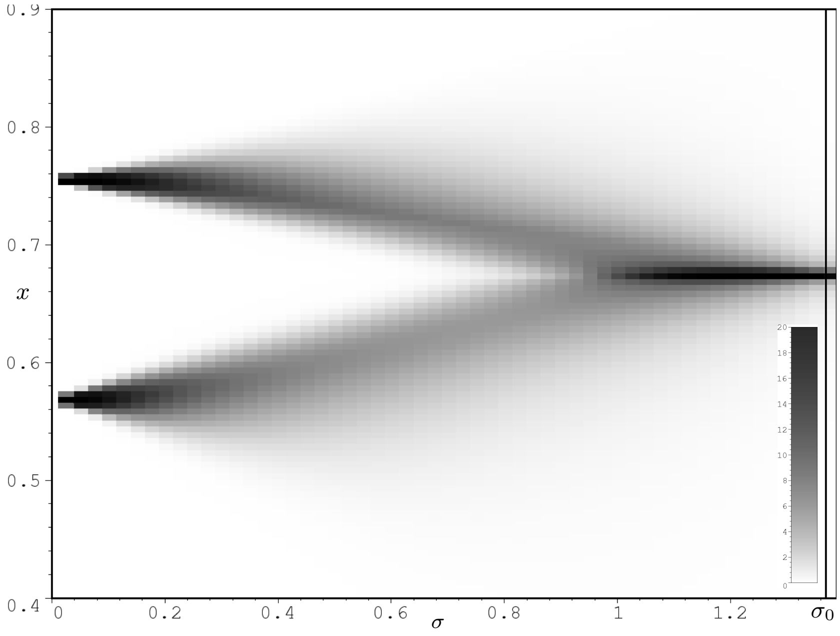

FIGURE 8. Invariant densities at $\rho = 3.08$ on a regular partition of $\{\sigma \in [0, 1.4]\} \times \{x \in [0.4, 0.9]\}$ with resolution $(0.025, 0.0025)$, for $N = 1000$, and $M = 4 \times 10^6$. Cells with absolute probabilities below 10^{-3} appear white.

$\sigma_0(3.08) \approx 1.3683$. As σ approaches 0, the system converges to the deterministic 2-cycle with attractor $\bigl\{\rho_\pm = \frac{1}{2\rho}(\rho + 1 \pm \sqrt{(\rho+1)(\rho-3)})\bigr\}$.

Figure 8 shows the development of the invariant densities for the BRS measures at $\rho = 3.08$ as σ varies. The data for a given σ represents the distribution of M i.i.d. starting points after evolving N time steps. This statistical strategy can be justified if we assume that the Perron–Frobenius–Ruelle theorem [25] holds. It asserts, among others, that the iterates $\mathbf{PF}^N(u)$ of any density u of a normalised measure which is absolutely continuous w.r.t. the BRS measure, converge to the density of the BRS measure as $N \to \infty$. However, just above the transition point (marked by the solid vertical line) the BRS measure $\mu_{\delta_{x^*}}$ is not well approximated with $N = 1000$, cf. Figure 4b). Below σ_0, the singularity of the invariant density disappears, and it develops two maxima around $\sigma \approx 1$, which become separated at $\sigma \approx 0.75$. Overall, the transition is a rather typical example for a stochastic bifurcation [8, Chapter 9].

Recently, it has been conjectured that stochastic bifurcations are in close correspondence to phase transitions in systems of statistical mechanics [30]. Although the theoretical basis for such a claim is presently not firm, the two classes of phenomena exhibit many common features. One of them is *symmetry breaking*, which is also present in the stochastic bifurcations of our system. It can be exhibited by looking at the evolution of the **return map** during the bifurcation. This map is defined by the sequence of the local expansion rates

$$y_t \stackrel{\text{def}}{=} \ln \bigl| f'_{\rho,\xi_t}(x_t) \bigr|,$$

and the statistical distribution of pairs (y_t, y_{t+1}) is its transition density, for which some examples are shown in Figure 9. The symmetry with respect to reflections on the diagonal that is present in the stable phase, is broken by the stochastic bifurcation.

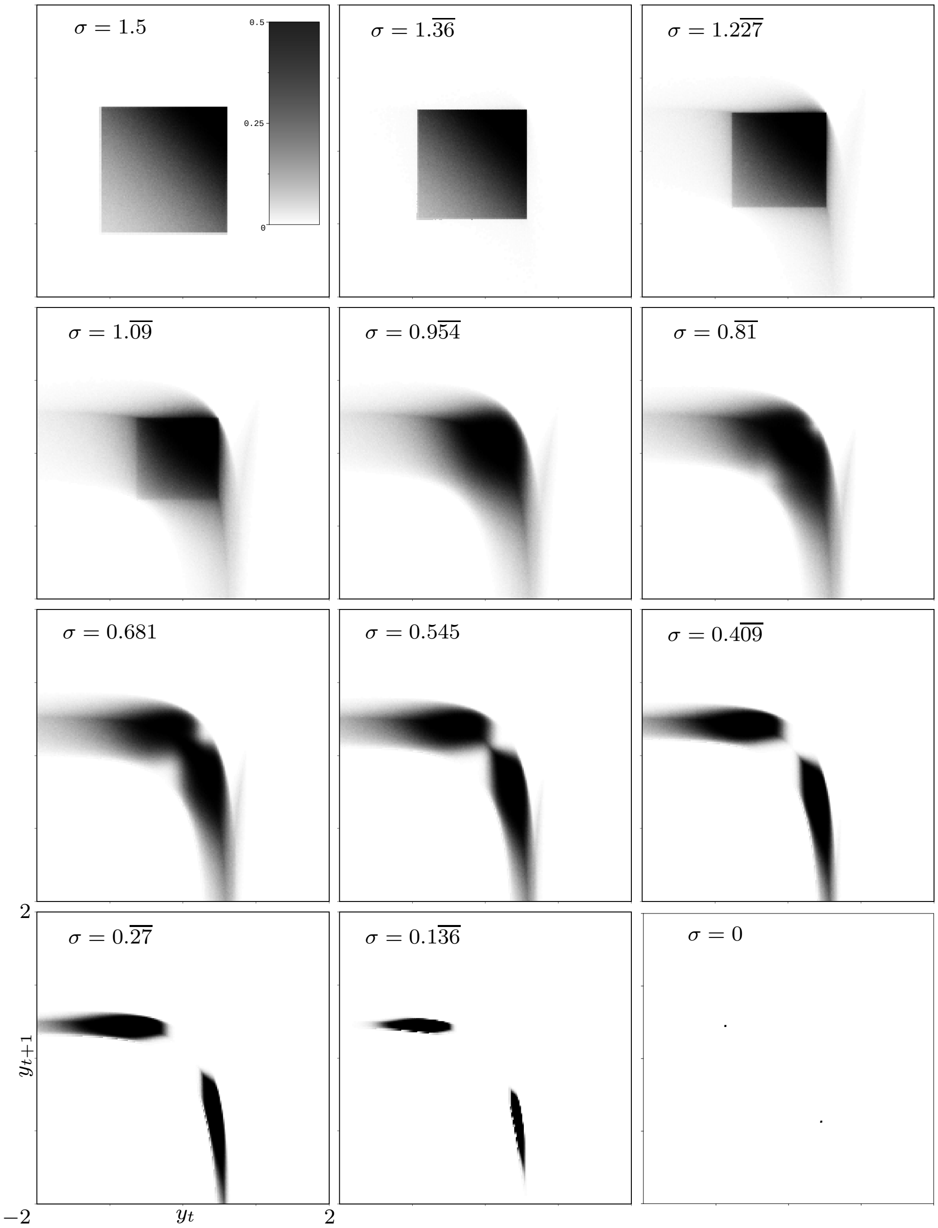

FIGURE 9. Transition densities of the return map at $\rho = 3.08$, approximated on a regular 200×200 partition of $[-2, 2]^2$ using 10×10^6 iterations, omitting 10^4 initial steps. Cells with absolute probabilities below 10^{-6} appear white.

5.5. **Critical Exponent.** Finally, we want to further stress the analogy between stochastic bifurcations and critical phenomena in statistical mechanics [31]. A characteristic of the latter is the vanishing or divergence of certain intensive observable quantities, called *order parameters*, at the critical point. Their values near the critical point are governed by universal scaling laws, which are quantitatively described by the so called *critical exponents*.

As a very simple example demonstrating the general scheme, consider the first bifurcation of the noiseless logistic map at $\rho_c = 3$. The only relevant

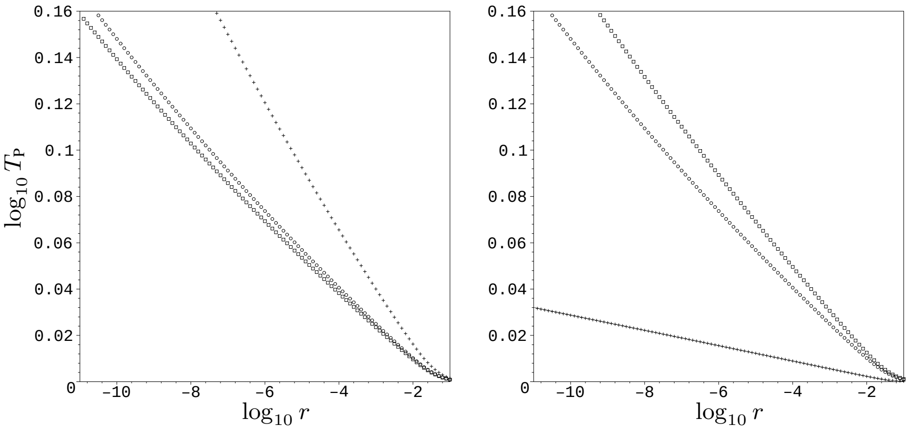

FIGURE 10. Left: Recurrence time statistics at $\tau = 0$; $N = 5 \times 10^3$, $m = 5 \times 10^4$ (+); $N = 10^5$, $m = 5 \times 10^4$ (○); $N = 5 \times 10^6$, $m = 10^3$ (□). Right: Correction for $N = 10^5$, $m = 5 \times 10^4$; $\tau \approx 0.00291$ (□); $\tau = 0$ (○); $\overline{T}_\mathrm{P}$ (+) and linear fit.

order parameter is the Lyapunov exponent, which vanishes at ρ_c with a critical exponent

$$\gamma_\pm \stackrel{\mathrm{def}}{=} \lim_{\tau \to 0_\pm} \frac{\ln \lambda(\tau)}{\ln |\tau|} = 1,$$

determined as the exponent of the leading power in the expansion of λ in the scale-free parameter $\tau \stackrel{\mathrm{def}}{=} (\rho - \rho_c)/\rho_c$.

In our RDS, there are a number of other order parameters apart from λ that can be considered, and we use the *Poincaré dimension* of the critical point x^*, which is a very natural measure for the collapse of the invariant density to δ_{x^*} at the critical point. It is determined by the *recurrence time statistics* as follows, cf. [32]. The **Pioncarè recurrence time**, which is the average time after which the system returns to a small domain $B_r(x^*) \stackrel{\mathrm{def}}{=} \{x \mid |x - x^*| \leq r\}$ of radius r around x^*, is given by

$$T_\mathrm{P}(r) \stackrel{\mathrm{def}}{=} \lim_{N \to \infty} \frac{N}{\#\{x_i \in B_r(x^*),\ 1 \leq i \leq N\}}.$$

Obviously, this is not a scale-free quantity, so to obtain the desired order parameter, one assumes that T_P behaves asymptotically as

$$T_\mathrm{P}(r) \propto r^{-D_\mathrm{P}} \quad (r \to 0).$$

The number D_P so defined is the **Poincaré dimension** of x^*. As the system enters the stable phase, with scale-free parameter $\tau \stackrel{\mathrm{def}}{=} (\sigma_0(\rho) - \sigma)/\sigma_0(\rho)$ tending to 0 from above, $T_\mathrm{P}(r)$ approaches 1 for all r, and we expect D_P to vanish.

In the following, we examine the transition at $\rho = 3.08$. Figure 10 shows various examples for recurrence time statistics gathered by numerical simulations with m independent runs of length N, always omitting 10^4 initial steps. In the left picture we see three graphs for T_P at the critical point $\tau = 0$ for various N. This exhibits the problem that, due to the combined effects of slow convergence of the system as such, and the additional numerical error which becomes relevant near the critical point, the measured Poincaré dimension does not seem to converge to zero as $N \longrightarrow \infty$ (and does not depend on the sample size for $m > 10^4$). To circumvent this difficulty we tentatively replace T_P at given τ and N with 'corrected' values $\overline{T}_\mathrm{P}$ gained by division by T_P at $\tau = 0$ and the same sampling time N, as shown in the right picture. A comparison

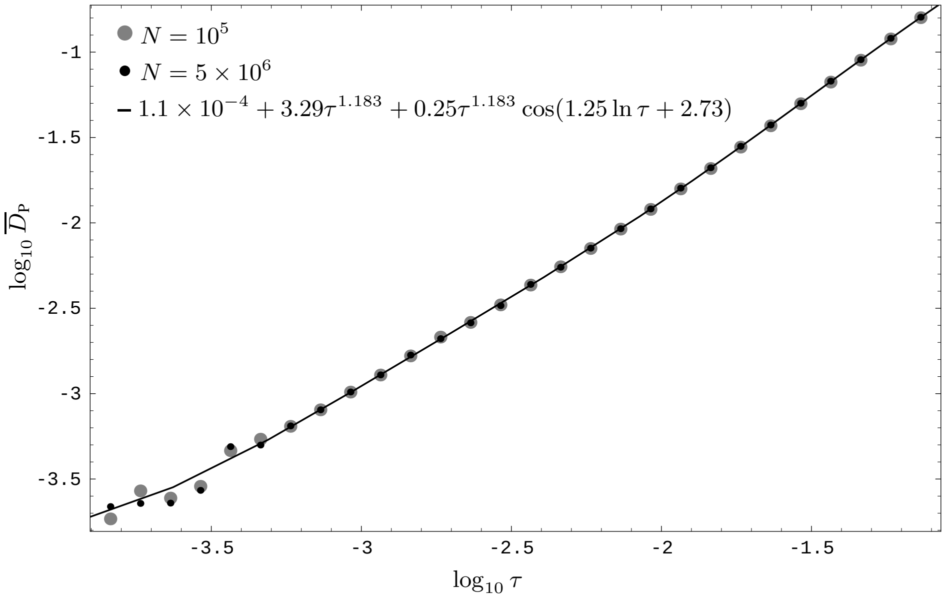

FIGURE 11. Poincaré-dimensions and log-periodic fit at $\rho = 3.08$.

of the 'corrected' Poincaré dimensions $\overline{D}_{\mathrm{P}}$ obtained from the $\overline{T}_{\mathrm{P}}$ for different values of N, see Figure 11, shows that these quantities do not depend strongly on N, lending some justification to this approach.

Remarkably it turns out that the behaviour of $\overline{D}_{\mathrm{P}}$ near $\tau = 0$ is not governed by a simple power law. Recently it has been observed in a number of areas that singularities in many natural phenomena exhibit *log-periodic oscillations* corresponding to complex critical exponents, see [33, Section 3] and references therein. In such cases, one generally expects the considered observable to behave asymptotically like $\mathrm{Re}\,\tau^{\beta+\mathrm{i}\omega} = \tau^{\beta}\cos(\omega\ln\tau)$. Therefore, we fit the Ansatz

$$\overline{D}_{\mathrm{P}}(\tau) \propto A + B\tau^{\beta} + C\tau^{\beta}\cos(\omega\ln\tau + \phi) \quad (\tau \to 0).$$

to the dataset for $\overline{D}_{\mathrm{P}}$. This yields the continuous line in Figure 11 as the best fit graph, and the corresponding value of the critical exponent is

$$\beta = 1.18 \pm 0.05.$$

Notice that the modification of the pure power law is rather small, as is expressed by the ratio $B/C \approx 3.3/0.25 = 13.2$. Similar analysis for $\rho = 3.12$ confirms that β is within the given error bounds.

6. CONCLUSIONS

Reiterating that general conclusions about prejudiced learning rules cannot, properly speaking, be derived by considering a single instance as we did above, we still want to note some of the indications which the detailed study of this special case provides us with.

The model for prejudiced learning and behaviour presented above is rather plain by its construction, which was focused on a few fundamental aspects we deemed characteristic. Its utter simplicity, although of conceptual beauty, appears as a drawback when its performance is critically assessed. For instance, it made the introduction of viability conditions necessary to ensure the desired functionality of the model, see Section 3.4. These conditions are unsavoury in rendering the three different classes of prejudiced learners incomparable. Yet, they are at least *intrinsic* conditions that can be satisfied by excluding certain values of the model's parameters from consideration. Totalling, the model itself calls for refinement to be applicable, and performant, in more realistic situations. Nevertheless, the basic performance results gathered in Section 4 support the heuristics of model building inasmuch as the class of adaptive agents learns at exponential speed, and can benefit from prejudice in a noisy environment by reducing their volatility significantly.

The most important phenomenological aspect of our model is doubtlessly the possible emergence of stubbornness as a *secondary* phenomenon of prejudice. In fact, for $\rho > 1$, it becomes the prevalent behaviour in noisy environments, aided by the mechanism of noise induced stability, a mechanism which is rather generic for nonlinear systems driven by noise and thus could well apply for other prejudiced learning rules than the logistic one. It must be emphasised that stubbornness cannot be ruled out as not being performant and therefore rare in realistic ensembles of prejudiced learners. The only condition for stubbornness to appear likely in individual agents is that the risk is initially underestimated (by the pertinent viability condition) and high, since $\rho = \alpha r$. Stubbornness thus is a high risk phenomenon. One scenario particularly catches the imagination. Assume for the moment that we have removed the awkward viability condition $\alpha < 1$ for adaptive agents, e.g., by replacing it with suitable boundary conditions. Then, agents of class A can be pushed into classes S or U by a sudden rise of risk, and consequently become stubborn.

The expectation to find the logistic prejudiced learning rule as such realised in learning systems and environments as complex as human beings and human societies is doubtful. Yet it might still be a candidate model for the description of certain behavioural aspects of biological systems from the level of single-cells to that of plants and lower animals. The thorough adherence to the principle of simplicity in the construction of the model, realised in the rather reduced game-theoretic framework, the minimal set of axioms, and finally the learning rule itself, in particular its memorylessness, are in favour of that view. Furthermore, the rich phenomenology of the model can aid its identification in such systems through the provision of many indicators — stubbornness being a prime one, alongside the characteristic adaptive behaviour with attenuated volatility.

Let us conclude by indicating some directions for further research, and posing some open questions.

In this first study, we have only considered static environments. One step to a more realistic model would be to include the effect of a statistical learning rule in it. This would generally lead to a slowly decreasing noise level, and in turn could let agents undergo behavioural changes. For example in class U, lowering the noise leads to cyclic behaviour ('evasive') through stochastic bifurcations.

It would further be interesting, along the lines sketched above, to improve the prejudiced learning rule itself by adding dynamical features. Within the framework of our model, making the prejudice parameter α dynamic naturally stands to reason, although too simplistic approaches would be inappropriate, cf. Section 3.4. In view of the heuristics for the performance of prejudiced learners developed in Section 4 it would seem promising, for example, to lower α when the agent has lived through a period of low volatility in the recent past. This modification would in particular pertain to S-agents, whose behaviour would then eventually become adaptive after some time in most cases, unlike S-agents with fixed α, and $\rho > 3$, which normally enter the uncertain (bifurcating) mode when the noise level decreases.

While we have concentrated here on the behaviour and performance of an individual prejudiced learner, studying the effects their presence will have at the community level in a society of learning agents is an important playground for further research. Good objects for such a study would be small-world networks [34, 35, 36, 37] as models for propagation of information. A prime question here is to what extent a proportion of prejudiced learners can stabilise beliefs in the given society. Studying these issues is work in progress.

This also, and finally, leads us to formulate some questions of an evolutionary kind. What quantitative relations between agents having various intrinsic parameters, i.e., belonging to different classes, would emanate through evolutionary pressure, over many generations of learners? The viability conditions of Section 3.4 reflect our presumptions on the net effect of evolutionary selection on the distribution of $(\alpha, r(0))$, and for A-agents we have seen in Section 4 that they have, in principle, a good chance to compete. Yet both the former presumptions and the latter claim still have to stand the test of more realistic models, including concrete performance measures and evolutionary selection schemes.

REFERENCES

[1] L. Tesfatsion: *Agent-Based Computational Economics*, ISU Economics Working Paper No. 1. Iowa State University, Ames, Iowa, 2003. http://www.econ.iastate.edu/tesfatsi/acewp1.pdf

[2] P. Fankhauser and T. Tesch, *Agents, a Broker, and Lies*, in Proceedings of the Ninth International Workshop on Research Issues in Data Engineering (Dimitrios Georgakopoulos, ed.), IEEE, 1999, pp. 56–63.

[3] ———, *Arbitration and Matchmaking for Agents with Conflicting Interests*, in Cooperative information agents III (Matthias Klusch, ed., Lecture Notes in Computer Science, Vol. 1652, Springer-Verlag, 1999, pp. 323–334.

[4] R. Grimm and P. Ochsenschläger, Binding Cooperation. A Formal Model for Electronic Commerce, *Computer Networks* **37** (2001), no. 2, 171–193.

[5] S. Rinaldi, *Love dynamics: the case of linear couples*, Appl. Math. Comput. **95** (1998), no. 2–3, 181–192. MR 99d:92057

[6] A. U. Schmidt, *A Model for Prejudiced Behaviour and a Noisy Dynamical System*, GMD Report 104, GMD – Forschungszentrum für Informationstechnik, 2000.

[7] J. P. Crutchfield and N. H. Packard, *Symbolic Dynamics of Noisy Chaos*, Physica D **7** (1983), 201–223.

[8] L. Arnold, *Random Dynamical Systems*, Springer Verlag, Berlin, Heidelberg, New York, 1998.

[9] R. Benzi, A. Sutera, and A. Vulpiani, *The mechanism of stochastic resonance*, J. Phys. A **14** (1981), no. 11, L453–L457.

[10] R. Benzi, G. Parisi, A. Sutera, and A. Vulpiani, *A theory of stochastic resonance in climatic change*, SIAM Journal on Applied Mathematics **43** (1983), no. 3, 565–587.

[11] L. Fronzoni, R. Mannella, P. V. E. McClintock, and F. Moss, *Postponement of Hopf Bifurcation by Multiplivative Coloured Noise*, Physical Reviews A **36** (1987), no. 2, 834–841.

[12] N. Sri Namachchivaya and S. T. Ariaratnam, *Stochastically perturbed hopf bifurcation*, International Journal of Non-Linear Mechanics **22** (1987), no. 5, 363–372.

[13] M. Yoshida, H. Hayashi, K. Tateno, and S. Ishizuka, *Stochastic resonance in the hippocampal ca3-ca1 model: a possible memory recall mechanism*, Neural Networks **15** (2002), no. 10, 1171–1183.

[14] M. I. Dykman, T. Horita, and J. Ross, *Statistical distribution and stochastic resonance in a periodically driven chemical system*, J. Chem. Phys. **103** (1995), no. 3, 966–972.

[15] A. Krawieckia and J. A. Holyst, *Stochastic resonance as a model for financial market crashes and bubbles*, Physica A **317** (2003), no. 3–4, 597–608.

[16] F. Chapeau-Blondeau and X. Godivier, *Theory of stochastic resonance in signal transmission by static nonlinear systems*, Phys. Rev. E **55** (1997), no. 2, 1478–1495.

[17] C. F. Lee and N. F. Johnson, *Exploiting randomness in quantum information processing*, Phys. Lett. A **301** (2002), no. 5–6, 343–349.

[18] R. D. Lee and H. Raiffa, *Games and Decisions*, Dover Publications Inc., New York, 1989, Reprint. Originally published: Wiley, New York, 1957.

[19] R. S. Sutton and A. G. Barto, *Reinforcement Learning*, MIT press, Cambridge, Massachusetts, London, England, 1998.

[20] V. N. Vapnik, *The nature of statistical learning theory*, Springer-Verlag, New York, 2000.

[21] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning. data mining, inference, and prediction.*, Springer-Verlag, New York, 2001.

[22] N. H. Bshouty, N. Eiron, and E. Kushilevitz, *PAC learning with nasty noise*, Theoretical Computer Science **288** (2002), no. 2, 255–275.

[23] G. Bonanno, *Information, Knowledge and Belief*, Bull. Econ. Res. **54** (2002), no. 1, 47–67. MR 1883624

[24] C. E. Hecht, *Statistical Thermodynamics and Kinetic Theory*, W. H. Freeman and Company, New York, 1990.

[25] A. Lasota and M. C. Mackey, *Probabilistic Properties of Deterministic Systems*, Cambridge University Press, Cambridge, New York , New Rochelle, Melbourne, Sydney, 1985.

[26] E. M. Bollt, L. Billings, and I. B. Schwartz, *A manifold independent approach to understanding transport in stochastic dynamical systems*, Physica D **173** (2002), no. 3–4, 153–177.

[27] M. Dellnitz and O. Junge, *On the Approximation of Complicated Dynamical Behavior*, SIAM Journal on Numerical Analysis **36** (1999), no. 2, 491–515.

[28] V. Loreto, G. Paladin, M. Pasquini, and A. Vulpiani, *Characterization of Chaos in Random Maps*, Physica A **232** (1996), 189–200.

[29] V. Loreto, G. Paladin, and A. Vulpiani, *Concept of Complexity in Random Dynamical Systems*, Physical Reviews E **53** (1996), no. 3, 2087–2098.

[30] M. Zinsmeister, *Thermodynamic Formalism and Holomorphic Dynamical Systems*, SMF/AMS Texts and Monographs, vol. 2, American Mathematical Society, 2000.

[31] J. J. Binney, N. J. Dorwick, A. J. Fisher, and M. E. J. Newman, *The Theory of Critical Phenomena*, Oxford University Press, Oxford, 1992.

[32] J. B. Gao, *Recurrence Time Statistics for Chaotic Systems and Their Applications*, Phyical Review Letters **83** (1999), no. 16, 3178–3181.

[33] A. Johansen and D. Sornette, *Finite-time singularity in the dynamics of the world population, economic and financial indices*, Phys. A **294** (2001), no. 3-4, 465–502.

[34] D. J. Watts and S. H. Strogatz, *Collective dynamics of 'small-world' networks*, Nature **393** (1998), 440–442.

[35] M. E. J. Newman and D. J. Watts, *Renormalization group analysis of the small-world network model*, Phys. Lett. A **263** (1999), no. 4-6, 341–346.

[36] M. E. J. Newman, C. Moore, and D. J. Watts, *Mean-Field Solution of the Small-World Network Model*, Phys. Rev. Lett. **84** (2000), no. 14, 3201–3204.

[37] F. Comellasa and M. Sampels, *Deterministic small-world networks*, Physica A **309** (2002), no. 1-2, 231–235.

FRAUNHOFER – INSTITUTE FOR SECURE INFORMATION TECHNOLOGY, DOLIVOSTRASSE 15, 64293 DARMSTADT, GERMANY

E-mail address: aschmidt@math.uni-frankfurt.de

URL: http://www.math.uni-frankfurt.de/~aschmidt