Conditions for memetic driving[1]

Hendrik J. Blok[2]

Centre for Applied Ethics,
University of British Columbia, Canada

Abstract: Assuming culture is transmitted horizontally (via imitation) a model was constructed to determine the conditions under which culture can dominate genetic evolution ("get off the leash" according to Blackmore (1999)). Two requirements were found: (1) culture must compete with genes (required only for the effect to be empirically testable); and (2) Interactions between individuals must be confined to small groups or neighbourhoods. The model was tested via analysis and simulation.

1 Introduction

1.1 Memetics

The theory of memetics is a recent conjecture gaining credibility within the scientific community. The main tenet is that there exists a second replicating entity—besides the gene—that reproduces and evolves according to its own criteria. This replicating entity is the *meme*, loosely defined as an idea that may be transmitted via mimicry. The act of imitating behaviour becomes the mode whereby memes are replicated—a single instance of the behaviour is reproduced by another individual. Within this framework the individuals which carry out the act of copying memes play the role of their hosts.

The point of meme theory is that this new replicator serves its own purpose, not merely to enhance genetic fitness. Memes replicate via imitation, and those that are imitated more frequently are by definition more *fit*. In order to qualify as a distinct replicator (rather than just a complicated mechanism by which genetic selection acts), the fitness function of this replicator must not be identical to genetic fitness. That is, how frequently a meme is imitated must be (at least somewhat) independent of how well adapted (genetically) the host is.

We emphasize this point to distinguish meme theory from what we will call *gene-culture coevolutionary theory*. Laland *et al.* (1995) demonstrated that elements of culture, which may be transmitted by imitation, can interact with genetics to influence the proportions of alleles within a population, thereby impacting genetic evolution. However, in their model the cultural elements were passed on via *vertical transmission*, that is, from parent to child. Hence, the success of a cultural element depended on the reproductive success of the parent so there exists no independent memetic fitness. The fitness of the cultural unit is equivalent to the fitness of the host carrying it. Laland doesn't discuss cultural elements as selfish replicators so this is not meme theory. It is just the interests of the genes that need to be accounted for—standard evolutionary theory. We define meme

[1] Submitted to *J. Memetics - Evol. Models Inf. Transm.* (http://jom-emit.cfpm.org), 2002. Preprint available from http://rikblok.cjb.net/lib/blok(01).html.

[2] email: rikblok@ethics.ubc.ca or rikblok@shaw.ca

theory by the recognition that what is good for genes may or may not be good for memes, and vice versa. A memetic account must ask the question "What is good for the memes?" More importantly, proponents claim that it must be asked to acquire a complete picture of evolution. Our objective is to discover the conditions under which a purely genetic account would indeed be inadequate.

1.2 Meme-gene interaction

By differentiating memetic from genetic fitness, we emerge with a co-evolving system of two "species" where, depending on the details, the interaction may be cooperative, competitive, or neutral (no interaction). Susan Blackmore (1999) presents the argument that the interactions between memes and genes are responsible for many of our genetic traits, such as our excessively large brains. In other words, it is claimed that memes can direct the course of genetic evolution.

1.2.1 Cooperation

How genes and memes interact may take many forms, as do the interactions of different species in an ecology. It may be that certain memes interact with certain genes to enhance each others' survival and reproduction. In this case we describe the interaction as *cooperative*. This is precisely the scenario described in Laland *et al.* (1995) in which case the genetic and memetic fitness are identical.

Unfortunately, this also makes the theory of memes difficult to test empirically because it is hard to determine the degree of influence of the memes. Historical examples of this kind are explainable by simply asking "What is good for the genes?"—a purely genetic account—without having to appeal to meme theory. So, cooperative evolution is not a good candidate for scientific scrutiny of memetics.

1.2.2 Competition

Alternatively, it may be that what is beneficial to memes is deleterious to genes and vice versa, a *competitive* interaction. The orthodox sociobiological view appears to be that memes that compete with genes cannot survive and are weeded out (Lumsden & Wilson, 1981).

However, Blackmore proposes that altruism is an example of a competitive interaction: altruism typically has a genetic cost but it can make an individual popular and therefore more frequently imitated, enhancing the memes. Recently, Bull *et al.* (2000) argued that, based on simulation results, in a battle between memes and genes the memes would "win" (because they could evolve faster). This conflicts with an earlier paper (Best, 1999) which found that memes, despite their adaptability, typically "lose" to genes when the two compete. The goal of this paper is to test the conditions under which memes may be able to dominate genes.

2 The Model

To test the necessary the most general conditions required for memes to drive genes two models presented in the literature were considered and rejected.

2.1 The NKCS model

Bull *et al.* (2000) implemented a variation of Kauffman's *NK* model which is designed to study the effects of varying the degree *K* of interaction between *N* genes (Kauffman & Johnsen, 1991). Their variation introduced a second "species," the memeplex, which coevolved with the genome via codependent linkages between the two. A mutation in a single allele of either species affects the fitness of *K* others of the same species and *C* of the other. Whether the mutation is accepted is decided by determining if it increases the original species' fitness, with no regard for the other. Thus a beneficial mutation in one species may lower the fitness of the other.

The only distinction between the two species was the timescales the two processes occurred on—precisely the relationship we intend to study in this paper. It was discovered that when the memeplex evolved much faster than the genome it rapidly achieved high fitness at the expense of the genes' fitness thereby demonstrating that memes can inhibit genetic adaptation.

However, the *NKCS* model is unsuitable for our purposes for two reasons: firstly, the effect of changing the character of the interaction between memes and genes (cooperation or competition) is not ascertainable because there is exists no corresponding exogenous parameter. Secondly, the model neglects intra-species variability: each species is assumed to evolve as a whole so that mutations are either adopted or discarded by every individual simultaneously. We were concerned that heterogeneity (in both memes and genes) might be an important factor in the dynamics so the *NKCS* model was deemed unsuitable.

2.2 Best's model

Best (1999) presented an alternative model which also explored meme-gene coevolution, but via a genetic algorithm approach. In his approach the genes are trying to find a single high fitness configuration in a huge phase space of low fitness configurations. To enhance the search process a social learning component is introduced, transmitted by imitation. This model satisfies the requirement for a memetic theory because the social component has its own fitness requirement and Best considers both the case of the two replicators having identical and diametrically opposing goals.

Best's results seem to contradict those of Bull *et al.* in that he finds that memes can only guide genetic evolution when they cooperate. Memes cannot compete with genes, consistently losing except under extreme conditions.

Best also imposes a separation of timescales in his model; the memetic replication rate is orders of magnitudes faster than genetic replication. However, this model does not definitively decide the issue because the memetic mechanism is starkly different from the genetic. Genetic replication is handled via standard replicator dynamics with crossover applied to the genetic code. However, the imitation process is more complicated, involving "models" (teachers), learners, and a transmission force parameter. The results depend critically on the simulation design and parameters. It is unclear whether his results are due to the timescale separation or other technical details.

Instead of trying to expand these models further to answer our question we chose to design a new, simpler model.

2.3 Model description

The model we propose is individual-, or *agent*-, based; each of the many agents is a host to two types of replicators: a memeplex and a genome. For simplicity, the smallest unit of information, a single bit, is used to represent each. (Simulations have also been explored wherein each is a sequence of multiple bits, but they will not be explored here because they are not analytically tractable and do not yield any further insight.) To focus on the interaction between memes and genes all other sources of information, such as environmental factors, are disregarded.

The dynamics consist of two types of interactions: intra- and inter-agent.

2.3.1 Intra-agent interactions: Fitness

Within a host the meme and gene interact such that their tendency to replicate, or *fitness*, is influenced by both. That is, the success of a host's meme depends not only on itself but also on the host's gene and vice versa. The canonical representation of such an interaction is via a two-player game. Let us define the bits as C=1 (cooperate) and D=0 (defect). Then the meme/gene (ordered) pairs can have four possible combinations: *CC*, *CD*, *DC*, or *DD*, with a separate fitness for the meme and gene for each of these combinations.

To prevent introducing bias in the model we require that the game be symmetric, so there is no inherent advantage to either memes or genes. Thus the space of possible fitness combinations is reduced to 4 independent dimensions, denoted by T (temptation), R (reward), S (sucker), and P (punishment), as shown in Table 2.1.

		Gene	
		D	*C*
Meme	*D*	*P*, *P*	*S*, *T*
	C	*T*, *S*	*R*, *R*

Table 2.1 Symmetric payoff table for two-player game applied to meme-gene interaction.

2.3.1.1 Parameter-space reduction

This approach gives us a four parameter space (T,R,S,P) to explore, which is somewhat overwhelming. Fortunately, this can be reduced further. As will be shown below, the only important detail in the payoff parameters is their ranking in relation to each other so we can require them to be ordinal values. Neglecting equalities there are only twelve possible distinct games (Rapaport *et al.*, 1976) as listed in Table 2.2. (There are 24

permutations of the payoffs but half of these are equivalent to the other half under the transformation $C \leftrightarrow D$. For example, $R>T>P>S$ is equivalent to $P>S>R>T$.)

Index	*Relation*	*Game*
1	$R>T>P>S$	Assurance Game (AG)
2	$R>P>T>S$	
3	$R>P>S>T$	
4	$R>S>P>T$	
5	$R>S>T>P$	
6	$R>T>S>P$	Privileged Game (PG)
7	$T>R>S>P$	Chicken Game (CG)
8	$T>S>R>P$	
9	$T>S>P>R$	
10	$T>P>S>R$	
11	$T>P>R>S$	Altruist's Dilemma (AD)
12	$T>R>P>S$	Prisoner's Dilemma (PD)

Table 2.2 The complete set of ordinally distinct symmetric 2×2 games.

To summarize our progress thus far, we are considering a population of agents which are merely hosts to pairs of replicating entities, memes and genes. We are interested in the interaction between the two replicators so the simplest possible interaction was invoked, a symmetric two-player game. Since only ordinal ranking is relevant the game space could be reduced to twelve distinct games. Note that all agents are forced to play the same game internally, set by the modeler in order to test a particular class of interaction. In some games the memes and genes can cooperate to their mutual advantage while in others they compete. These interactions determine their ``fitness''which determines their potential to replicate.

2.3.2 Inter-agent interactions: Replication

The second type of interaction in the model is replication between agents. Each replication event can be either memetic or genetic. The processes are distinguished by their rates: the memes are assumed to replicate much faster than the genes. To model this, each process is taken to be Poisson in nature with a rate unity for genetic replications and a rate $\rho>1$ for memetic replications (so that, on average, there are ρ memetic replications per individual per generation). By varying ρ in the model we are then able to test the effect of timescale separation. Hence, this component of the model corresponds to simple replicator dynamics operating independently on two separate timescales.

The mechanism whereby genes and memes replicate is assumed to be mathematically equivalent even though the interpretation may differ. The process begins with the random selection of two agents for interaction. The relative fitness of the agents' replicators (memeplex or genome, as the case may be) are compared and the agent with the higher fitness *wins*, transmitting its replicator to the *loser*. In the case of a draw symmetry is broken by a fair coin toss to decide the winner. (It is assumed that genes and

memes replicate independently so that, for instance, a host's genome is unchanged during a memetic replication.)

The process described above is a reasonable idealization of asexual reproduction. For genetic replication we might interpret it as a *push*: the winning genome forcibly displaces the loser. For memetic replication, on the other hand, this is better interpreted as a *pull*: the losing agent prefers the winner's memeplex to its own and chooses to mimic it. Nevertheless, the mechanisms are identical by design.

The reader may object that the use of asexual genetic reproduction makes this model irrelevant to the study of humanity. It may indeed be more accurate to test sexual replication, with coupling and genetic crossover, but the goal of this project is to test the viability of memetic driving in the simplest scenario. First we should establish whether memes can drive genes even when the only difference between the two is the timescales and then we may extend the model by adding more realism.

Each replication also entails a risk of error during the transcription process—a random mutation. We incorporate this as another new model parameter μ, the mutation rate. μ indicates the chance that a replicating bit is flipped ($D \leftrightarrow C$) when it is copied. If $\mu=0$ then there are no transcription errors whereas if $\mu=1/2$ then each bit has an equal chance of being copied faithfully or erroneously—complete randomness. Thus we require two model parameters, $\rho>1$ and $0<\mu<1/2$, to describe inter-agent interactions.

In summary, the model consists of individual agents consisting of meme and gene bit-values, yielding the four varieties of agents listed in Table 2.3. Two varieties (*CC* and *DD*) are symmetric, with equal fitnesses for both memes and genes while two are asymmetric. (The labels given to the asymmetric varieties in the table are actually misnomers for games 3–5 since the cooperator achieves the higher payoff ($S>T$) but, as we shall see, they are insignificant in these games anyway.)

Meme	*Gene*	*Agent Variety*
C	*C*	Cooperator
C	*D*	Gene-dominant
D	*C*	Meme-dominant
D	*D*	Defector

Table 2.3 Four possible varieties of agents. The asymmetric (*CD* and *DC*) labels (gene-/meme-dominant) refer to the six competitive games (indices 7–12 of Table 2.2) in which *T*>*S* so the defecting replicator achieves the greater utility.

Together, the intra-agent meme-gene interactions and inter-agent replications generate the model dynamics. Table 2.4 demonstrates how they combine to yield transitions between agent varieties for the case of Chicken (competitive meme-gene interactions). To reconstruct the table, recall in Chicken the payoffs are ordered $T>R>S>P$ so from a memetic (genetic) perspective the varieties are ranked $DC>CC>CD>DD$ ($CD>CC>DC>DD$) and the higher ranking meme (gene) replaces the lower when two agents meet.

	DD	*DC*	*CD*	*CC*
DD	•	•	***C****D*+*CD*	***C****D*+*CC*
		*D**C***+*DC*	•	*D**C***+*CC*
DC	•	•	*DC*+***D****D*	*DC*+***D****C*
	DC+*D**C***		*D**D***+*CD*	•
CD	*CD*+***C****D*	***D****D*+*DC*	•	•
	•	*CD*+*D**D***		*CD*+*C**D***
CC	*CC*+***C****D*	***D****C*+*DC*	•	•
	CC+*D**C***	•	*C**D***+*CD*	

Table 2.4 Transitions on replicative interactions between two agents for Chicken-type meme-gene interactions. The upper/clear (lower/colored) line in each row indicates a memetic (genetic) replication event. Bullets indicate no change in either agent. Transcription errors are neglected.

The model as described is the simplest that we could construct which captured the following salient properties: (1) memes and genes are independent replicating entities with (2) a non-trivial interaction. The advantage of this simple approach is twofold. First, it allows the demonstration of some interesting results clearly. Each of the few ingredients are well understood so it is unlikely that any interesting dynamics which may emerge are artifacts of dubious mechanisms. Secondly, even with this degree of simplicity the model is only marginally analytic (Blok & Bergersen, 2002). Any further complicating factors would almost surely push it beyond the threshold of tractability.

3 Results

The extensive use of game theory above is not coincidental. Game theory also allows us to make some predictions of the model dynamics. This may seem odd since we are not modeling rational agents trying to optimize some utility but unthinking, mechanical replicators. Nevertheless, Skyrms (2000) has shown that game theoretic arguments often apply to evolutionary dynamics since the replication mechanism preferentially selects the fittest replicators.

3.1 Nash equilibria

The *Nash equilibrium*, which is the point in the payoff table (Table 2.1) such that neither player (replicator) could improve its payoff (fitness) by changing its strategy (bit), helps us predict the dominant variety of agents to arise from the dynamics. For the symmetric two-player game there may be one or two Nash equilibria for each combination of payoffs, as shown in Table 3.1.

Game Index	*Relation*	*Nash Equilibria*	*Class*
1–3	*R>T, P>S*	*CC* or *DD*	Assurance
4–6	*R>T, S>P*	*CC*	Privileged
7–9	*T>R, S>P*	*CD* or *DC*	Chicken
10–12	*T>R, P>S*	*DD*	Dilemma

Table 3.1 The four classes of games and their Nash equilibria. The game indices are from Table 2.2. The class names are borrowed from popular games they contain.

Note that the Privileged and Dilemma classes (games 4–6 and 10–12) only have one Nash equilibrium so it is inevitable that these varieties will be selected for in the replicator dynamics and eventually dominate the system. The other two classes are more interesting: in the Assurance class (games 1–3) both symmetric varieties of agents are Nash equilibria even though the *CC* is fitter than *DD* (*R>P*). The Nash equilibrium alone is insufficient to predict the outcome.

Similarly, in the Chicken class (games 7–9) both asymmetric varieties are Nash equilibria. In this case, since *T>S*, the *CD* (*DC*) variety will be genetically (memetically) fitter so the Chicken class provides an excellent framework for exploring competitive interactions between memes and genes.

The next sections will study both of these classes in more detail using analytic techniques and computer simulation[3]. The simulation algorithm is identical to the model as described except that instead of Poisson processes, time is iterated random sequentially, closely mimicking Poisson-distributed events (Schönfisch & de Roos, 1999).

3.2 Assurance class

The Assurance class of games (*R>T, P>S*) has two Nash equilibria, *DD* and *CC*, of which the latter is superior, allowing both replicators to achieve a higher fitness (*R>P*). Mean-field analysis (a method borrowed from statistical physics which neglects fluctuations and correlations) indicates that only the *CC* variety is stable (Blok & Bergersen, 2002). Thus, given a non-vanishing mutation rate, the system is always able to escape the suboptimal Nash equilibrium, as is confirmed by simulations (see Figure 3.1).

[3] A 32-bit MS-Windows executable simulation of the model may be found at http://rikblok.cjb.net/lib/blok(01).html, with full source code. The simulation requires R2DToo, a free, generic simulation tool available from the same site.

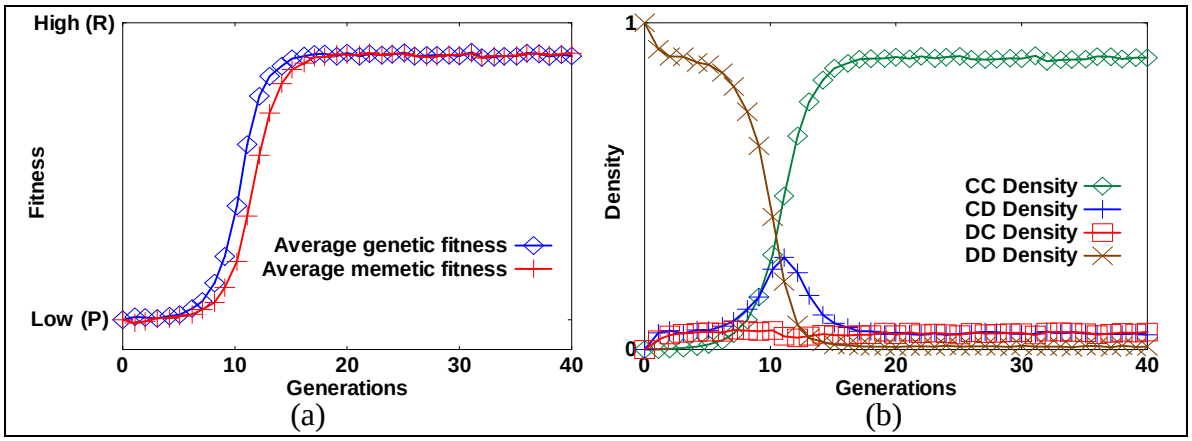

Figure 3.1 The suboptimal Nash equilibrium (*DD*) is unstable in simulations of the Assurance game, as demonstrated by replicator fitness (a) and agent variety densities (b). [Non-local, N= 64×64, μ=0.1, ρ=10, Game index=1, initial conditions: all *DD*.]

3.3 Chicken class

In all but the Chicken class of games the Nash equilibria are symmetric with the meme and gene sharing the same value (*C* or *D*). Thus they will both achieve the same fitness. However, in the Chicken class ($T>R$, $S>P$, including the game we will limit our attention to, Chicken) the Nash equilibria are asymmetric—only one replicator can achieve optimum fitness, at the expense of the other. Since $T>S$, the *CD* (*DC*) equilibrium would be to the advantage of the genes (memes). Thus, Chicken demonstrates competition between memes and genes.

Mean field analysis of Chicken indicates that the system is bistable: the dynamics will fixate on one or the other Nash equilibrium depending on initial conditions (Blok & Bergersen, 2002). In particular, if genes initially dominate (such that *CD* represents more than half the population) then they will do so forever; fitter memes will never be able to invade. This is confirmed by the simulations, as demonstrated in Figure 3.2 which shows that the steady-state population proportions depends on the initial conditions.

As an aside, note that the game of Chicken actually has three Nash equilibria, not two: if mixed strategies are allowed then the dominant strategy is to randomly choose *C* or *D* with equal probabilities. In an evolutionary setting with only pure strategies one expects this to emerge as a mixed population of both types (Skyrms, 2000). However, the mixed state is not observed in our model because the details of the interactions make it dynamically unstable. From Table 2.4 we can see that mixing between *CD*s and *DC*s tends to eliminate one or the other, driving the system away from the mixed equilibrium.

In these first set of runs we assume the population is "well-mixed" so that any host interacts with any other with equal probability. The agents exist in a *global* community, with no preference for interacting locally. This is the experimental equivalent of mean-field theory which ignores correlations between neighbors.

The important point here is that neither steady state can be invaded—both *CD* and *DC* are evolutionary stable strategies (within the constraints imposed). In particular, this means that *once genetic advantage has been established it cannot be invaded* by competitive memes if interactions are global and vice versa. (By genetic advantage we mean that the average genetic fitness is better than memetic.)

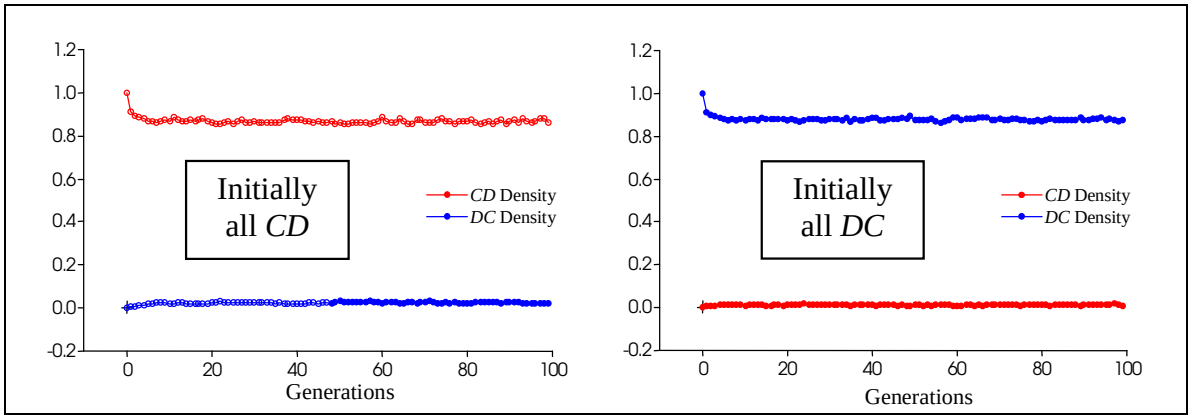

Figure 3.2 Experiments confirm that both mean-field fixed points are preserved in the Chicken game with the outcomes depending on initial conditions. [N=64×64, μ=0.1, ρ=10, Game index=7.] (Only the first 100 generations are plotted but no deviations were observed even on much longer runs.)

3.4 Localized Chicken

We now explore the robustness of our analytic and experimental results under the constraint of *locality*. Each agent is assigned a location and only allowed to interact with agents at nearby locations. (The agents do not move.) For the runs that follow we will focus on a two-dimensional Moore neighborhood: each agent occupies a vertex of a square, two-dimensional lattice and only interacts (reproductively) with its eight nearest neighbors (four directly opposite and four diagonal). (The results below generalize to other forms of locality, including 2-D von Neumann and one-dimensional neighborhoods.) Periodic boundary conditions are imposed to minimize edge effects. The important point is that beneficial memes and genes can only be transmitted locally, creating spatial correlations which the mean-field analysis does not take into account.

In most of the games considered no variation in the dynamics are observed indicating that spatial structure is irrelevant in these cases. The exception is the competitive scenario ($T>R$, $S>P$, including Chicken). In this case the gene-dominant fixed point of the nonspatial case destabilizes so that the only stable fixed point is the meme dominant *DC*. Even if genes dominate in the short term, they are eventually undermined by memes. With local interactions the *CD* configuration is no longer an evolutionary stable strategy, it can be invaded by individuals with worse genetic fitness (but better memetic fitness). An example of such an invasion is shown in Figure 3.3. Thus Darwin's law of "Survival of the fittest" appears to be violated when only genetic factors are accounted for.

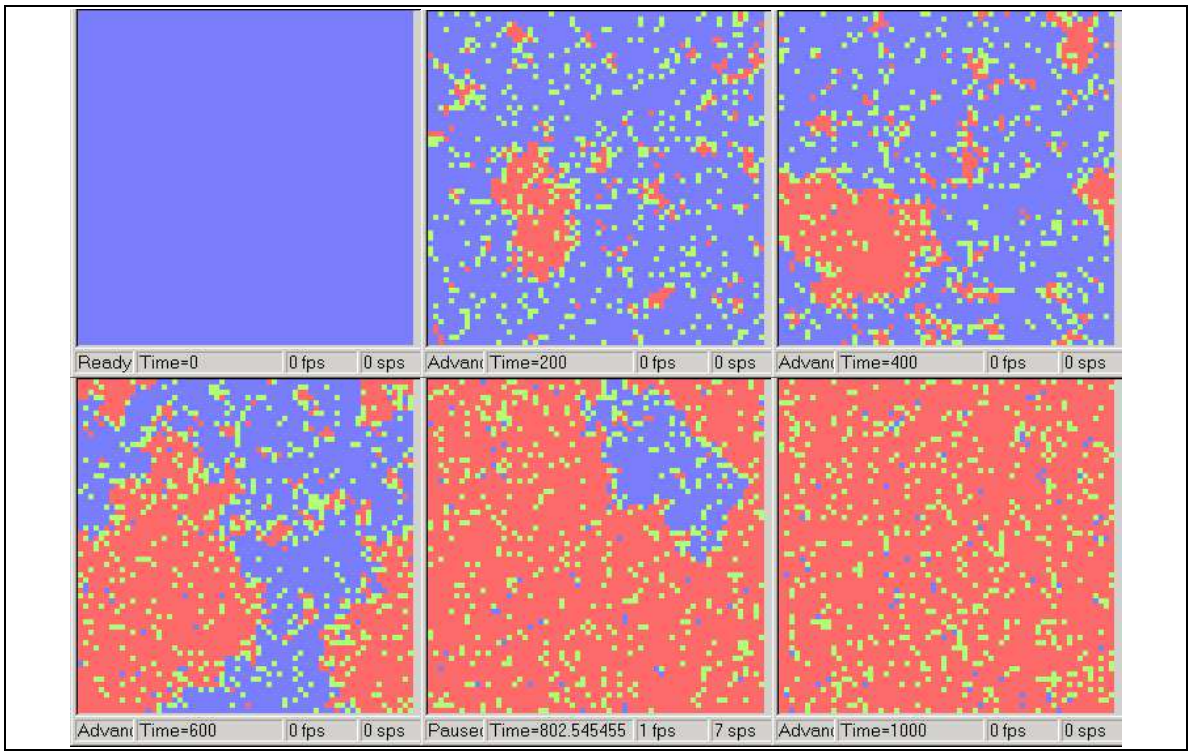

Figure 3.3 A typical run on a two-dimensional Moore neighbourhood with competitive interactions. Snapshots taken at generation 0 (top left) through generation 1,000 (bottom right), left to right. Even though the system is initialized with all *CD* (blue) it is steadily invaded by mutant *DC*s (red). [N=64×64, μ=0.1, ρ=10, Game index=7.] (Blue=*CD*, Green=*CC/DD*, Red=*DC*)

The mechanism whereby meme-dominant individuals can invade a gene-dominant population can be understood by considering small systems of N individuals. The analysis carried out in a companion paper (Blok & Bergersen, 2002) derives an *effective potential* curve which can be metaphorically viewed as a landscape with two valleys separated by a hill (see Figure 3.4). The system macro-state is then analogous to a marble which rolls down any local slope to the bottom of the valley. Random mutations buffet the marble, causing it to jitter, and occasionally the disturbance is large enough to send it over the top of the hill and into the other valley. The deeper valley, corresponding to the meme-dominant configuration is more stable since it requires a much larger disturbance for the marble to escape.

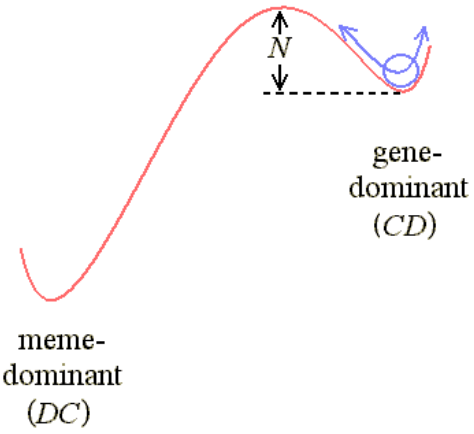

Figure 3.4 The *effective potential* landscape serves as a metaphor to explain the dynamics: the system behaves as a marble, jostled around by random noise, which seeks the valleys of the curve. Occasionally the bumps are large enough to drive it over the hump separating the two valleys into a new equilibrium.

Note that the depth of the gene-dominant valley is proportional to the number of agents *N*. Thus, as *N* grows the depth of the valley grows and so does the time required for the marble to escape. As *N* gets very large, such as in the nonspatial simulations in this paper, the system is effectively trapped forever within the gene-dominant valley as demonstrated in Figure 3.2.

However, the spatial case is different: since each site only interacts with a few neighbours each is effectively a small system coupled to other small systems. Since the local system size is small it doesn't take long for mutations to drive the marble over the hump and into the meme-dominant regime. This seed, in turn, provides an impetus for systems coupled to it to also climb up the hill and increases their likelihood of crossing over. Hence, the meme-dominant strategy invades the population by spreading from an initial random mutation, as shown in Figure 3.3.

The conclusion we can draw from this is that *competitive memes can undermine optimal genes if interactions are local*, in stark contrast to the global result and other work (Best, 1999). This is a somewhat surprising result given that within the model memes and genes differ in their rates of reproduction alone.

4 Summary

In this paper we constructed an agent-based model of genetic-memetic interaction via a game theoretic approach. For simplicity memes and genes were both constructed as selfish replicators, differentiated only by their rate of replication (memes were assumed to replicate more often). The goal was to determine the minimum set of requirements for memes to have an impact on genetic evolution. We discovered two necessary and sufficient conditions: (1) individuals must interact locally, within small neighbourhoods rather than mixing with the population at large; and (2) for the impact to measurable memes and genes must compete or the effect cannot be distinguished from a purely

genetic process. The main lesson is that a purely genetic perspective of evolution may not account for all adaptations and a memetic account may be required in some cases.

The possible extensions of the model are boundless but a few will be discussed here. The next logical step would be to remove the symmetry constraint imposed on memes and genes. One could consider heterogeneous mutation rates for the two replicators, or represent them via more complex data structures. Also, it would be interesting to explore the introduction of sexual genetic reproduction which may stabilize the genome by protecting against deleterious mutations (West *et al.*, 1999). Finally, we leave with the question: Is it necessary that both memetic *and* genetic interactions be local in order for memes to "get off the leash?"

This work was supported by a grant from the National Sciences and Engineering Research Council of Canada. Many thanks to Peter Danielson, Birger Bergersen and Michael Baumann for helpful discussions.

References

Best, M. L. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior* 7:289–306.

Blackmore, S. (1999). *The Meme Machine*. Oxford University Press.

Blok, H. J. and Bergersen B. (2002) Fast processes can drive slow. *Work in progress, preprint available from* http://rikblok.cjb.net/lib/blok(02).html.

Bull, L., Holland, O., and Blackmore, S. (2000) On meme-gene coevolution. *Artificial Life* 6:227–235.

Kauffman, S. A. and Johnsen, S. (1991). Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches. *J. Theor. Biol.* 149:467–505.

Laland, K. N., Kumm, J., and Feldman, M. W. (1995) Gene-culture coevolutionary theory: A test case. *Current Anthropology* 36:131–156.

Lumsden, C. J. and Wilson, E. O. (1981) *Genes, Mind, and Culture*. Harvard University Press.

Rapaport, A, Guyer, M. J., and Gordon, D. G. (1976) *The 2×2 Game*. University of Michigan Press.

Schönfisch, B. and de Roos, A. (1999). Synchronous and asynchronous updating in cellular automata. *BioSystems* 51:123–143.

Skyrms, B (2000). Game theory, rationality and evolution of the social contract. *J. Conscious. Studies* 7:269–84.

West, A. A., Lively, C. M., and Read, A. F. (1999). A pluralist approach to sex and recombination. *J. Evol. Biol.* 12:1003–1012.