

A Cybernetic Perspective on the Role of Noise in the Iterated Prisoner's Dilemma.

CSRP 506

Anil K Seth

Centre for Computational Neuroscience and Robotics
and School of Cognitive and Computing Sciences

University of Sussex

Brighton, BN1 9QH, UK

`anils@cogs.susx.ac.uk`

January 21, 1999

Abstract

An interpretation of the evolution of complexity in the Iterated Prisoner's Dilemma (IPD) is developed, based on Ashby's "law of requisite variety". It is demonstrated that the influence of *noise* on the evolutionary dynamics of this system is critically dependent on the locus of this noise. It is also argued that noise in such an evolving system is not merely, (or necessarily) a source of variation that must be adapted to, but, in certain circumstances, can facilitate the evolutionary exploration of increased areas of genotype space.

This convergence between artificial life/game theory and cybernetics holds implications for how we understand the general relationship between environmental complexity and agent complexity, and for our understanding of the role of noise in evolving systems.

1 Introduction

The general principle that there is organism complexity by virtue of environmental complexity has been well explored philosophically, [6],[13],[1],[4], and there have been several recent contributions, within artificial life, to an empirical foundation for this position, [5],[8],[9],[11],[10].

This report extends work presented in [9] demonstrating that the introduction of *noise* into a co-evolving Iterated Prisoner's Dilemma (IPD) ecology promoted the evolution of more complex strategies, with strategy complexity being measured by strategy "memory".

Here, we undertake an account of the evolution of complexity in an IPD ecology from a cybernetic, rather than game-theoretic, perspective. On the basis of this perspective, novel predictions are proposed for how the *locus* of noise in the IPD ecology may influence the evolutionary dynamics in the system. Evidence is also presented to suggest that noise may facilitate the evolutionary exploration of genotype space, independently of providing a source of variation to be adapted to, or "coped with".

The motivation behind this research is twofold. Firstly, cybernetic theory can indicate general principles for how adaptive systems should be understood; in the present case, in terms of the importance of the locus of variability. This research aims to demonstrate the *quantitative* translation of this principle to a computational context of general interest in artificial life. I will argue that this bridge between artificial life and cybernetics supports some implications that go beyond the IPD model itself, extending to more general considerations of the relationship between agent-side complexity and environmental complexity. For example, I will suggest later on that the common notion of "behaviour-based mechanism" (see, for example, [3]) may sit uncomfortably in this light.

A second motivation derives from the fact that the influence of *noise* within evolving systems, both in nature and in artificial evolution, is inescapable and imperfectly understood. I will argue that the two roles of noise identified here - in precipitating adaptive responses and in facilitating genetic search - are of sufficient interest to stimulate further research directed at assessing their generality.

We begin with a discussion of the particular aspect of cybernetic theory

that pertains to this research; Ashby’s “Law of Requisite Variety” (LRV), [1]. Then, in Section 2, the IPD model is rehearsed and its relation to the LRV is explained. Section 3 presents results describing the various ways in which the locus of noise *does* influence the evolutionary dynamics of the system. Section 4 presents evidence concerning the influence of noise on genetic mobility, and these results are discussed, and conclusions drawn in Sections 5 and 6.

1.1 The Law of Requisite Variety

Ashby [1] provides a general reason to suppose that agent complexity may evolve in a complex environment, in terms of his *Law of Requisite Variety (LRV)*. For an agent to maintain relative stability in certain (internal) essential variables (for example heart-rate, body core temperature), it must *prevent* the transmission of environmental variability through to these essential variables. In the same way that a good thermostat prevents the transmission of environmental variations in temperature through to a particular object (for example, the interior of a fridge should remain at a constant cool temperature despite the fluctuating temperature of a kitchen on a summers day), a good agent prevents the transmission of certain environmental variables (such as prevalence of food, proximity of predator) through to such essential variables as blood sugar or heart rate.¹

Given such a general situation, Ashby’s LRV is conceptually very simple. Consider a set of possible environmental disturbances, D , a set of possible responses on the part of the agent, R , and a set of possible outcomes, O . Consider also that for each D_i , there is a distinct outcome O_i , and a particular response R_i . There is also the ‘system’ S , which transforms a given (D_i, R_i) into the appropriate O_i . Stability in the essential variables requires minimising the variation in O , and this then requires that the variety in D is matched by the variety in R .² To quote:

If R ’s move is unvarying, then the variety in the outcomes will be as large as the variety in D ’s moves; only variety in R ’s moves can force down the variety in the outcomes. ([1], p.206).

¹Of course, these variables can and do vary within strict limits, but trespass beyond these limits is severely maladaptive.

²The LRV is also well stated in information theoretic terms, and corresponds to a theorem of Shannon, [12]: if noise appears in a message, the amount of noise that can be removed by a correction channel is limited to the amount of information that can be carried by that channel.

Ashby also reminds us that not all environmental variability need threaten the survival of the agent, and argues that variety comes in two fundamental forms:

There is that which threatens the survival of the gene-pattern. This part must be blocked at all costs. And there is that which, while it may threaten the gene-pattern, can be transformed (or re-coded) through the regulator R and used to block the effect of the remainder. (ibid, p.212).

According to Ashby, then, environmental variability can be either potentially beneficial or downright dangerous. However, in Section 4 we see that environmental variability can *facilitate genetic mobility*. This role for variety is not necessarily *either* beneficial *or* dangerous, and is not anticipated by any interpretation of the LRV.

2 The IPD Model

All the experiments described in this paper employ implementations of co-evolving IPD models. Their most immediate contribution is therefore towards a deeper understanding of the IPD model itself. The results on their own cannot, and do not, claim to represent adaptive, evolving systems in general. However, the simple nature of the IPD model, the broad scope of the cybernetic theory from which the hypotheses derive, and the ubiquity of noise in evolving systems, suggest that the *possibility* of there being wide application of these results should be taken seriously. We will return to this issue of generality in Section 6.3.

2.1 Structure of the Model

The IPD is a non-zero-sum game, for two or more players, in which each player chooses either to *cooperate* or *defect* on any given iteration. Each player is ignorant of the present move (though not necessarily of the history of moves) of its opponent, with payoffs being distributed according to Table 1. IPD models are of course renowned for providing insight into how cooperation can evolve in a population of selfish individuals [2], and this initial work has provided a versatile platform for further research. In [9] it was demonstrated that the introduction of *noise* into a co-evolving IPD ecology led to the

	<i>player 2 cooperates</i>	<i>player 2 defects</i>
<i>player 1 cooperates</i>	1:R=3 2:R=3	1:S=0 2:T=5
<i>player 1 defects</i>	1:T=5 2:S=0	1:P=1 2:P=1

Table 1: IPD payoff matrix; note that the actual scores are unimportant as long as $T > R > P > S$ and $2R > T + S$.

evolution of *complex* strategies. The ecology consisted of a population of variable length genotypes, with each genotype coding for a particular strategy for playing the IPD, and with the length of each genotype determining the number of previous iterations it could take account of when delivering its next move. This strategy “memory” can be taken to be a metric of complexity³. Noise was incorporated by introducing a certain probability that the move-to-be-made on any given iteration (as specified by the genotype) was “flipped” before the payoffs were calculated.

Two models were developed (for details, see [9]); a *compulsory* model, in which every member was forced to interact with every other member, and a *choice and refusal* model (the IPD/CR), in which each member maintained a set of expected payoff values for every other member, and could choose and refuse who to interact with. This latter model is drawn from [14] and is redescribed in [9]. A tournament-style genetic algorithm was employed in both models, with equal probabilities for genotype doubling, splitting, and point mutations (see Appendix 1 for parameter values).

2.2 The IPD and the LRV

Recalling the description of Ashby’s LRV in Section 1.1, we can think of an IPD ecology in terms of a disturbance D , response R , outcome O , and system S . D would be the set of all moves made over all iterations of the game, by all the opponents (of a given agent). R would be the strategy of the agent, S would be the IPD payoff table (Table 1), and O would be the set of scores awarded to the agent.

In a stable, cooperating population, the overall fitness of each agent (over many generations) will be maximised with minimum variation in O away from repeated cooperation payoffs (since mutual cooperation, in the IPD, maximises overall payoff in the population). Any environmental variation

³See [9] and [10] for a justification of this equivalence; also for an explanation of the genotype encoding scheme.

away from stable cooperation can then be expected to have a deleterious effect on fitness, unless it is countered by strategy that can effectively cope with such variation. For example, an occasional “accidental” defection will throw a population of “tit-for-tat” players into continual mutual defection but a “tit-for-two-tats” population will “absorb” such a defection, permitting general cooperative behaviour to persist throughout the population.

In terms of Ashby’s LRV, any variety in D can only be prevented from affecting O if it is countered through the action of R on S ; that is, if the variety in R (given S) matches that of D . This construal of the IPD allows us to predict that noise (variety) on D may lead to the evolution of more complex strategies (supporting variety in R) to counter this disturbance, but that noise on O will not. Noise on the outcome, after all, is precisely what the strategies would be expected to prevent, and if the noise is applied directly to O , then, by definition, *no* strategy can provide an effective response.

We can therefore distinguish two types of noise:

- *M-noise*: on each iteration, and for each player, there is a certain probability, ($p = 0.01$ in these experiments), that the move specified by the genotype is flipped, and only then are the payoff scores calculated with reference to the payoff table.
- *O-noise*: on each iteration, and for each player, there is a certain probability, ($p = 0.01$ in these experiments), that the payoff awarded is altered. Each time this occurs, the actual payoff awarded is selected randomly from the four possible payoff values.

We can now predict that M-noise may lead to the evolution of more complex strategies, but O-noise will not. This hypothesis will henceforth be referred to as the *LRV hypothesis*. Furthermore, we may expect that this hypothesis will only be confirmed in stable cooperating IPD populations, and *not* in unstable populations.

3 The LRV Hypothesis

The first set of experiments were conducted with a compulsory IPD model, the parameters of which are given in Appendix 1. Ten evolutionary runs,

of 10,000 generations each, were performed in each of 3 conditions; M-noise, O-noise, and zero noise.⁴

Figure 1 shows that without any noise, complex strategies never really evolve; the evolved memory tends to stay either at the initial level of 1, or drop to zero. Figure 2 illustrates that with M-noise, complex strategies *do* evolve; not always, and not always to the maximum, but it does happen. However, contrary to the LRV hypothesis, figure 3 illustrates that O-noise has the *same* effect as M-noise. Applying t-tests to the sets of means from each condition, these observations are statistically significant. There is significantly more evolved memory in the M-noise and O-noise conditions than in the zero noise conditions ($t = 2.52, df = 18, p < 0.05$ and $t = 2.51, df = 18, p < 0.05$ respectively). But there is *no* significant difference between the M-noise and the O-noise conditions ($t = 0.22, df = 18, p > 0.5$).

The second set of experiments were conducted with an IPD model incorporating choice and refusal (IPD/CR). In order to differentiate the conditions in this model, a small *cost* on complexity was applied (in all the conditions) by levying a fitness penalty on genotype length. Again, 10 runs were performed, of 5,000 generations each, in each condition. Here we clearly see that with either zero noise (figure 4) or O-noise (figure 6), complex strategies do not often arise. However, with M-noise (figure 5), the evolution of complex strategies is considerably more noticeable. Again, t-tests reveal these observations to be highly significant. There is significantly more evolved memory in the M-noise condition than with O-noise or zero noise ($t = 3.20, df = 18, p < 0.005$ and $t = 3.79, df = 18, p < 0.005$ respectively). But there is *no* significant difference between the zero noise and the O-noise conditions ($t = 1.42, df = 18, p > 0.1$).

The LRV hypothesis, in this model, is seen to hold true. The evolution of complexity only responds to variety in the environment, not to variety on the outcome.

So why is it that only the IPD/CR model produces results consistent with the LRV hypothesis? Table 2 clearly indicates that *only* the IPD/CR model

⁴Each run, of 10000 generations, took approximately 1 hour on a 143MHz single user Sun UltraSparc. The 5000 generations of the IPD/CR model required 3 hours of CPU time for each run.

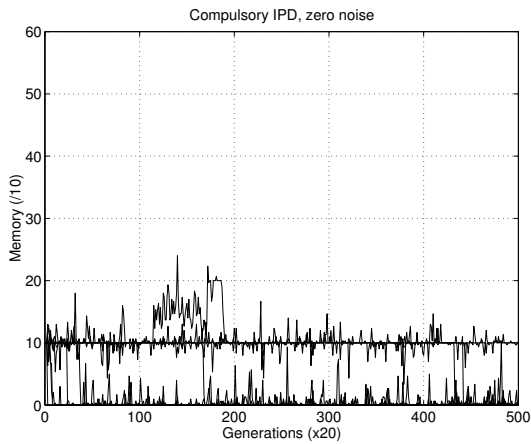


Figure 1: *Compulsory IPD; evolution of complexity with zero noise. Complexity does not evolve. Results from 10 evolutionary runs are superimposed.*

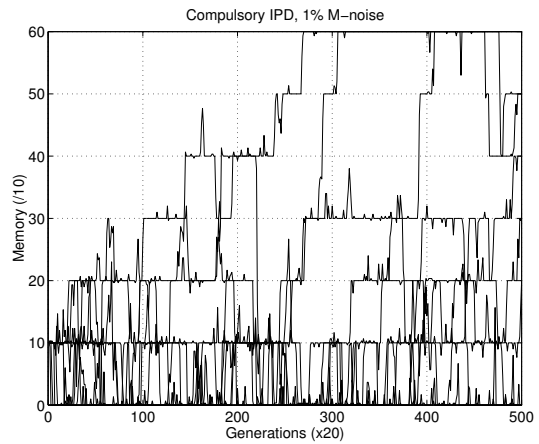


Figure 2: *Compulsory IPD; evolution of complexity with M-noise. Complexity does evolve. Results from 10 evolutionary runs are superimposed.*

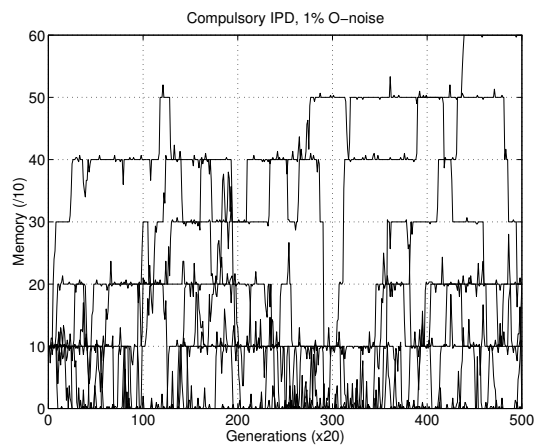


Figure 3: *Compulsory IPD; evolution of complexity with O-noise. Complexity does evolve. Results from 10 evolutionary runs are superimposed.*

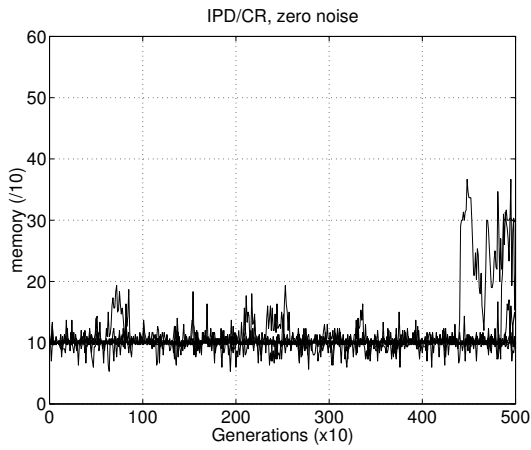


Figure 4: *IPD/CR; evolution of complexity with zero noise. Complexity does not evolve. Results from 10 evolutionary runs are superimposed.*




Figure 5: *IPD/CR; evolution of complexity with M-noise. Complexity does evolve. Results from 10 evolutionary runs are superimposed.*




Figure 6: *IPD/CR; evolution of complexity with O-noise. Complexity does not evolve. Results from 10 evolutionary runs are superimposed.*




Figure 7: *Compulsory IPD; cooperation dynamics in the three noise conditions, zero noise, M-noise, and O-noise. Cooperation is more prevalent in both M-noise and O-noise conditions. Only one evolutionary run is shown.*

	<i>IPD/CR</i>	<i>Compulsory IPD</i>
<i>mean cooperation</i>	86.36	13.95
<i>standard deviation</i>	0.37	10.48

Table 2: *Cooperation in compulsory model and in IPD/CR, with averages taken over all ten runs in all three conditions (M-noise, O-noise, and zero noise). Cooperation is much more prevalent and much less variable in the IPD/CR.*

delivers the stable cooperating population required for the LRV hypothesis to apply (see also figure 7 for typical cooperation dynamics in the compulsory model). In the compulsory model, there is already a great deal of instability in the outcome, and so there is no reason to expect complex strategies to evolve to cope with environmental variability. The question therefore arises; why do complex strategies *ever* evolve in the compulsory model in *either* O-noise or M-noise conditions? As we will see in the following section, this has to do with the facilitation of genetic search.

4 The Facilitation of Genetic Mobility

This section considers how O-noise and M-noise may influence the dynamics of the compulsory IPD model, given the unstable nature of cooperation in

this model. We first consider evidence that the evolution of complexity that we observe with both M-noise and O-noise in the compulsory IPD model is an example of *enhanced genetic mobility* rather than adaptation to “cope with” environmental variability. Then we consider a possible mechanism in terms of noise influencing the dynamics of cooperation, and evidence is presented to show that the exploration of genotype space (at least along the dimension reflected by strategy memory) is dependent on the presence of cooperation, or the presence of transients between cooperation and defection, in the population. Finally, it is argued that the locus of noise is still important in understanding the evolutionary dynamics of the system - although, rather than being reflected in the differences with regard to the evolution of complexity (as in the IPD/CR model), the difference here is reflected in the dynamics of cooperation.

4.1 Non-functional Evolution of Complexity

A compulsory IPD model is considered in which, although the heritable complexity of the strategies could vary as before, a maximum functional memory of 1 was imposed during every game. In this model, then, there is no possible functional difference, with regard to the IPD, between a strategy of memory 1 and a strategy of memory 6. Nevertheless, as figures 8 and 9 illustrate, both M-noise and O-noise conditions result in the evolution of complexity as before. This suggests that the more complex strategies are *not* evolving in response to any additional environmental complexities, since they could not possibly be employed for any such response. This is an initial hint that noise is indeed facilitating a more general exploration of genotype space rather than promoting a particular adaptive response.

A further set of experiments considered a compulsory IPD model in which the memory of the strategies was *genetically* limited to being either 0 or 1 (the average population memory could then range continuously from 0 to 1).⁵ The motivation here was to investigate differences between the three noise conditions that are independent of the evolution of complex strategies. In other words, to try and distinguish between the causes of the evolved complexity, and a mixture of the causes and consequences. If complexity is given

⁵Of course, this implies a corresponding functional limitation, as in the previous experiment set.

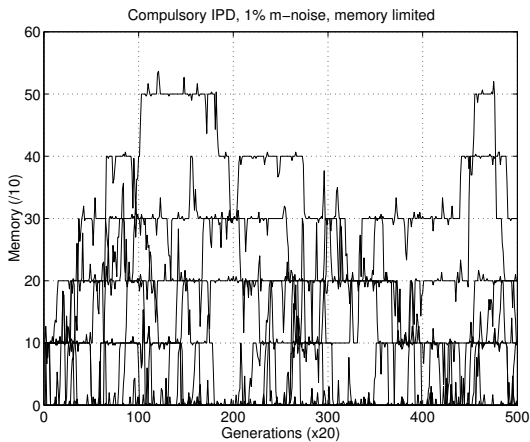


Figure 8: *Compulsory IPD with functional memory limit of 1; evolution of complexity with M-noise. Complexity does evolve. Results from 10 evolutionary runs are superimposed.*




Figure 9: *Compulsory IPD with functional memory limit of 1; evolution of complexity with O-noise. Complexity does evolve. Results from 10 evolutionary runs are superimposed.*

free reign to evolve, the causes and consequences of its evolution cannot be rigorously distinguished. The following experiments allow such a distinction to be drawn.

Twelve evolutionary runs were performed in each of the three conditions; zero noise, O-noise, and M-noise. Fig 10 illustrates that in both the O-noise and M-noise conditions, the average evolved memory is near to the expected average (0.5). But in the zero noise condition, in each case the average evolved memory either stays near the initial level of 1, or is very close to 0. This indicates that both O-noise and M-noise are permitting evolution to wander easily within the prescribed area of genotype space (determined by the maximum memory limit of 1), but with zero noise, the evolving population always becomes “stuck” on either memory 1 or memory 0. Furthermore, the means of these average memories are not significantly different between any of the conditions. The average (of the average memory) in the M-noise condition is 3.67; in the O-noise, 4.25; and with zero noise, 5.25. Applying t-tests to these data sets indicate that these overall means are not significantly different from each other (between M and O, $t = -1.33$, $df = 22$, $p > 0.1$, between M and zero noise, $t = -1.25$, $df = 22$, $p > 0.1$, and between O and

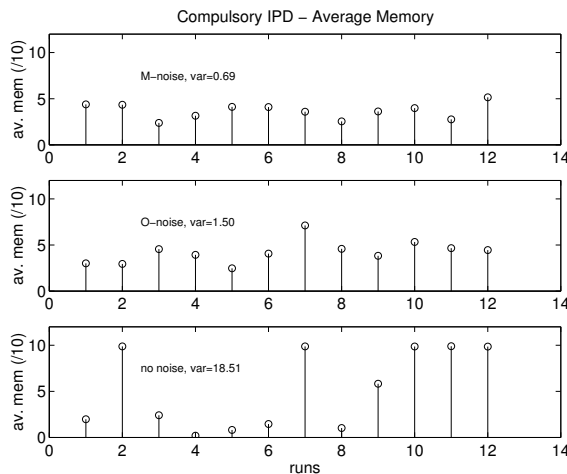


Figure 10: *Compulsory IPD; memory limited to 0-1 range. Both M-noise and O-noise permit average memory to approach the expected average; without noise, the average memory appears to be “stuck” at one extreme or the other.*

zero noise $t = -0.78, df = 22, p > 0.1$). So neither type of noise is actually encouraging (or discouraging) the evolution of complexity *per se* in this model.

4.2 Influence of Noise on Cooperation

It was also observed that both O-noise and M-noise increase the amount of cooperation in this memory-limited compulsory IPD model. Both M-noise and O-noise might be expected to have such an effect since, given a predominantly defecting population, both kinds of noise would permit a cooperative move to occasionally score well. Fig 7 indicates that typical cooperation in the zero-noise compulsory IPD model is very low, and indeed the average prevalence of cooperation over all ten runs was only 4.43% (standard deviation 2.43), and also seems to indicate that both M-noise and O-noise enhance population cooperation.

This intuition was tested in the memory-limited model, and fig 11 provides summary data, collected from the 12 runs in each condition, indeed indicating that both O-noise and M-noise significantly increase the level of population cooperation from that present with zero noise ($t = 3.39, df = 22, p < 0.001$, and $t = 5.49, df = 22, p < 0.001$ respectively), with the




Figure 11: *Compulsory IPD; memory limited to 0-1 range; O-noise enhances population cooperativity and M-noise enhances population L-Z complexity. Standard deviations indicated by error-bars.*

effect being significantly more pronounced for the O-noise condition ($t = 4.62, df = 22, p < 0.001$). This graph also illustrates how population *stability* varies across the different conditions. Stability was measured by calculating the *Lempel-Ziv* complexity of the cooperation for each run, (see [7]). This calculation involved noticing that in the compulsory IPD model, the population, at any time, was generally either completely cooperating or completely defecting. Thus, the population cooperation over x generations can be recast as a binary string of length x . The Lempel-Ziv (L-Z) complexity measure delivers the size of the minimum program required to generate a binary string, and this is minimal for a uniform string and maximal for a completely random string. Fig 11 is therefore indicating that both M-noise and O-noise lead to significantly greater levels of population cooperation instability than the zero noise condition ($t = 10.16, df = 22, p < 0.001$, and $t = 4.19, df = 22, p < 0.001$ respectively), with the effect being significantly more pronounced in the M-noise condition ($t = 5.60, df = 22, p < 0.001$).

So both O-noise and M-noise influence the population stability dynamics, but in different ways. Both increase the overall level of cooperation, but whereas for O-noise this effect is characterised by longer periods of steady cooperation, for M-noise increased population instability is the more pronounced effect. Note that this effect is also visible in the standard (non-

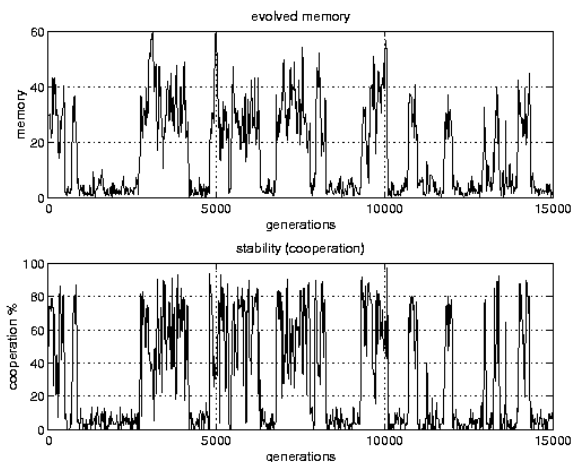


Figure 12: *Compulsory IPD model; there is a high correlation between complexity and stability (high mutation rate used).*

memory-limited) compulsory IPD model (see fig 7).

4.3 Dependence of Complexity on Cooperation

We have seen that both M-noise and O-noise influence both the exploration of genotype space (at least with regard to memory) and the cooperation stability dynamics. Evidence that these two phenomena are related is provided by a highly visible correlation between the mean population cooperation and the evolved memory in the compulsory IPD model. Fig 12 show results with M-noise; the model has been run with a particularly high mutation rate (0.05) in order to increase the population instability. This figure suggests that for genetic mobility to be present (in terms of memory), cooperation and/or transients between defection and cooperation need to be present.

Evidence that the correlation in fig 12 is a general phenomenon is provided by a further correlation between the L-Z measure of population cooperation instability (reflecting the number of transients), and the *variance* in the evolved memory, over many different evolutionary runs. Fig 13 plots summary data from 16 evolutionary runs of the standard compulsory IPD model; 8 with M-noise, and 8 with O-noise, in which the memory was allowed to vary between the normal limits (0 to 6), and in which the overall variance




Figure 13: *Compulsory IPD with 0.01 M and O-noise; L-Z instability (complexity) of the environment correlates with variance in the evolved memory. The correlation co-efficient, disregarding the two outliers, is 0.86.*

in evolved memory was also measured. The correlation, whilst not perfect, is clear.

Clearly, mobility in genotype space is linked to the structure of the fitness landscape (as indicated by the prevalence and stability of population cooperation). Both M-noise and O-noise are able to influence this structure, and they do so in different ways.

5 Functional Adaptation in the IPD/CR

Let us finish by returning to the IPD/CR model, in which stable cooperation was always observed, and in which the LRV hypothesis was observed to hold true. In the light of the previous sections, it is important to confirm that the evolution of complexity observed in the IPD/CR model *is* functionally significant - otherwise there would be no sense in arguing that the longer memories had evolved *in response to* some characteristic of the environment dependent on the introduction of M-noise. To confirm this, a functional memory limit was applied to an IPD/CR model with M-noise, just as was applied to the compulsory model in Section 4.1.

Fig 14 illustrates that the evolution of complexity *is* abolished when the effective memory is limited (let us call this the ML condition). Applying t-tests, there are no significant differences between the ML condition and either the O-noise or zero noise conditions ($t = 1.21, df = 18, p > 0.1$, and $t = 0.68, df = 18, p > 0.1$ respectively). But the difference between the M-




Figure 14: *IPD/CR; evolution of complexity with M-noise; effective (but not heritable) memory limited to 0-1 range.*

noise condition and the ML condition is highly significant ($t = 3.54, df = 18, p < 0.005$). These results clearly contrast with those obtained from the compulsory model, and strongly suggest that in the IPD/CR model, the evolved complexity is indeed functionally significant, just as predicted by the LRV hypothesis.

6 Discussion

Two IPD models have been described; a compulsory model, where cooperation is unstable, and an IPD/CR model, where cooperation is stable. Only in the latter model would we expect the LRV hypothesis to apply, and only in the latter model was it upheld. However, the compulsory model has provided evidence that noise can facilitate genetic mobility. We will consider each point in turn, and then turn to some implications.

6.1 The LRV Hypothesis Revisited

In the IPD/CR model, M-noise promotes the evolution of complexity but O-noise does not, and, furthermore, this complexity is functionally significant. The locus of noise is therefore critically important in understanding how the evolutionary dynamics of the system unfold. However, it may still not be the

case that the variety of the response is *matching* the variety of the M-noise, in the strict sense implied by the LRV (Section 1.1). Indeed, figure 2 illustrates that the complexity response to a steady level of M-noise is generally unpredictable. It is therefore not possible to say that the complex strategies are adapting *to* the M-noise, and to that alone. What we can say is that the complex strategies are adapting to some aspect of the environment dependent on the introduction of M-noise, and which is almost certainly heavily influenced by the co-evolutionary nature of the IPD model. The exact nature of the variety of the environment presented by (and to) any given population will not be predictable (thanks to co-evolutionary dynamics), and so a stable and predictable response to the introduction of M-noise will not be observed. So, although the LRV undoubtedly applies to this situation, these experiments also counsel against any simple interpretation that considers noise merely as a pernicious aspect of an environment to be coped with in a regular way.

6.2 Genetic Mobility

Evidence from the compulsory IPD model indicates that both M-noise and O-noise promote genetic mobility (at least in the dimension reflected by strategy memory), and can influence population cooperation dynamics. There is also evidence that these phenomena are linked. This is an illustration of a novel role for noise in evolving systems, and does not follow from any interpretation of the LRV (where, if variety is significant to an agent, it is always *functionally* significant). However, it is still the case that M-noise and O-noise have different effects with respect to the dynamics of cooperation; that is, the locus of noise is still important, as the LRV would have us believe, even though the influence of the noise on the subsequent evolutionary dynamics does *not* follow from the law.

6.3 The Issue of Generality

It should be emphasised once again that the experiments described throughout this paper are based on particular implementations of the IPD. Nevertheless, in this discussion, I would like to suggest that the results obtained *do* have some claim to a wider interpretation. In the final analysis, this claim can only be justified through further research. It is my belief that such research will be successful to the extent that the following three intuitions are sound; that there is *a*) some common ground between the IPD model and

other adaptive evolving systems in use in artificial life research, *b*) a broad scope to the principles of cybernetic theory, in particular the LRV, and *c*) a distribution of noise in a broad class of evolving systems.

In other words, the models developed here serve best as a bridge between artificial life and cybernetics from which some new theoretical and empirical territory, concerning the role(s) of noise, can be surveyed. Of course, in a more immediate sense this work also constitutes a contribution to the understanding of the IPD model, of interest in its own right.

6.4 Conclusions and Implications

Firstly, we have seen that the locus of variability is critical for understanding the dynamics of an evolving, adaptive system. Given the above caveat regarding generality, I believe that this has implications for the understanding of systems in which there is a distribution of variability across multiple loci, and in which agent-side complexity is considered in terms of a response to environmental complexity (in the spirit of [6]). For example, if one were to consider the “internal mechanism” of an agent to be constituted from various different behaviours and some arbitration mechanism (see, for example, [3]), then the distinctions between disturbance, response, system, and outcome are lost; after all, any given “behaviour” (such as “fleeing-from-predator”) incorporates all four of these aspects. And if these distinctions are lost, then the LRV perspective cannot be applied, and the evolutionary influence of environmental variability on the internal mechanism of an agent cannot be properly understood.

The second lesson is that whilst noise can indeed precipitate functional, adaptive change, it can also facilitate evolutionary mobility around genotype space. This is a novel role for noise in evolving systems, and goes beyond the two cybernetic interpretations of variability as either directly “threatening” to an agent, or as a source of information that can be transformed into an effective response. The role of noise in altering the structure of fitness landscapes to facilitate evolutionary search clearly invites further inquiry. Whether or not such a phenomenon is “useful” or not (either theoretically, or for design purposes) is, of course, still an open question.

	<i>Compulsory Model</i>	<i>IPD/CR Model</i>
Pop. Size	30	30
Iterations per game	60	100
Noise levels	1%	1%
Generations	10,000	5,000
Crossover Rate	0.95	0.95
All Mutations	0.001	0.005
Initial Memory	1	1
Cost on Complexity	0	0.003

Table 3: Parameters for the Compulsory and IPD/CR models (Appendix 1).

Acknowledgements

Thanks be to those of the CCNR, also to Phil Husbands and Hilary Buxton, for their generous and incisive criticism.

References

- [1] W.R. Ashby. *An Introduction to Cybernetics*. Chapman Hall, 1956.
- [2] R. Axelrod. *The Evolution of Cooperation*. New York : Basic Books, 1984.
- [3] R.A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47(1/3):139–159, 1991.
- [4] J. Dewey. *Experience and nature*. New York: Dover (rev. edn.), 1929.
- [5] J.A. Fletcher, M. Zwick, and M.A. Bedau. Dependence of adaptability on environmental structure. *Adaptive Behaviour*, 4(3-4):275–307, 1996.
- [6] P. Godfrey-Smith. *Complexity and the function of mind in nature*. Cambridge University Press, 1996.
- [7] F. Kasper and H. Schuster. Easily calculable measures for the complexity of spatio-temporal patterns. *Physical Review A*, 36(2):842–848, 1987.
- [8] F. Menczer and R.K. Belew. From complex environments to complex behaviours. *Adaptive Behaviour*, 4(3-4):309–355, 1996.
- [9] A.K. Seth. Interaction, uncertainty, and the evolution of complexity. In P. Husbands and I. Harvey, editors, *Proceedings of the 4th European Conference on Artificial Life (ECAL 97)*, pages 521–530. MIT Press, 1997.

- [10] A.K. Seth. The evolution of complexity and the value of variability. In C. Adami, R.K. Belew, H. Kitano, and C.E. Taylor, editors, *Proceedings of the 6th International conference on Artificial Life (ALIFE VI)*, pages 209–221. MIT Press, 1998.
- [11] A.K. Seth. Noise and the evolution of complexity; a study in evolutionary robotics. In P. Husbands and J.A. Meyer, editors, *Proceedings of the 1st European Workshop on Evolutionary Robotics (EVOROB 98)*. Springer Verlag, 1998.
- [12] C.E. Shannon and W. Weaver. *The mathematical theory of communication*. University of Illinois Press, 1949.
- [13] H. Simon. *The Sciences of the Artificial (3rd edition)*. MIT Press, 1988.
- [14] E.A. Stanley, D. Ashlock, and M.D. Smucker. Iterated prisoner’s dilemma with choice and refusal of partners: Evolutionary results. In F. Moran, A. Moreno, J.J. Merelo, and P. Chacon, editors, *Advances in Artificial Life : Lecture Notes in Artificial Intelligence*. Springer-Verlag, 1995.