Natural Rationality

V. M. Darley
Stuart A. Kauffman

SFI WORKING PAPER: 1996-08-071

SANTA FE INSTITUTE

Natural Rationality

V.M.Darley*
Division of Applied Sciences
Harvard University
Cambridge, MA 02138

S.Kauffman†
Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501

Abstract

We propose a method for modelling economic systems in which outcomes depend locally on the predictions that agents make of other agents. We develop population games in which each agent adaptively searches for a good model of its environment. We demonstrate that such systems can exhibit persistent dynamics, cyclic oscillations between low and high complexity states, and other complex, yet endogenous, phenomena. We propose these 'adaptively rational' agents as a natural extension of rational expectations, suitable when mutual consistency is not realistic. We discuss the connections between our work and the programs of bounded rationality, evolutionary game theory and models motivated by statistical mechanics.

1 Introduction

We seek a definition of *homo economicus* which relaxes the assumption that globally optimal, mutually consistent strategies are selected by a system of autonomous agents. We believe the development, dynamics, stability and change of this optimality should be modelled, not graven in stone. We accept Nash's 1950 judgement that:

> "It is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal."

Ours is part of a growing body of research studying the complex relationship between population dynamics and Nash equilibria. In fact, an explicit population in modelling games was first introduced in Schelling's 1971 model of segregation in housing markets. Such models often go further than Nash's admittance that "actually, of course, we can only expect some sort of approximate equilibrium, since the information, its utilization, and the stability of the average frequencies will be imperfect." — with an explicit population we can address when Nash's 'approximate equilibrium' is only metastable or unstable (and therefore of lesser relevance) and model the actual dynamics in such cases. We adopt this approach and, for instance, we show that even for simple cooperative games, equilibrium is not assured: punctuated equilibria, meta-stability, limit cycles and other emergent properties may all be observed.

Sargent (1993) discusses a large variety of boundedly rational macroeconomic models, and neatly summarises the basic philosophy underlying that program of study:

* vince@das.harvard.edu
† stu@santafe.edu

> "Rational expectations imposes two requirements on economic models: individual rationality, and mutual consistency of perceptions about the environment ... I interpret a proposal to build models with 'boundedly rational agents' as a call to retreat from the second piece of rational expectations (mutual consistency of perceptions) by expelling rational agents from our model environments and replacing them with 'artificially intelligent' agents who behave like econometricians. These 'econometricians' theorize, estimate, and adapt in attempting to learn about probability distributions which, under rational expectations they already know."

We take a jointly boundedly rational, population based modelling approach, and present a class of evolutionary games, based upon the following assumptions:

- large numbers of agents interact locally in a multi-dimensional lattice;
- they use predictive models from a class specified by:
  - the past history of observations to use for predictions;
  - the complexity of model to construct using that data.
- in their interactions, the agents form predictions of one another and base their behaviours upon those predictions, in an attempt to optimise their utility — a function of the forthcoming and historical behaviour of the system;
- they change their choice of model according to past predictive success via a given decision rule (the evolutionary learning process).

So our model considers local interactions only, and deals with games in which there are a continuum of possible responses. See Blume (1993) for a related approach to $K \times K$ local interaction games, with an asymptotic analysis using techniques from statistical physics (which we shall discuss in §5), and Ellison (1993) for results comparing local with global interactions. In particular Ellison highlights the importance of transients and rates of convergence — if the rate of convergence of a system is slow compared with its expected unperturbed lifetime, then the history-dependent transient is more important than the equilibrium state. He proves the expected result that local interactions can give rise to much (e.g. 10^{100} times) faster convergence than global interactions.

Lane (1993) argues for the importance of such transients, and that the concept of a metastable state which may contain the seeds of its own destruction may be more important than pure stability. Our current work supports this stress on the importance of transients; indeed the concept of a metastable state is crucial to an understanding of the systems we consider.

We begin in §2 with some simple adaptive agents — our agents are not bayesian optimisers; rather they make predictions of other agents using their predictive model. They then optimise their utility in a myopic fashion based upon those predictions. The also adjust their predictive models over time, picking the best allowed adjustment that exists (details will be given in §3). We show how these simple adaptively optimising dynamics can drive agents playing such games away from steady state equilibria.

Given that simple equilibria are unstable, what is then possible? We expect a very interesting dynamic, based upon a refinement of the argument given in Kauffman (1993). We introduce our results in §4.1, but give an intuitive description here. In order to predict one another's behaviour optimally, complex adaptive agents will build optimally complex, and hence boundedly rational models of one another. Such adaptive agents might well coevolve to the 'edge of chaos' (the boundary region between a disordered and ordered world state), via the following framework: first, given finite data, models which can optimise the capacity to generalise accurately must be of optimal, intermediate complexity (neither overly complex nor overly simple models can generalise effectively). Second, when adaptive agents make models of one another as part of their mutual ongoing behaviour, the eventual failure of any finite, approximate model of another's behaviour drives substitution of a "nearby", optimally complex model of the other's behaviour which now appears to be the best fit to the other's behaviour. The point is that, given any finite set of data, multiple models of about the same complexity will fit the data roughly as well. As the data stream evolves, overlapping patches of the data are optimally fit by nearby models drawn from a set of models of the same complexity. Third, adaptive agents may persistently alter their models of one another's behaviour. Since a change in agent behaviour follows from a change in model, such agents must coevolve with one another using changing models of one another's behaviour. Fourth, presumably such coevolving behaviour can be chaotic, ordered, or at the edge of chaos. The ordered extreme in which changes to models are rare (so models are mutually consistent) corresponds closely to a rational expectations state, whilst the chaotic extreme corresponds to rapidly changing models of the agents. At the edge of chaos, models of one another would be poised, tending to change, unleashing avalanches of changes throughout the system of interacting agents. An appearance of punctuated equilibria seems likely. Fifth, a qualitative argument suggests that in a persistent attempt to optimise prediction about the behaviour of other agents, adaptive agents will alter their finite optimally complex models of one another so that the entire system approaches the edge of chaos.

We expect the following picture: if the dynamics are very stable and mutually consistent, then each agent has an abundance of reliable data about the behaviour of the other agents. Given more data, each agent naturally attempts to improve his capacity to generalise about the other agents' behaviour by constructing a *more precise model* of the others' actions. In our systems, this model is more sensitive to small alterations in other agents' behaviour. Thus as agents adopt more precise models to predict better, the coevolving system of agents tends to be driven from the ordered regime toward the chaotic regime. Conversely in the chaotic regime, each agent has very limited *reliable* data about the other agents' behaviour. Thus in order to optimise the capacity to generalise each agent is driven to build a *less precise model* of the other agents' behaviour. These less precise models are less sensitive to the vagaries of others' behaviour, and so the system is driven from the chaotic regime towards the ordered regime. So we expect an attracting state on or around the edge of chaos. The precise form of this attractor is discussed in §4.2.2. Furthermore we expect model precision and model complexity to be dynamically anticorrelated. This is because a stabilising world will be better predicted by a simpler model and vice-versa.

Thus we expect that the natural definition of *homo economics* which we seek is one in which agent complexity, information-use, forward-planning horizons, and recursive modelling depth (my model of your model of my model of ...) are dynamically constrained within a finite 'bubble' of activity, on average poised between the ordered and chaotic regimes of behaviour. This is the first step within that search, in which we address model complexity, precision and information-use.

After introducing our model and results, we present a mean-field analysis of the aforementioned state, and some extensions of that analysis based upon dynamical considerations. Finally we discuss extensions to our work and add a few general conclusions.

2 Adaptive agents

Given a system of interacting agents, "rational expectations" is taken to mean that the conditional expectation of the future each agent has is consistent with the future generated by actions based upon those very expectations. "Perfectly rational agents" act so as to maximise their expected utility under any given information constraints. Assumptions such as these have been used to build a rigorous mathematical formalism for economics. Now, most such rational expectations models implicitly or explicitly consider either

1. the representative agent framework, in which the entire economy is modelled with one aggregate agent;

2. only interactions between pairs of agents, or an agent and an averaged whole.

This research forms part of a growing body which seeks to redefine the assumptions of economics by modelling uncertainty, limited information and bounded rationality. A simple example will suffice to make our point. Given a system with a large number of possible Nash equilibria, to select a particular rational expectations state one requires additional parameters indexing agents' beliefs[1]. If one lets agents adaptively select not only amongst those parameters, but also amongst a larger class of beliefs of which rational expectations is but a subset, then what will the agents do? Will they evolve to a particular rational expectations equilibrium (possibly path dependent)?

If the answer is 'no', they must pick something quite different. It is within this context that we seek a definition of the natural rationality of agents whose beliefs are not arbitrarily constrained by assumptions such as rational expectations.

Our models can be considered as one of a string of possible relaxations of RE, moving from agents who adjust a fixed model in an attempt to learn the correct distribution (which is assumed to be of that form), to agents who adjust both model and model parameters, to populations of agents whose models predict *each other* and in which, therefore, the dynamics are completely endogenous.

We shall situate our models by describing some of the links in this chain, before arriving at §3 where we define our model using all of these ingredients.

Bray (1982) has studied a simple model of adaptation to a rational expectations state. In this model the equilibrium price p_t for a single commodity is given by:

$$p_t = a + bp^e_{t+1} + u_t \tag{1}$$

This is a market clearing condition with p^e_{t+1} the price that market participants expect at time $t+1$. Now the steady-state rational expectations solution is $p^e_{t+1} = \beta \quad \forall t$, where $\beta = a/(1-b)$. Bray addressed the question of what happens if we relax rational expectations and assume people form the expectation p^e_{t+1} by an average over past prices p_t. Bray's dealt directly with a least-squares learning process of a simple model, which effectively implied an equation of motion for $\beta_t = p^e_{t+1}$ (the expected price at time $t+1$), of the form:

$$\beta_t = \beta_{t-1} + \frac{1}{t}(p_{t-1} - \beta_{t-1}) \tag{2}$$

Bray showed that a representative agent learning in a system whose dynamics are governed by equations (1-2) will converge to the rational expectations equilibrium with probability one iff $b < 1$. The models we use are functionally more complex than Eq.(2), but they do not use information from the entire past. Applied to Bray's representative agent, our model is:

$$\beta_t = \alpha_0 + \sum_{i=1}^{c} \alpha_i p_{t-i} \tag{3}$$

for some positive integer c, where we will write $\boldsymbol{\alpha}$ as a shorthand for $(\alpha_0, \alpha_1, \ldots, \alpha_c)$. This class of models can approximate a wide variety of functional forms, and its range of dynamics include complex oscillatory timeseries. This begins to approach the natural goal of replacing the global expectation of a representative agent, with multiple expectations and predictions formed by a small group of agents.

[1]See (Kreps, 1990) for related discussion of equilibrium selection.

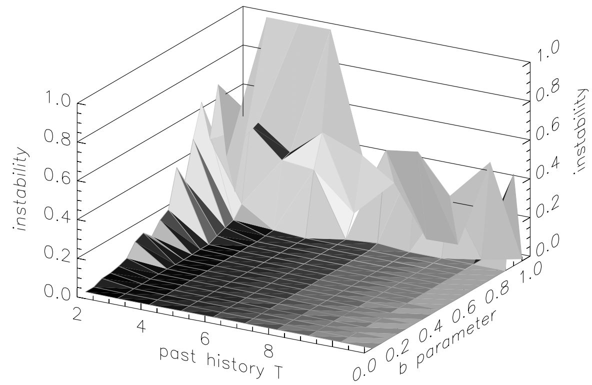

Figure 1: A phase portrait for update rule (6)

Model (3) is quite arbitrary as it stands. However if the agent optimises $\boldsymbol{\alpha}$ as a function of empirical observations, then it can build an accurate predictive model over time, and the model is then quite appropriate — especially for the case in which agents to not have prior knowledge of the expected equilibrium state (e.g. steady state, oscillatory, or more complex dynamics may all be encompassed within the same framework). Hence the dynamics of the price p_t are driven by the representative agent's behaviour in the following way:

$$p_t = a + b(\alpha_0 + \sum_{i=1}^{c} \alpha_i p_{t-i}) + u_t \tag{4}$$

where

$$\boldsymbol{\alpha} = \arg\min \frac{1}{T} \sum_{i=1}^{T} (p_{t-i} - \alpha_0 - \sum_{j=1}^{c} \alpha_j p_{t-i-j})^2, \tag{5}$$

where we have introduced another parameter T, the past history (number of lags) over which the agent optimises the mean square error in the predictions of its model. Now (4) is a simple linear rule, and (5) is an easy linear programming problem. However the combination of the two into an update rule of the form $p_t = F(p_{t-1}, \ldots, p_{t-T})$ is highly non-linear, and will generally possess a more rich dynamical structure — which in our models leads to instability.

If we consider the simplest non-trivial case, $c = 1$, then $\boldsymbol{\alpha} = (\alpha_0, \alpha_1)$ and:

$$p_t = a + b\frac{1}{\mathrm{Var}(p_{t-1})} \left(\langle p_t \rangle \langle p_{t-1} p_{t-1} \rangle - \langle p_{t-1} \rangle \langle p_t p_{t-1} \rangle - p_{t-1}.(\langle p_t \rangle \langle p_{t-1} \rangle - \langle p_t p_{t-1} \rangle) \right) \tag{6}$$

where $\langle p_{t'} \rangle = \frac{1}{T} \sum_{i=1}^{T} p_{t'-i}$, $\langle p_{t'} p_{t''} \rangle = \frac{1}{T} \sum_{i=1}^{T} p_{t'-i} p_{t''-i}$ and $\mathrm{Var}(p_{t'})$ is the variance of p_t over the range $t'-1, t'-2, \ldots, t'-T$.

Intuitively for large T, this system will be more stable, although stability is always tempered by the value of b. Figure 1 shows how the stability of this system varies as a function of T and b. This is a short-time-scale approximation to the true phase portrait which is independent of initial conditions in the presence of noise. In the figure, the 'instability' is a measure of how long the system spends away from the unique equilibrium point. The 'least squares optimisation' adaptive approach is one ingredient of the class of adaptive models we shall introduce in §3. There are many other examples in the literature, from bayesian learning, to genetic algorithms, neural networks and classifier systems. Sometimes such systems use a representative agent, sometimes a population of agents with heterogenous beliefs.[2]

[2]In the latter case, Arifovic (1991) has shown the convergence to a rational expectations equilibrium in Bray's model may take place even when $b \not< 1$.

More formally, for the simple case above, without loss of generality set $a = 0$, and assume we are near equilibrium, so $\langle p_t \rangle \approx 0$. In fact let us assume $\langle p_t \rangle = O(1/T)$. Then

$$p_t = bp_{t-1}\frac{\langle p_t p_{t-1}\rangle}{p_{t-1}^2} + O(\frac{1}{T}) + \text{noise} \tag{7}$$

Approximate solutions to this equation are of the form $p_t = e^{zt} + O(1/T)$, where $z = \frac{1}{2}\log b$. Therefore provided T is large, so the $O(1/T)$ terms are actually small, the update equation (6) will converge to the unique fixed point $a/(1-b)$ iff $0 \le b < 1$. For smaller T we anticipate values of b close to, but smaller than 1 to be unstable. This is precisely the picture calculated empirically in figure 1.

Finally, we should note that stabilising simple conditions on $\boldsymbol{\alpha}$ could be imposed ($\sum \alpha_i = 1$ is an obvious candidate). But in the short term such conditions would lead to a worsening of the agent's predictive ability. It is only if we wish to consider very foresightful agents who plan far into the future (and do not discount that future too rapidly) that such conditions might be realistic. We do not pursue that line of research here.

3 Model definition

We use the ingredients introduced earlier to model predictive economic systems[3] in which each agent bases its own behaviour upon predictions of a small group of other agents. Allowing the agents to change their beliefs over time, we observe these systems in computer simulations, and perform approximate theoretical analyses, to understand the nature of the system's attractor. The attractor may be a static equilibrium, or an attracting dynamic, and the conditions under which a given system evolves to one state or another must be determined. The predictive techniques used are similar to Equations (4,5) on page 5, yet even with such simple models, if we reject artificial stabilising constraints on agent models, very diverse dynamics are observed and convergence to static equilibrium is not to be expected.

We shall present a formal specification of our population models, including the predictive models the agents use, and the manner in which they update those models, and how their model predictions dictate their behaviour. This includes discussion of locality of neighbourhood, risk-awareness, heterogeneity, cooperation versus competition, and the 'stabilised' and 'forward' predictive scenarios we study.

The model we shall study for the rest of this paper may be formalised as follows:

Definition 1 *A predictive system $\mathcal{P}$ is a sextuplet $\langle N, L, A, B, M, g\rangle$, where:*

N is the number of agents in the system,
L is a regular lattice with exactly N sites,
$A = \{a_1, a_2, \dots, a_N\}$ is the set of agents,
B is a Euclidean space of possible behaviours,
M is the space of possible predictive models,
$u : B \times B^{|\mathcal{N}|} \to \mathbb{R}$ is the utility.

Here, the lattice L prescribes a function $\mathcal{N} : A \to A^r \subseteq 2^A$ which gives an agent's neighbours. The neighbourhood size r is constant, and we shall often write it as $|\mathcal{N}|$. The neighbourhood

[3]Our results may also be applied under some circumstances to biological contexts, in which individual bacteria, say, interact and compete and cooperate in the production of mutually useful chemicals and the absorption of resources from the environment.

relation will be reflexive, but not transitive, and we shall use multidimensional square lattices. Each agent has the same utility function. Throughout this paper we shall take B to be the set of reals $[0, 1000]$. Agents who try to exceed these limits are constrained to act within them.

The entire state of the system at any time t maybe specified by the time series of behaviours $b_t(a) = (b_t, b_{t-1}, b_{t-2}, \dots, 0)$ each agent a has generated, and the current model $\mu_t(a) \in M$ each agent uses. It is assumed the time series of behaviours are common knowledge but the agents' models are not. For simulation purposes a short artificial past $\{b_t; t < 0\}$ is generated randomly to initialise the system. We assume the information necessary for the evaluation of u, i.e. the past time-series, is common knowledge. This is a first order decision strategy.

The utility is used in two ways in our systems: (i) via forward expectations and maximisation to dictate an agent's behaviour; (ii) retrospectively to give the quality of an agent's model.

Both of these uses deal with local interactions and local information. This, and our wish to investigate global dynamics and coordination problems, is why we restrict u to be a function from the agent's own behaviour B and the local neighbourhood of behaviours $B^{|\mathcal{N}|}$ at a single time period (it is myopic). Thus we only use point expectations.

3.1 Utility Functions

Our agents act via a best-response dynamic. Given predictions of their local environment, they maximise their utility function (all agents have the same function, although they will obviously use that function on a different information set). There is a unique best response for the utility functions we consider, of which there are three general types (phrased in terms of best-response production quantities):

Coordination $b_{\text{br}} = \frac{1}{N}\sum_j b_j$; each agent must try to output precisely the average of the quantities that its neighbours output (so the utility function is '$-|b - \frac{1}{N}\sum_j b_j|$')

Substitution $b_{\text{br}} = 1000 - \sum_j b_j$; each agent acts so that the sum of its production and that of its neighbours is a fixed constant. Clearly each agent is taking part in a particular market of fixed size.

Coordination with Preferences $b_{\text{br}} = \lambda D + (1-\lambda).\frac{1}{N}\sum_j b_j$, so each agent must output a quantity biased between a fixed value and the average of its neighbours. The fixed value becomes the average for the whole system, and becomes a dominant attractor, with attraction given by the level of bias. A good interpretation of this is: each agent has a private valuation 'D', with a certain level of preference λ. Its action is a weighted sum of its personal preferences (to carry out 'D') and the influence of its neighbours (to coordinate). This kind of strategy is more naturally considered in the realm of social interaction and discrete choice (see (Brock and Durlauf, 1995) for such a model with local interactions).

This last game may also be considered a 'dominant Nash equilibrium' case of the coordination game. It is mathematically identical to the game in Bray's study of adaptive price formation (with $\lambda = b$ and $a = (1-\lambda)D$).

These three game types are special cases of the following best-response rule:

$$b_{\text{br}} = \mu + \lambda(\langle b \rangle - \mu) \tag{8}$$

This is a simple linear combination of the average of the locally observed behaviours ($\langle b \rangle = \frac{1}{N}\sum_j b_j$), with a single fixed point $b = \mu$ except for the case $\lambda = 1$ for which there are infinitely many fixed points, independent of μ. The utility function implied by (8) is $u(b, \{b_j\}) =$

$-|b - \mu - \lambda(\frac{1}{N}\sum_j b_j - \mu)|$. According to the value of λ we classify the best-response rule into the following categories:

$\lambda > 1$	unstable coordination game
$\lambda = 1$	basic coordination game
$0 < \lambda < 1$	dominant Nash/ coordination with preferences game
$\lambda = 0$	zero game
$-1 < \lambda < 0$	dominant substitution game
$\lambda = -1$	exact substitution game
$\lambda < -1$	unstable substitution game

We primarily study coordination games with $0 < \lambda \leq 1$, but we do also address our results to substitution games. These utility functions imply assumptions of risk-neutrality; a more general pay-off function would include observed predictive errors over some past history, weighted by both their variance and mean.[4]. A pay-off function which heavily penalises the variance in errors in preference to the mean is known as 'risk averse'. A risk neutral pay-off seeks only to minimise the mean. The pay-off function we use is risk-neutral and blind to the past.

3.2 Predictive models

At date 't', agent a_i is completely specified by its predictive model $\mu_t(a)$ and its past timeseries $b_t(a)$, where the space of predictive models satisfies the following properties:

Definition 2 *A predictive model $\mu \in M$ defines a triplet T, c and p, where:*

$T = T(\mu) \in \mathbb{N}$, *the length of history to use as data for predictions;*

$c = c(\mu) \in \mathbb{N}$, *the complexity of the predictive model;*

$p = p(\mu) : B^{T(1+|\mathcal{N}|)} \to B^{|\mathcal{N}|}$, *a mapping from local historical observations to future predictions.*

so that μ allows an agent to form a prediction of the subsequent behaviour of its neighbours, based upon their past behaviour for the previous T consecutive time-steps[5]. Clearly both T and c constrain the choice of p.

We use the predictor, $\boldsymbol{\alpha}$, as introduced earlier, except that it operates endogenously now (there is no global p_t). The basic model specifies how to calculate $\boldsymbol{\alpha}$ to predict an expected behaviour $b_{t+1,e}$ from its preceding c behaviours: b_t, b_{t-1}, ... , b_{t-c+1}, via the following linear recurrence relation:

$$b_{t+1,e} = \alpha_0^* + \sum_{t'=1}^{c} \alpha_{t'}^* b_{t+1-t'} \tag{9}$$

In order to calculate $\boldsymbol{\alpha}^*$, we must minimise the total error over the past history, of length T. For computational reasons the standard technique is to minimise the least squares error over some past[6]:

$$\text{error}_j^*(T,c) = \sum_{t=0}^{T-c-1} \{b_t^j - \alpha_0^* - \sum_{t'=1}^{c} \alpha_{t'}^* b_{t-t'}^j\}^2 \tag{10}$$

[4]In which case it may be necessary to separate the decision-making u from the 'pay-off' utility u.

[5]Other examples of the complexity could be the number of Fourier modes or the size of a correlation matrix used to fit the series.

[6]However we also consider random initialisation of the α followed by random adjustment over time in which better predictive models are retained (this is known as 'hill-climbing'), for which our experimental results are unchanged.

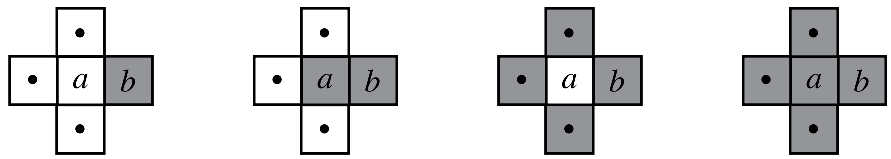

(a) Just b **(b)** a and b **(c)** All neighbours **(d)** a and all neighbours

Figure 2: For agent a to predict a given neighbour b using local information it can form a predictive model using any of the four data-sets shown. Furthermore, for (c) and (d) the agent may either use one or $|\mathcal{N}|$ models.

The sums here are given for the j'th agent, a_j with timeseries $\{b_t^j\}$. There are a large number of choices to consider when extending such a predictive model to a population in which local neighbours are modelled. The basic choices concern whether we believe different timeseries are related; whether our neighbours should be modelled individually or together, and whether an agent should include its own timeseries in its model estimation. It transpires that such choices are usually not important, so we will fix on the case in which each agent has a single model with $1 + c|\mathcal{N}|$ coefficients, which are calculated with a simultaneous regression on all the neighbours' timeseries. This model is then used to predict the average $\langle b_t \rangle$ directly. We now calculate the optimal $\boldsymbol{\alpha}^* = \{\alpha_0^*\} \cup \{\alpha_{j,1}^*, \dots, \alpha_{j,c}^*; j \in \mathcal{N}_i\}$, so that the prediction is given by:

$$b_{t+1}^{\text{pred}} = \alpha_0^* + \sum_{j \in \mathcal{N}_i} \sum_{t'=1}^{c} \alpha_{j,t'}^* b_{j,t+1-t'} \tag{11}$$

and the coefficients $\boldsymbol{\alpha}^*$ are given by a least-squares minimisation over the set of neighbouring lagged timeseries of length T. Further details on the more general cases in which agents may have separate models of each neighbour or use their own timeseries as regressors are given in Darley (1997). These scenarios are portrayed in figure 2.

The techniques used to perform the minimisations are singular-value decomposition (details to be given (Darley, 1997)). Let us note, however, that numerical stability factors (Press et al., 1992) mean that it is actually better to solve the minimisation problem rather than differentiating the above equations and solving the subsequent equalities exactly.

Under any of these techniques, once neighbouring predictions are calculated, the best-response rule/utility function dictate the agent's own behaviour.

3.3 Dynamical update rule

The system's dynamics stem from the ability of the agents to tune their predictive models, based upon differences between observed and expected behaviours, in an attempt to optimise their own behaviour with respect to the changing environment. We do not impose any exogenous shocks on our models. This procedure can be summarised by the following sequence of actions, carried out at every time-step, in parallel, by every agent a_i:

Prediction$(i, \mathcal{P})$ at time t:

i. Calculate private predictions $b_{t+1,e}^j \in B$ giving the expected behaviour of all other agents ($j \neq i$) at time $t+1$.

ii. Find $b^* = \arg\max_{b \in B} u(b, \{b_{t+1,e}^j\})$, agent a_i's predicted optimal response.

iii. Carry out action b^*.

iv. Observe actual behaviours b^j_{t+1} and, using g, calculate agent a_i's utility $u^* = u(b^*, \{b^j_{t+1}\})$.

v. If $\exists \mu'_i$ with $T' = T \pm 1$ or $c' = c \pm 1$ s.t. $u(b^*_{\mu'_i}, \{b^j_{t+1}\}) > u^*$ then pick the *best* such new model μ'_i. This is the model update rule.

Step (v) is the 'predictive model update rule' which dictates the main dynamical properties of our systems. Each agent compares its utility under its current model μ with the utility it would have had under a perturbed model μ'. If any of the alternative models would have been better the *best* such model is selected. Under our perturbation scheme, movement in the space of predictive models, M, is local, by steps of length 1 in either the T or c direction. Hence, discrete changes to an agent's model may be considered as moves in the two-dimensional discrete lattice of history-complexity pairs, and agents can be thought of as performing local search in this space as they seek the best model of the world around them.

As remarked earlier, the pay-off function is blind to the past, and the only effect of a particularly poor prediction is presumably a change in one's predictive model to a better pair of parameters. The only level of evolution in the system is the survival of certain models (and therefore certain parameter values). As a consequence we do not directly investigate phenomena pertaining to the *accumulation* of utility (rule (v) only compares instantaneous utilities).

Finally we shall note two facts: first, the sequence of update steps (i)-(v) is completely deterministic, with no exogenous perturbations or shocks; second, the entire model is evolving in synchrony. This differs from the approach of Blume (1993) who considers 'strategy revision opportunities' independently exponentially distributed, such that the probability of two or more agents updating simultaneously is zero.

3.3.1 Heterogeneity and the use of Information

Our agents are heterogenous because of two facts: that they may choose different models (or groups of models); and that the information upon which they base their predictions may be different.

Given some slight variation in the world (originally brought about at least by non-uniform initial conditions), each agent will operate with u upon a slightly different set of behaviours giving rise to a variety of actions. We shall investigate whether this initial heterogeneity grows, diminishes, or even vanishes with the evolution of the system.

It is worth pointing out that this research does not currently differentiate fully between the two sources of heterogeneity (models and information): at a first approximation the information used to model a given neighbour (its behaviour time-series) is common to all agents modelling that neighbour, so models are the only source of local heterogeneity. However, the manner in which correlations between agents are modelled is constrained to utilise slightly different neighbourhood information sources. This is because we currently constrain agents to use local information only to make their predictions. This means the four agents who make models of a given common neighbour are all constrained to use different information sets to construct their models if they wish to model correlations.

Further research will allow agents to use non-local information. Such more sophisticated models will enable us to pinpoint more accurately the assumptions from which heterogenous non-stationary worlds can be derived.

3.4 What to Predict?

Finally, given a class of predictive models, an economic strategy, and a set of neighbourhood relations, we must decide what the agents should predict and the manner in which that determines their own behaviours. We shall consider two scenarios:

(i) the 'forward predictive scenario' in which each agent predicts what its neighbours will do tomorrow and uses the utility u to determine its own behaviour. This is the natural, obvious predictive method.

(ii) the 'stabilised predictive scenario' in which each agent ignores the just realised predictions. The dynamics are as follows: all agents make predictions, adjust their models based upon the success of those predictions, but then those predictions are forgotten and ignored. The agents generate a new set of predictions and repeat. Hence the models are updated as above, but the behaviours and predictions undergo nothing more than a process of iterated refinement.

The first scenario is the natural one. We also study the second because it is effectively a stabilised version of the first. Rather than predicting forward in time, the agent effectively re-predicts what it should have just done, and then carries out that action (the process could be considered one of perpetual refinement of action, not unlike the simple adaptive processes considered earlier).

We will find that the dynamics which arise from these different scenarios can be quite different in character. Intuitively scenario (i) may lead to excessively unstable dynamics, as the agents' forward predictions diverge from one another. A more sophisticated predictive model may be required to follow the dynamics. Scenario (ii) on the other hand may be too stable, with predictions from one date to the next hardly varying at all.

A third and fourth scenario which we leave to future research are the following: (iii) each agent predicts using a discounted sum of expected future utility based upon neighbour predictions over a given planning horizon (and seeks to optimise the length of that horizon); (iv) agents build a hierarchy of meta-models of each other — my model of your model of my model of, ... and optimise the cut-off height of this hierarchy.

We shall present our results for scenario (ii) first, in §4.1, which illustrates the coherently organised coupling between model-update events. Then in §4.2 we give our results for scenario (i) in rather more detail.

4 Observation and Simulation

A system such as ours has a large number of possible rational expectations states. There are clearly an infinite number of possible Nash equilibria (using the coordination strategy at least), and many more periodic equilibrium states are possible. One important consideration when given multiple equilibria is to try and understand the problem of equilibrium selection. The parameters which characterise the choice of equilibrium index beliefs the agents have (individually and collectively) about their world. One common use of adaptive models is in equilibrium selection in just such a scenario. From our perspective we would like to understand more than just 'selection'; we would like to know what happens when models are not forced to be stable. This implies that our agents do not have as a priori (and rather ad hoc) beliefs that the world is heading inexorably to a simple static equilibrium (this is implicit in Bray's model and our extensions to it). Our agents will naturally pick from a class of models (which includes a set implying stationarity) so as to maximise their immediate gain.

So one question we must address is: "Is a static (coordinated) equilibrium selected for?". If not then we will concern ourselves with understanding and explaining whatever non-stationary dynamics are observed. In particular we attempt to formulate a categorisation of the *natural* rationality of agents whose mutual interactions form their world. We shall compare the results of our analysis and experimentation with the intuitive ideas introduced earlier.

In order to initiate the simulations, agents are given an artificial randomised past time-series, and history and complexity parameters. All the results are robust to changes in the initialisation technique (gaussian, sinusoidal and uniformly random initialisations have all been tested). We now present our experimental results for the two predictive scenarios.

4.1 Observations of the Stabilised Scenario

The following results are all for predictive scenario (ii). As remarked earlier, the stabilised scenario allows us to investigate feedback and propagation of information between models in a more restricted stable world. For this scenario, there is little difference in agent *behaviour* dynamics between the small and large population cases (the behaviours are simple and stable across the system). However the large case shows interesting spatial order in the space of agent *models*, so we shall consider that case exclusively.

We observe two regimes of behaviour, which we label 'coherent' and 'random', each preceded by a short transitory phase. These two regimes are qualitatively different.

During the first, coherent regime predictive errors decrease exponentially fast, whilst variance in agent behaviour ('system heterogeneity') collapses down exponentially fast onto the system mean. These are diffusive spatial dynamics in which any behavioural heterogeneity disperses rapidly. Such dynamics can be generated by a wide class of models in which agents try and imitate each other using simple adaptive models; such models are presented elsewhere (Darley, 1997). During the second, random regime, predictive errors and system heterogeneity have reached a lower bound at which they remain. These occurrence of these regimes is explained below.

Model update dynamics are more interesting: define an 'avalanche' to be the number of consecutive time-steps an agent spends adjusting its predictive parameters c, T. Then in the coherent regime, we observe a power-law relationship between avalanche frequency and size: $f = c/l^k$, where $k = 0.97 \pm 0.07$. Hence the model-update process has organised model change into a *critical* state in which large, long-duration avalanches may occur. This requires coordination to build up endogenously across significant sub-groups of the population of agents, so we refer to this state as 'self-organised'. In the random regime the avalanche frequency distribution develops into an exponential fall-off: $f = p^{l-1}(1-p)$, where $p = 0.323 \pm 0.005$. Furthermore, an examination of spatial (rather than temporal) avalanches in the lattice gives the *same* 'power-law then exponential' result. In that sense the system is scale-invariant in both space and time.

We should point out that there is a large *practical* difference between a power-law and an exponential fall off of avalanche size. For the exponential case, large avalanches effectively *never* occur, whereas in the coherent regime we have much data for disturbances of size nearing 100, for instance. This is an important distinction.

The reason for the dramatic change between regimes is as follows: in the first regime behaviour variance and predictive errors are both converging to zero exponentially fast (the former at least is expected for a diffusive dynamic). Once the differences between models' predictions are less than that discernible under the numerical representation used, numerical rounding errors contribute more than model difference. At that point model selection becomes a random process. This hypothesis has been confirmed using test experiments of varying floating point accuracy. The curve fit $p^{l-1}(1-p)$ in the random regime is just the expected number of consecutive moves in a random walk with probability p of moving.

The random regime is therefore of lesser importance, and the self-organised critical behaviour can be considered the predominant characteristic of the internal dynamics of the stabilised scenario.

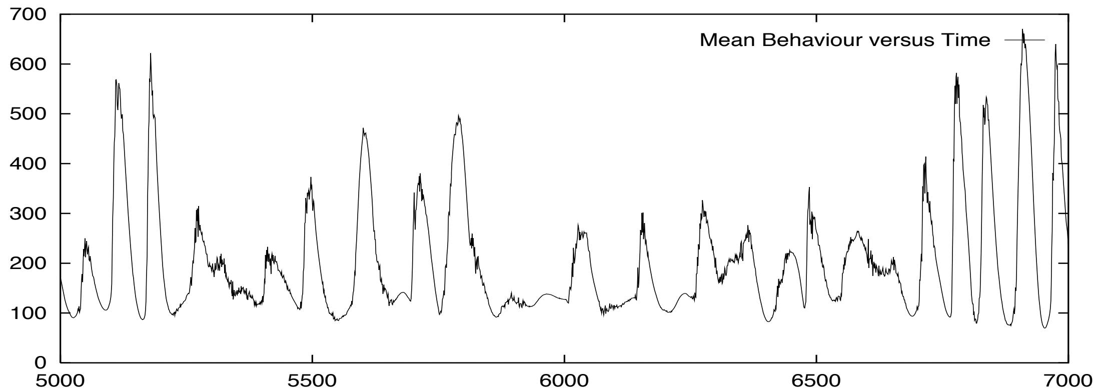

Figure 3: Mean Agent Behaviour in a 10×10 world. The maxima/minima represent almost completely homogenous behaviour; in between behaviour is highly disordered.

This critical behaviour exists in a wide variety of observables, indicating that the system does truly self-tune to a boundary intermediate between order and dis-order. Although coherence is eventually lost due to the overwhelming stability we have imposed, leading to a degenerate dynamic, it it clear that basic co-evolutionary forces between the agents' models have profound influence upon the global dynamics, and the macroscopic behaviour can be captured with a relatively simpler picture.

An extension of the above observations can be found if each agent models each of its neighbours using a totally separate model. We still find the same two regimes (and avalanche characteristics) as before, but now we get exponential convergence to a *non-uniform* state. Each agent has a different behaviour, but such behaviours are coordinated so the agents are still in a high-utility state. So whereas the old system converged to a system-wide fixed mean, zero variance; the new system converges to a system-wide fixed mean, but non-zero variance.

A basic issue which this raises is: persistent local diversity requires at least the capability of modelling that diversity — an agent which uses a single model for each of its neighbours (or, more generally, for the entire information set it observes) believes the system's equilibrium states are much simpler than the agent with multiple models who has no such presupposition. In systems such as ours in which the dynamics are endogenous, and agents' beliefs are reflected in those dynamics, and hence it is important to take those beliefs into consideration.

4.2 Observations of the Forward Predictive Scenario

The following results are all for predictive scenario (i). This is a more natural predictive situation, in which we can expect both models and behaviours to exhibit interesting dynamics.

The most clear difference between this and the stabilised scenario is in observations of the agents' behaviours. They no longer always settle down over time. There are two generic cases:

Small system or stabilising models — agents coordinate across the system on relatively simple strategies. The variance in behaviour across the system is very low, and its mean follows a simple monotonic path.

Large system — an interesting interplay between between periods of co-ordinated behaviour, and periods of dis-ordered rapidly changing behaviours is observed.

Figure 3 shows the basic behaviour-space dynamics for large systems. Comparing this with figure 6, we can see that the coordinated time-phases are precisely linked with periods of very low variance in agent behaviour, whereas the uncoordinated periods show very high variance levels.

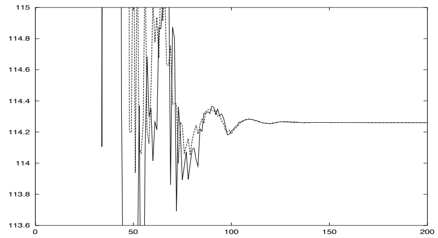

(a) 2 agents, unit neighbourhood — the models become consistent

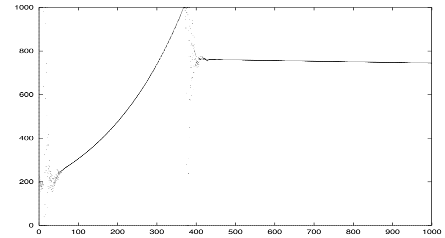

(b) 2 × 2 agents, $\mathcal{N} = 4$; the models become consistent

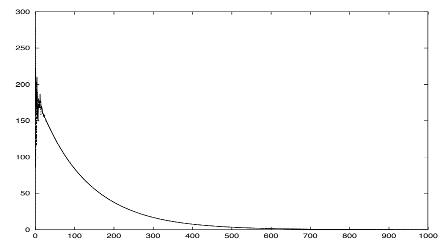

(c) 4 × 4 agents, $\mathcal{N} = 4$; again consistent models develop (very quickly in this case)

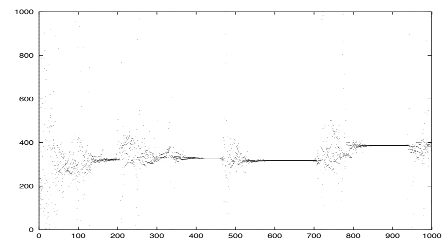

(d) 10 agents, unit neighbourhood. Meta-stability and punctuated equilibria are now observed generically.

Figure 4: Agent behaviours in small populations. The horizontal axes are time; vertical is behaviour. Each agent's behaviour at each time step is plotted as a dot, which are connected in the first figure.

Hence the agents *all* converge to a particular selection, retain that behaviour for some time, with only quite minor fluctuations, and finally the coordination breaks down and the agents pass through a dis-ordered bubble of activity, selecting any of a great range of behaviours before settling upon a new meta-stable state.

Consider the coordination and substitution games we introduced. We can summarise our results for these very succinctly: the 'unstable' variants are indeed unstable because best-response dynamics drive the system away from equilibrium; dominant-Nash games are identical to coordination games (which are the specific case with no domination, $\lambda = 1$) in the short term, but have more robust equilibria in the long-term. Their short term dynamics and the general dynamics of the coordination game exhibit interesting punctuated equilibria, cyclic oscillations, ... which we shall analyse in detail below.

We shall discuss the case with small populations first, before considering the large population case.

4.2.1 Small populations

The large scale dynamics described above are sufficiently complex that it is insightful in this case to consider small populations in which we can better see how the agents influence each other, and in which there is less room for variation, and fewer degrees of freedom to go unstable.

The first thing to learn from figure 4 is that the agents can learn to coordinate on particular behaviours, and that that coordination can take place in a relatively rapid, damped oscillatory form. Plots (b) and (c) show that coordination need not be on a static equilibrium, if the agents have a sufficiently large percentage of the total information available, i.e. if each agent is modelling a reasonably large fraction of the total number of agents, then global coordination on a t-dependent

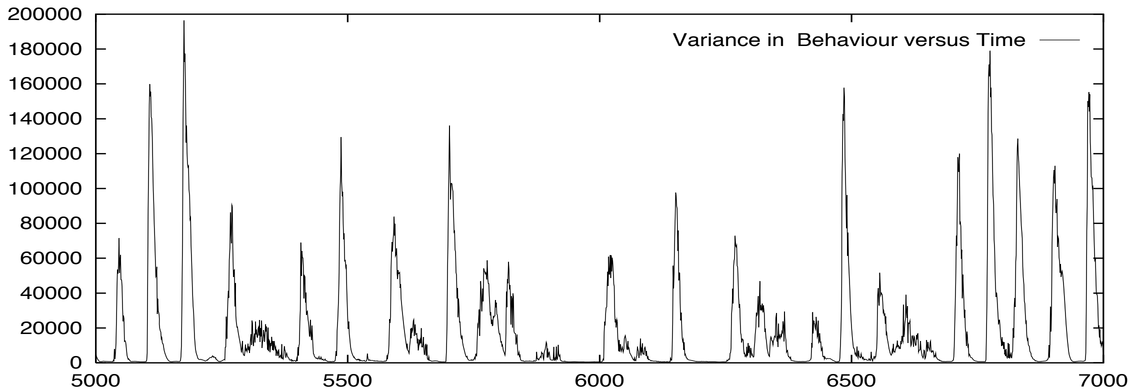

Figure 6: Variance in Agent Behaviour — it is clear that when variance is low it is vanishingly small.

path may be achieved. Notice that, since the agents are constrained to the range $[0, 1000]$ that this will causes a temporary instability on any path which hits the bounds.

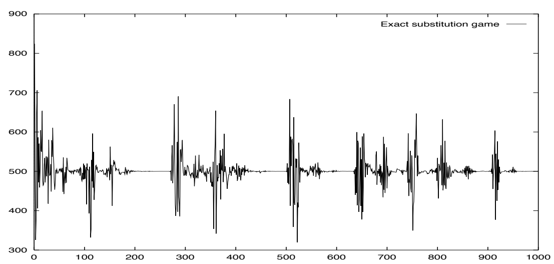

Figure 5: Behaviour of a sample agent taken from a ring of 10 playing the exact substitution game.

The exact paths in figures (a), (b) and (c) are completely history-dependent, although the range of possible qualitative dynamics is rather small. When we consider systems in which the proportion of agents modelled is low, a characteristic dynamic emerges. Equilibrium is only possible on a constant value, and those equilibria are meta-stable. The system will leap from one such equilibrium to the next via dis-ordered transitions. This *qualitative* dynamic is very robust, and the only historical influence is in the actual equilibrium values.

These results are all identical for the exact substitution game, for which we give a sample plot in figure 5. The stabler variants of these games, with $|\lambda| < 1$ damp down to the unique equilibrium after a λ-dependent time, for both coordination and substitution games. For large populations this still occurs, but only after significant transient periods.

4.2.2 Large populations

As explained above, the only characteristic simple coordinated states for larger populations are static equilibria. These occur with great regularity, and are highly uniform across the entire system. Figure 6 shows how the variance of the agents' behaviours drops almost to zero at these times — a spatial plot of such a state is shown in figure 7(a). Such states are only metastable, however, and hence oscillatory dynamics are observed[7].

All the large systems we consider show a punctuated equilibrium dynamic. Successive equilibria are destabilised endogenously, leading to wild fluctuations before the agents settle to another equilibrium.

Furthermore these dynamics are not transient for the exact coordination or substitution game. However, interestingly enough, given some small level of dominance in the game ($|\lambda| = 0.9$, say), the dynamics do become transient after sufficiently long periods of time. The agents eventually correlate their behaviours and models sufficiently that they reach of state of small fluctuations about a static equilibrium. The surprising dynamics of history and complexity by which this may be achieved

[7]The interested reader, with fast Internet access, should point their web browser at `http://www.fas.harvard.edu/~darley/Vince-Thesis.html` for some some movies of the 2-dimensional evolution of such a system.

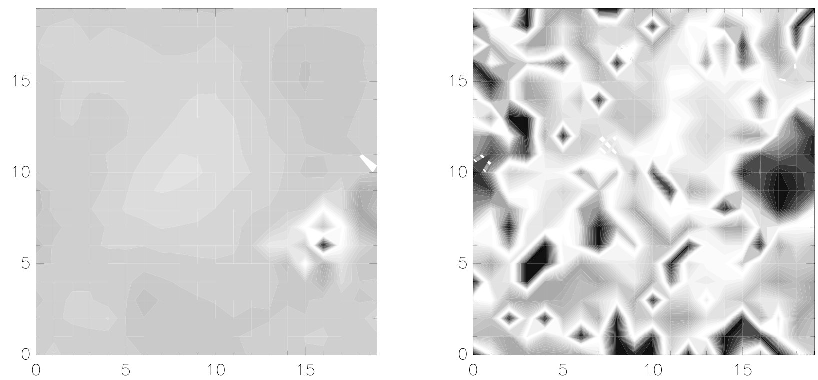

(a) A coordinated regime **(b)** An uncoordinated regime

Figure 7: Spatial distribution of behaviours for the forward predictive scenario. This is for a 20×20 square lattice of agents, with smoothing applied for representational purposes.

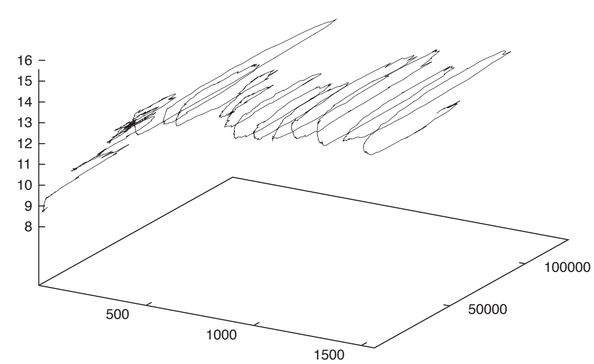

Figure 8: Limit cycle behaviour in $(t, \sigma_b^2, \langle c \rangle)$ space — the vertical axis is mean model complexity; the left-right axis is time and the depth axis is variance in agent behaviour (a measure of system heterogeneity). This is for a system with 900 agents in a toroidal 2-d lattice with a neighbourhood size of 4.

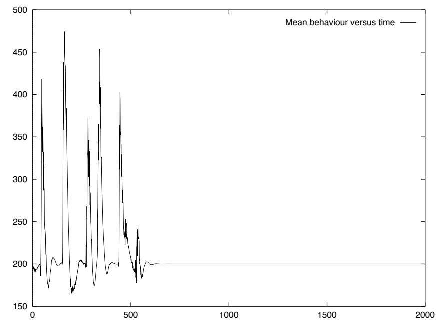

(a) Mean behaviour

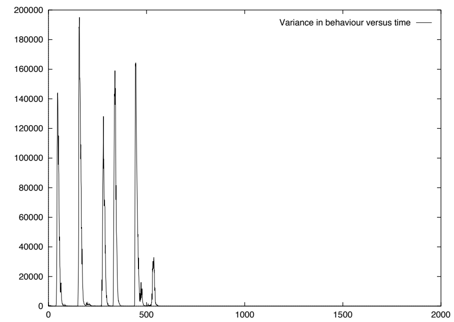

(b) Variance in behaviour

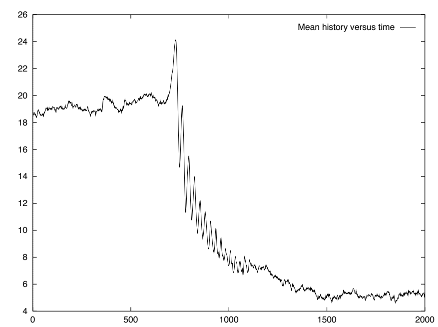

(c) Mean history

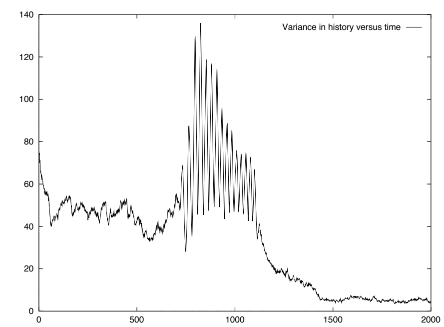

(d) Variance in history

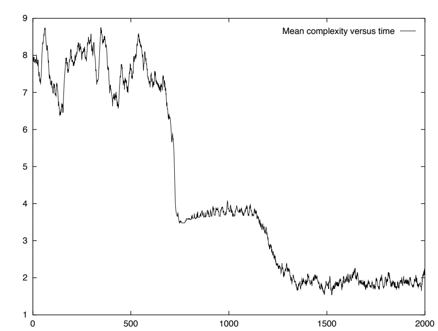

(e) Mean complexity

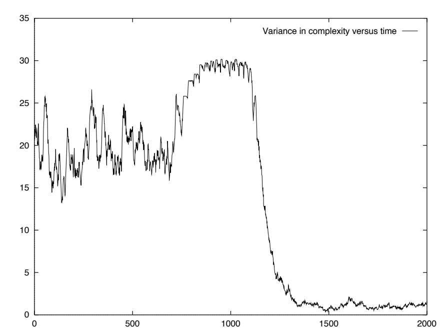

(f) Variance in complexity

Figure 9: Long transients for a dominant Nash game, followed by a particular transition to a stable equilibrium. The equilibrium persists for at least 10000 time steps. The form of the transition is in fact typical for this type of system.

are shown in figure 9. Both the early punctuated equilibrium regime, and the transition last for significant periods of time. Note, however, that external perturbations can de-stabilise that system, so that the 'natural' dynamics may be either the punctuated or static equilibrium depending upon the natural frequency with which exogenous perturbations occur. As remarked in our introduction, it is always important to remember that distinction when analysing systems with long transients, since that will dictate which of the possible behaviours is actually observed.

We will postpone a detailed analysis of the transition in figure 9 for future research, but we will point out a number of simple facts. That model history and complexity both eventually evolve to small values ($\langle c \rangle \approx 2$ and $\langle T \rangle \approx 5$) shows that the agents do learn to change their models (i.e. implicit beliefs) so that those beliefs are in accordance with the dynamics of the world around them. In this case the world has become very simple, and the agents learn that. This is an important point which is not so readily discernible in the presence of more unstable dynamics — for instance in the limit-cycle, figure 8, the mean complexity certainly reduces with system heterogeneity, but the world destabilises before the complexity is reduced too far. A possible intuitive explanation for why the reduction takes so long in figure 9 is that more complex models or models with longer histories are only discarded because of subtle effects of over-fitting the data. However, this is work in progress, and we expect to give more precise, rigorous reasons in a future analytical work.

Why does the system spontaneously de-stabilise? These instabilities start at a point, and propagate rapidly across the system[8]. They are caused by the following dynamic: a locally simple world state allows agents to build successively more precise models of the local dynamics (i.e. modelling the exact shape of local oscillations, with small expected error). Once an agent begins modelling the shape of the dynamics sufficiently closely, rather than some generic average, small 'extraneous' influences from not-so-nearby agents cause neighbouring changes which under the agent's model result in very large changes in local predictions. Within our systems, a very precise model (i.e. with a very small expected predictive error) is much more sensitive to local perturbations than a more generic average model. Hence these 'extraneous' perturbations will cause an agent with a precise model to predict dramatically different behaviours, and hence its own behaviour will leap accordingly. In a self-fulfilling manner, the neighbouring agents will observe that leap and change their subsequent predictions accordingly. Such large changes immediately propagate outward and a destabilised wave is observed, propagating across the entire uniform state of the system, and disrupting that equilibrium. The result is that all agents will pick models with more degrees of freedom but larger expected errors. In the extremely disordered case, these models will predict a basic weighted average, since their observations are too chaotic for precision to be valuable. Such 'averaging' models will then stabilise the system again and the cycle repeats.

5 Analysis

Condensed matter physics has a long history of analysing lattice models in general. These will often be abstractions of ferro- and para-magnetic materials, taking the form of Ising spin models (and numerous variants). In general, for such models, there exist parameters such as the temperature and external magnetic field which have a direct, profound influence upon the given system's dynamics. These models are very useful for describing and explaining experimental observations of physical systems. For example, a phase portrait of the equilibrium state under different values of the temperature and external field can be determined experimentally. In such a portrait, different regimes of behaviour are noted, separated by phase changes which are accompanied by critical transitions (exhibiting power-law scaling in correlations between changes to variables) and

[8] Again the interested reader is referred to the movies which have been made available on the internet to help gain an intuitive feel for these models.

discontinuities in certain observable variables (e.g. lattice magnetisation). By slowly varying the temperature or field we can move the system through these different regimes, and can watch the behaviour of the system's observables.

The predictive models we consider are fundamentally different to these basic spin models, in that there are no such obvious parameters to be tuned by the experimenter. In particular the temperature, which plays a crucial role in all statistical physics, has no immediate and clear interpretation (or role) in our models. Although we could perhaps observe an analogue of the temperature or at least of the entropy of the system, those values emerge from the system itself - they are self-tuned, and hence the whole system has the capability of moving itself through different regimes and phase changes. As an experimenter, we do not have control over these variables. As such, we call these systems *self-organising.*

Self-tuning models often exhibit interesting dynamical attractors, rather than static equilibria. Unlike an experimenter-tuned model, in which for a given parameter set a given equilibrium (or equivalence class of equilibria) is determined (provided there are no energy barriers preventing it from being reached in the time-scales under consideration), a self-tuning system relies upon its own dynamics to dictate the environment in which those very dynamics operate. Provided the system is not overweeningly stable, it can only really settle down to a dynamic equilibrium, a coherent mode of *active* behaviour which is self-reinforcing. The inherent dynamicity of the system will normally de-stabilise any potential static equilibria.

Such phenomena are observed in many realms in the real world, and an interesting paradigm for their study has developed from the work of Per Bak (Bak et al., 1987).

5.1 Mean-field approach

A standard approach from which to begin an analysis of such lattice systems is the mean-field approximation. We can apply this whenever the system's constituent parts interact in a primarily local fashion. We approximate by considering that each agent is only really interacting with a hypothetical 'mean-field', the average of the forces (i.e. field) generated by the actions of the agents in the rest of the system, perturbed slightly by its own actions (or that of a small cluster of agents).

5.1.1 Approximating assumptions

We shall make certain assumptions designed to contain a minimal specification of what we consider to be the essential properties of our models. We do not yet attempt to derive these approximations from the system's precise micro-dynamics:

(1) that one can derive from each agent a real number representing the sophistication of the agent's model 'σ'. Moreover, that each sophistication, denoted $s(\sigma)$, labels an equivalence class of models. Within each class there is variation in model, given in this simple case by the directional degree(s) of freedom of σ. As such we shall assume the set Σ of all possible σ is a vector space over the reals. Σ will be chosen as a succinct characterisation of the space of possible models.

Formally then, for every agent, a, we have a corresponding model given by σ_a, where we define the usual inner product over Σ:

$$\langle \sigma_a, \sigma_{a'} \rangle \to \mathbb{R} \tag{12}$$

So that the 'sophistication' of a given agent's model is: $s(\sigma_a) = \sqrt[+]{\langle \sigma_a, \sigma_a \rangle}$. We shall often write $|\sigma_a|^2$ for $s(\sigma_a)^2$. Small elements of Σ correspond to simpler models, and the directional degrees of freedom represent different types of behaviour. This defines a sophistication equivalence relation over Σ:

$$\sigma_a \sim \sigma_{a'} \Leftrightarrow s(\sigma_a) = s(\sigma_{a'}) \tag{13}$$

Two models (and therefore agents) are in the same equivalence class iff the models have the same sophistication.

(2) From the observable behaviours in the world, we can derive a vector field, such that an agent may be considered to be in a local behaviour field β' which is some externally imposed behaviour pattern β (assumed to be very small) plus a local field provided by the neighbouring agents.

$$\beta' = \beta'(\beta, \text{local agents}) \tag{14}$$

(3) There is a global difficulty of prediction, $\gamma \in \mathbb{R}^+$, derived from observations of the character and complexity of timeseries behaviour observed in the predictive system.

$$\gamma = \gamma(\{b(a)\}_{a \in A}) \tag{15}$$

5.1.2 Mean-field assumptions

Given the above assumptions, it is reasonable to assume that the average effective model σ_m in the field β' will be given by the following law:

$$\sigma_m \propto \frac{\text{behaviour field } \beta'}{\text{difficulty of prediction}} \tag{16}$$

This states the average agent's model is proportional to the local behaviour field and inversely proportional to the difficulty of prediction in the system as a whole. We assume that if prediction is hard the agents will tend to pick simpler models.

The mean field assumption is that the behaviour field due to neighbouring agents is a function of the average model σ_m. We shall consider σ_m to be small so we can expand β' in a power-series about β. Note that β is the sole source of exogenous effects; other terms are purely endogenous. The first term is given by:

$$\beta' = \beta + l\sigma_m \tag{17}$$

This relates the observed behaviour field β' to the character of the exogenous forcing and average local model. Note that it implicitly addresses the interactions from the point of view of a single individual. It can be extended to consider small clusters of individuals, but the qualitative consequences of the results are not changed. Now, writing γ for the difficulty of prediction,

$$\begin{aligned}\sigma_m &= c(\beta + l\sigma_m)/\gamma \\ &= \frac{c\beta}{\gamma - lc} = \frac{c\beta}{\gamma - \gamma_c}\end{aligned}$$

Where c, l are constants and $\gamma_c = lc$. This solution is only valid for $\gamma > \gamma_c$, otherwise σ_m points in the opposite direction to β which is not meaningful. For $\gamma < \gamma_c$ we must expand β' to third order (second order terms are not present for reasons of symmetry) in σ_m:

$$\beta' = \beta + (l - b\langle\sigma_m, \sigma_m\rangle)\sigma_m \tag{18}$$

$$\beta' = \beta + (l - b|\sigma_m|^2)\sigma_m \tag{19}$$

and we find:

$$\sigma_m \propto \beta^{\frac{1}{3}} \text{ for } \gamma = \gamma_c \tag{20}$$

For $\gamma < \gamma_c$,

$$\sigma_m^2 = (\gamma_c - \gamma)/cb \tag{21}$$

Following assumption (1), since σ_m is not just a real number, this implies that the mean sophistication $s(\sigma_m)$ is fixed by (21), but that the choice of σ_m within that equivalence class is *not* fixed, and a preferred direction will exist.

5.1.3 Interpretation

It is clear from the above analysis that the correlate of temperature in physical systems is the 'difficulty of prediction', and that there are two qualitatively different regimes of behaviour, given by $\gamma \gtrless \gamma_c$.

Therefore:

$\boxed{\gamma > \gamma_c}$ If the difficulty of prediction is high (the 'high temperature' regime), in the absence of external forcing ($\beta = 0$, as in the endogenously generated worlds we study) agents will:

1. Pick simple models (sophistication is proportional to the reciprocal of the difficulty).
2. Since the average of all models is the 'zero' model, there will be no preferred choice of σ within any given equivalence class of sophistications. Agents' models of the world are not mutually consistent (the significance of this point will be expanded upon later).

If there were to be external forcing, signified by β, then a preferred choice of model would exist. Currently this situation has no real interpretation in our predictive systems, $\mathcal{P}$, as they stand. However, were we to impose a certain pattern of cyclical variation in the world (incorporating, for instance, an economic 'sunspot' effect (Barnett et al., 1989)), we would expect the particular choice of agents' models to reflect the character of that forcing behaviour.

$\boxed{\gamma < \gamma_c}$ If the difficulty of prediction is low, even without any external forcing, the agents spontaneously symmetry break within their equivalence class of models, σ, to pick out a preferred type of model. This occurs because the directional degree of freedom of the average model σ_m is non-zero. That degree of freedom gives the preferred model.

The interpretation here is that by selecting a preferred model, the symmetry breaking ensures that the agents' models are *mutually consistent*. The system naturally reaches a state exhibiting an important ingredient of what are normally the assumptions of rational expectations.

Of course this description and analysis only apply directly to a static equilibrium scenario, which we do not expect to arise in these systems. However, one expects a dynamic equilibrium situation to situate those dynamics around any marginally unstable static equilibrium. Indeed the first step in taking the above analysis further is to consider the character of perturbations about the static equilibrium we have so far derived.

There are two important points still to address: that of the assumed 'difficulty of prediction' and that of interpreting the of agents' models as being mutually consistent or inconsistent.

Difficulty of Prediction Despite the form of assumption 3 on page 20, in the analysis which followed, the difficulty γ was treated as a tunable parameter, independent of the agents. In general it is certainly not independent of the agents, as assumption (3) states, it is a function of the agents' behaviours: $\gamma = \gamma(\{b(a)\}_{a \in A})$. Certainly if the agents behave in a reasonably varied fashion it will be easier to discern their underlying models and predict than if all agents behave similarly (given some underlying level of noise and that we wish to consider relative ability of predictions). So the difficulty of prediction will be an emergent observable, tuned by the dynamics of the system. Its dynamics over time will depend upon the following forces:

(i) the relative ability of a successful versus unsuccessful agent will be greater in a higher variance world;

(ii) agents with fancy models will be more susceptible to subtle changes in the behaviour of the world. i.e. a smaller discrepancy is capable of disproving their that models;

(iii) if the systems are sufficiently dynamic in the sense that agents must adjust their models over time in order to do relatively well in the world, then non-stationarity is maintained.

(iv) if the world is 'too' static or simple, an agent which encourages more varied behaviour in its environment will lead to the potential for relative success to be possible in its neighbourhood (following (i)), leading to destabilisation of a small region.

Consistency of Model The two regimes over γ, are characterised by:

$$\gamma > \gamma_c \Rightarrow \text{Agents have inconsistent models}$$
$$\gamma < \gamma_c \Rightarrow \text{Agents have consistent models}$$

What does this actually mean? Consider an equivalence class of models with the same sophistication $s(\sigma)$. Any given choice of σ within that class reflects a *hypothesis of the character of the observed dynamics in the world.*

A simple analogy would be that a given equivalence class selects a fixed number, $s(\sigma)$, of fourier modes (or splines or ...) to be used to model a given timeseries. The set of models within that equivalence class will be the set of all models which use exactly s independent fourier modes. Clearly the models within that class which utilise different fourier modes to describe a series reflect wildly differing hypotheses about the dynamics of the underlying process generating the given timeseries.

So, for the former case, the dynamics of the world prevent information of agent's models from transferring through the subsidiary medium of generated behaviours. Agents will pick inconsistent models of the world. In the latter case, information transfer is achieved, and agents will pick consistent models.

So, in the high difficulty regime, agents' hypotheses about the underlying dynamics of the world are mutually inconsistent (the dynamics of the world prevent information of agent's models from transferring through the subsidiary medium of generated behaviours). In the low difficulty regime agents will actually select, on average, a single model from a given class (information transfer is achieved). Thus the agents will have a common hypothesis for the world's dynamics.

Observations An estimate for the concept of 'difficulty of prediction', upon which the analysis is based, may be obtained from the observed predictive errors. Figure 10 shows how these vary over time, and a very clear correlation between time-periods of high-difficulty and uncoordinated behaviour can be discerned — compare with figures 3 and 6.

Although we have as yet no statistical estimates, the transitions between coordinated/uncoordinated and small/large errors are very sudden. This lends credence to the concept of an underlying phase transition.

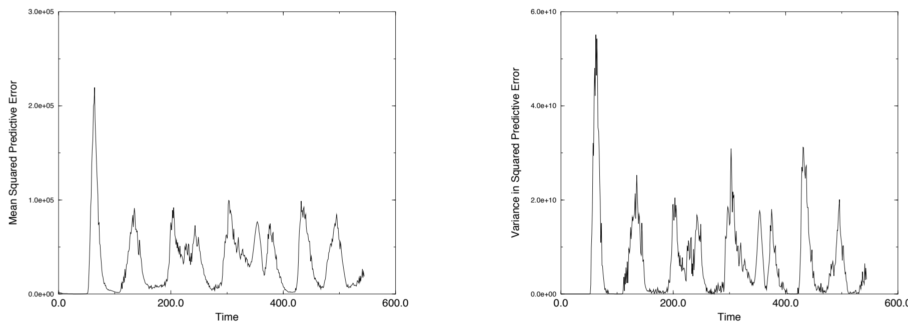

(a) Mean error **(b)** Variance in error

Figure 10: Averaged predictive errors over a 20×20 world. These can be considered approximate measures of observed predictive difficulty. Errors are uniformly small or uniformly large.

5.2 Statics and Dynamics

The difficulty of prediction γ is analogous in this exposition to the temperature parameter in magnetic phase transitions, with one crucial difference. It is not a exogenous parameter — it is an endogenous observable, driven by the very dynamics it implicitly tunes. At this point we leave as an open theoretical question whether its emergent dynamics drive the system to the consistent or inconsistent regime, or to fluctuate between the two. The above analysis can only be considered a static stability analysis, and must be extended to derive dynamical results.

Let us first note, however, if models are simple we would expect γ to be small. Hence a large γ encourages a small $|\sigma|$ which in turn produces simpler dynamics and a smaller γ. Conversely, as γ becomes smaller, $\sigma \propto \sqrt{\gamma_c - \gamma}$ increases and so we expect γ to increase. These opposing forces will therefore push γ away from extremal values. This is in agreement with the argument presented in the introduction and with empirical observations — we fluctuate between an approximate rational expectations state (in which agents' implicit expectations of the world are in almost perfect agreement), and a highly disordered state.

6 Discussion and Conclusions

We shall first discuss a number of limitations and assumptions on which our work is based, and what perhaps might be done to alleviate some of the more artificial constraints. We then summarise our goals, progress and results.

6.1 Utility Constraints and Generalisations

Throughout this study, we have confined ourselves to utility functions exhibiting a number of constraints (for the forward predictive scenario):

Historical Dependency — first note that, given some level of random noise, all the results of our simulations are 'independent' of the initial conditions. In particular the initial distribution of history and complexity values has no influence upon the long-term dynamics. By 'independent' we mean that the qualitative nature of the system's evolution is not dependent on initial conditions. In the absence of noise, one can find a few pathological initial conditions which are non-generic (completely uniform starting values for history, complexity, for example).

Convergence — the distinctions between coordination, substitution, and varying the λ parameter have already been discussed.

Continuity — all the games we use are smooth, such that a small difference in predictions causes only a small difference in behaviour. If the space of behaviours (or those behaviours selected by the game) is discretised, then most of the observations and predictions we make are no longer valid. We have not investigated whether this can be alleviated via the choice of a predictive model more suited to discrete predictions.

Variation — when making predictions one can use single models or multiple models; predict the average or the individuals. We have already discussed the trade-off between model stability and system size as it effects the system's dynamics.

Frequency — the trade-off between dynamicity of the system and the stabilising forces of model selection and prediction can be altered by changing the frequency with which the agents can adjust their predictive models (i.e. only carrying out the predictive update step (v) every n'th time-step). Our results are very robust against changes to this parameter: even if agents only adjust their model every, say, every 10 time steps, no difference is observed. This is further evidence (though no more than that) of the independence of our results of synchronous versus asynchronous modelling.

Topology — the neighbourhood size affects the rate of information propagation, and we have noted that system stability increases with neighbourhood size (for a fixed number of agents). The type of local neighbourhood is not important. If each agent's neighbours are given by a fixed random assignment from the set of N, the results are identical. This accords with our mean-field analysis.

6.2 Cooperative vs. Competitive Games

A significant limitation to our work is the following: the dynamics we have observed in these predictive systems are limited to games which are 'cooperative' in the sense that the best response for a given agent is not excessively bad for its neighbours. Some counter-examples such as the 'unstable substitution game' seem of little importance since we understand quite clearly why they are unstable. However other 'non-linear' games such as the undercut game (in which each agent tries to pick as a behaviour a number just a little bit smaller than its neighbours, unless they pick very small numbers in which case it is desirable to pick a very large one) seem to require a more sophisticated modelling and strategic approach than we provide. We have not investigated whether out results can be extended in these directions via a more competitive iterated system, in which multi-step forward planning, is incorporated, and non-pure strategies are allowed. Procedures like this allow cooperation to become reasonably stable in the iterated prisoners' dilemma, for example.

A direct extrapolation of our results suggests the intuition that such a forward planning horizon would be evolutionarily bounded in the same manner in which history and complexity are in our current systems.

6.3 Predictive Regime Changes

Examining the evolution of behaviours through time, as in figure 3, and in connection with the theoretical results, it is clear that no single, simple model should try to analyse and predict the coordinated and uncoordinated regimes together. A model which seeks to perform better should categorise the data from the past into blocks from each regime. It should use just the data from coordinated periods for predictions when the world is reasonably stable, and the remaining data for the uncoordinated periods. This requires the identification and selection of non-contiguous blocks of data, and hence a much more sophisticated predictive approach. These kind of value-judgements must be made when modelling real-world economic/market data — in a healthy economic climate, data affected by political turbulence, depressions and market crashes will not be used for basic modelling.

Our current research makes no attempt to address these issues. It is an interesting question whether that type of modelling would in fact stabilise our systems.

6.4 The accumulation of pay-offs

Ordinarily one associates with a pay-off function the gain of something of value (money, energy) which one would expect in any economic market or system of biological survival. However, these predictive systems operate solely via the use of a relative pay-off, in which the only use of the pay-off is in deciding which model should be used subsequently. The reason for this is that in our models there is no place for the accumulation of pay-offs; a more complex class of model would allow the introduction of such a scheme. In particular it will be possible to study the dynamics of price formation in market systems, an area in which economic theory is notoriously unsuccessful. Work in progress is beginning to achieve an understanding of how predictive agents interact to form prices.

Other extensions to this work, upon which we do not report in detail in this paper, are the following: (i) agents use a probability distribution over an entire space of models rather than a single model, and update that distribution (using a Bayesian rule) according to model success, creating a probability flow over model space. Our results there show that all agents evolve to attach non-zero weight to only a small, finite fraction of the space. In this sense, model *selection* is also bounded; (ii) extensions to Bray's model in which single or multiple agents imitate, learn and model each other in simple environments; (iii) forward planning; (iv) hierarchies of models.

6.5 Conclusions

It has long been observed (for a good discussion, see "More is different", Anderson (1972)) that qualitatively different emergent phenomena arise in large systems; usually accompanied by macroscopic events and correlations brought about by the *accumulation* of what are essentially simple, local microscopic dynamics. The spirit of this paper is that such collective behaviour can be harnessed in a dynamic theory of economic systems. Furthermore, we believe that interesting global, dynamical phenomena can arise out of systems in which the agents are relatively simple, with homogenous behaviour patterns based upon simple predictive rules. This is no simple case of 'complexity begets complexity'; we do not demonstrate that any given complex phenomenon can be recreated with complex, heterogenous underlying rules, rather we wish to show that, given a reasonable, simple class of agents, *they* can generate a wide-class of global, punctuated, metastable phenomena.

So, as a first step in that direction, we selected a class of predictive systems, $\mathcal{P}$, designed to relax rational expectations and stationarity assumptions towards a more natural unconstrained dynamic, in which agents' beliefs are allowed to evolve and hence select the most advantageous state. The systems were chosen to be as simple as possible, without exogenous shocks, but the games and rules were selected so as to bring about a large amount of *feedback* in the manner in which the interactions accumulate. It is because an individual's environment is no more than its neighbours, that changes in its behaviour may cause changes in its neighbours' behaviours which will then both propagate and feedback, generating an interesting dynamic.

6.5.1 Natural Rationality

We based our systems, and choice of agents, upon the need to investigate assumptions and questions of local interactions, coordination, optimal use of information, and the resulting dynamics of rationality. We gave heuristic arguments for why dynamics of this form will cause the system to drive itself towards interesting non-stationary dynamic equilibria, intermediate between conceivable ordered and dis-ordered regimes in its character, separated by a critical point.

A mean-field analysis of our systems shows very strong agreement with the qualitative nature of the simulated system's dynamics. There are two regimes: coordinated and uncoordinated, and the system is driven from one to another according to the level of difficulty of prediction in the world (although the mean-field model does not predict the fluctuations in detail), under a self-organising, endogenous dynamic.

Hence the system cycles between selecting for a static rational expectations state and breaking down into a dis-ordered state. We studied games with both an infinite number of Nash equilibria and a single equilibrium, expecting that a single equilibrium would stabilise the system so that persistent diversity would not be observed. However this was shown to be wrong. Punctuated equilibrium may be observed for very long transient periods even in systems with such a single fixed point. Under the influence of noise, such transients can become the predominant dynamic.

The manner in which these dynamics pass from coordinated phases to dis-ordered, uncoordinated bubbles of dynamics is very reminiscent of real-world market phenomena. In that sense our results show that as systems of autonomous agents grow in size, periods of highly disordered activity (stock market crashes or bubbles?) are to be *expected*, and are not an anomaly. These results hold even in the absence of noise or exogenous perturbations[9].

There are some connections which can be drawn with business cycles and the exploitation of market niches. For example, if we consider the static state to represent a market in which all the firms have the same strategy, then a firm which offers a slightly different kind of product/service may well be able to exploit the rest of the system's unsophistication. However our results suggest that such exploitation can *destabilise* the market, leading to a highly dis-ordered regime (bankruptcies etc.) from which a new standard emerges. It could be argued that the beginnings of such a process can be observed in the airline industry. On a rather different scale, the nature of the 'wheel of retailing' (in which firms evolve from innovative, small operations through a more sophisticated growth phase and finally become a larger, stabler more conservative retailer, opening a niche for yet another low-cost, innovative operation to enter the market) could perhaps be analysed via the same techniques (see McNair (1958); Brown (1988) for elaboration on this 'wheel')

In summary, our results demonstrate that agents evolve towards a state in which they use a limited selection of models of intermediate history and complexity, unless subjected to a dominant equilibrium when they learn to use simple models. It is *not* advantageous to pick models of ever increasing sophistication, nor to attach non-zero weight to all available information. As such, our agents in their mutual creation of an economic world, do not evolve towards a static rational expectations equilibrium, nor to what one would naively consider a perfectly rational state. Rather in their co-evolutionary creation of a non-stationary world, they evolve to use a level of rationality which is strictly bounded, constrained within a finite bubble of complexity, information-use, model-selection. As they fluctuate within those constraints, these agents can still achieve system-wide coordination, but the coordinated states are meta-stable. The dynamics are driven by the need to create models of greater/lesser precision to match the world's observed dynamics.

We interpret this state as an extension of rationality away from the simple static case to encompass the *natural rationality of non-stationary worlds*. Thus the natural (evolutionarily advantageous) rationality of these systems is a state in which information usage, model sophistication

[9] At a very low level there is some noise caused by rounding and other numerical phenomena.

and the selection of models are all bounded. It is in that sense that the natural rationality of our agents is bounded. Adaptive agents evolve to such a state, mutually creating a system which fluctuates between meta-stable coordinated and uncoordinated dynamics. This is the first step to understanding what agents *must* do in systems in which their mutual interactions coevolve to create a non-stationary world, and is therefore the first step to an understanding of *homo economicus* in such systems.

7 Acknowledgements

This research was undertaken under support from Harvard University and the Santa Fe Institute (with core funding from the John D. and Catherine T. MacArthur Foundation, the National Science Foundation and the U.S. Department of Energy).

This research has benefited enormously from discussions and time at the Santa Fe Institute, in particular with Steve Durlauf, and also with Eric Maskin at Harvard University, and Maja Matarić at the Volen Center for Complex Systems, Brandeis University. Thanks are also due to Ann Bell, James Nicholson, Dave Kane, Phil Auerswald and Bill Macready.

References

P. W. Anderson. More is different. *Science*, 177(4047):393–396, 1972.

Jasmina Arifovic. *Learning by Genetic Algorithms in Economic Environments.* PhD thesis, University of Chicago, 1991.

P. Bak, Chao Tang, and K. Wiesenfeld. Self-organized criticality: an explanation of 1/f noise. *Physical Review Letters*, 59(4):381–4, 1987.

William A. Barnett, John Geweke, and Karl Shell, editors. *International Symposium in Economic Theory and Econometrics (4th: 1987: Austin, Tex.)*, International symposia in economic theory and econometrics. Cambridge University Press, 1989.

Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.

Margaret M. Bray. Learning, estimation, and stability of rational expectations. *Journal of Economic Theory*, 26:318–39, 1982.

W. Brock and Steve Durlauf. Discrete choice with social interactions I: Theory. Technical report, University of Wisconsin, 1995.

Stephen Brown. The wheel of the wheel of retailing. *International Journal of Retailing*, 3(1):16–37, 1988.

Vincent M. Darley. Towards a theory of optimising, autonomous agents. Forthcoming PhD thesis, 1997.

Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.

Stuart A. Kauffman. *The origins of order: self organization and selection in evolution.* Oxford University Press, 1993.

David M. Kreps. *Game theory and economic modelling.* Oxford: Clarendon Press; New York: Oxford University Press, 1990.

David A. Lane. Artificial worlds and economics. parts 1 and 2. *Journal of Evolutionary Economics*, 3:89–108,177–197, 1993.

M.P. McNair. Significant trends and developments in the post-war period. In A.B. Smith, editor, *Competitive Distribution in a Free High Level Economcy and its Implications for the University*, pages 1–25. University of Pittsburg Press, 1958.

John Nash. *Non-cooperative games.* PhD thesis, Mathematics Department, Princeton University, 1950.

W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. F. Flannery. *Numerical Recipes in C.* Cambridge, 2nd edition, 1992.

Tom Sargent. *Bounded Rationality in Macroeconomics.* Oxford: Clarendon Press, 1993.

Thomas Schelling. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1:143–186, 1971.