

Global Adaptation in Networks of Selfish Components: Emergent Associative Memory at the System Scale.

Richard A. Watson^a, Rob Mills^a, C.L. Buckley^b

^aNatural Systems group, Electronics and Computer Science, University of Southampton, SO17 1BJ, U.K..
Tel. 44 (0)23 8059 2690. Fax. +44 (0)2380 59 9179. email: raw@ecs.soton.ac.uk. (corresponding author)

^bInformatics, Sussex University, BN1 9RH, U.K.

ABSTRACT

In some circumstances complex adaptive systems composed of numerous self-interested agents can self-organise into structures that enhance global adaptation, efficiency or function. However, the general conditions for such an outcome are poorly understood and present a fundamental open question for domains as varied as ecology, sociology, economics, organismic biology and technological infrastructure design. In contrast, sufficient conditions for artificial neural networks to form structures that perform collective computational processes such as associative memory/recall, classification, generalisation and optimisation, are well-understood. Such global functions within a single agent or organism are not wholly surprising since the mechanisms (e.g. Hebbian learning) that create these neural organisations may be selected for this purpose, but agents in a multi-agent system have no obvious reason to adhere to such a structuring protocol or produce such global behaviours when acting from individual self-interest. However, Hebbian learning is actually a very simple and fully-distributed habituation or positive feedback principle. Here we show that when self-interested agents can modify how they are affected by other agents (e.g. when they can influence which other agents they interact with) then, in adapting these inter-agent relationships to maximise their own utility, they will necessarily alter them in a manner homologous with Hebbian learning. Multi-agent systems with adaptable relationships will thereby exhibit the same system-level behaviours as neural networks under Hebbian learning. For example, improved global efficiency in multi-agent systems can be explained by the inherent ability of associative memory to generalise by idealising stored patterns and/or creating new combinations of sub-patterns. Thus distributed multi-agent systems can spontaneously exhibit adaptive global behaviours in the same sense, and by the same mechanism, as the organisational principles familiar in connectionist models of organismic learning.

Keywords: self-organisation, adaptive networks, Hebbian learning, multi-agent systems, social networks, emergent computation, games on networks.

Selfish changes to connections and global adaptation

One of the key open questions in the field of artificial life and adaptive systems is how, if at all, it is possible that a complex system that is not evolved can exhibit adaptation or increased functionality without design. Self-organisation [33] has the potential to play a role in the formation of pre-biotic organisations, or in moving from one level of biological organisation to another [42], for example, but theory to understand exactly what this means or how it might work is limited. In contrast, theory to understand distributed adaptive processes in neural networks is well-developed, but generally assumed to be relevant only to brains and nervous systems ‘programmed’ to exhibit such adaptation. Here we show that organisational principles familiar in learning neural networks emerge spontaneously in distributed networks of self-interested agents. This enables us to utilise the theoretical framework of neural networks to understand how self-organisation produces global adaptation in complex adaptive systems.

The behaviour that is rational for a self-interested agent (e.g. a person in a social network, or a species in an ecosystem) is generally sensitive to its environment or context – i.e. the best response for an agent often depends on the strategy adopted by the other agents it interacts with [41]. How the strategy of one agent affects another, and then feeds back to influence the first, can be significantly affected by the topology of a network, and different interaction structures can create conditions that are more or less conducive to high-levels of cooperation or system efficiency in a multi-agent system [51,52,86,69,72]. Often, although an agent may not be able to control the strategy of others, it may have some influence over which agents it interacts with [75,15,2], and by altering this an agent can produce contexts that subsequently change its own social behaviour or function [54,59,62,55]. Accordingly, many types of agents do not merely change their own behaviours in response to context but in fact engineer their context by changing who they interact with, how frequently they interact with them, and how much they depend on one agent or another. Network models where agents can alter the topology of the network on which they interact therefore provide a flexible way to model the self-organisation of interaction structure [33,55,74,56,75,73,31,29,30] and have been used to illustrate examples where this enhances cooperation [55,72]. More generally, the conditions where multi-agent systems can produce self-organised structure that enhances system-level or holistic function are not well understood. In this paper we wish to better understand the relationship between self-organisation in interaction networks amongst selfish agents and the potential for enhanced global adaptation [33,17,55,92,62,91,72].

Although both the structure of the network and the behaviours on the network are dynamic variables, a conceptual separation between agent behaviours that affect the structure of interactions from other ‘ordinary’ agent behaviours (e.g. whether an agent plays defect or cooperate strategy) is natural in many types of systems where there is a separation of timescales. In adaptive networks [11] it is common to conceptually separate the dynamical behaviour of agents on a network from the structure of the network even though agents can often (slowly) alter the network structure. For example, ecological population dynamics are controlled by a relatively slow-changing network of evolved inter-species relationships [60,65], selection on individual genetic traits is altered by a relatively slow-changing network of evolved pleiotropic and epistatic interactions [76,57]; and social behaviours are affected by a relatively slow-changing network of social connections [55,74,56]. Understanding how behaviours on networks promote changes in the structure of the network and, in turn, how the modified structure of the network affects global function and efficiency [51,52], is important for domains as varied as social networks [86,87], the evolution of evolvability [76], the evolution of altruism [88], technological infrastructures [67], epidemiology [34,50,89] and ecosystem complexity [90,70,20]. However, the strategy with which agents alter network structure may be domain specific; which makes a general theory to understand the coupling between behaviours on networks and changes to network structure problematic.

In this paper we adopt a simple and general assumption; that agents modify relationships with one another (i.e. connections in a network) *selfishly*. That is, the changes to network structure are carried out by the agents on the network rationally with respect to their individual utility [55,75,30,31]. By this we simply mean that behaviours of agents on a network and behaviours of agents that *alter* that network should have the same motivations – i.e. be controlled by maximisation of the same individual utility function. Specifically, we assume, as is conventional, that the individual utility of an agent is a function of its behaviour and its connections to other agents, and that agents can adopt a behaviour to maximise their individual utility given the current network structure and the behaviours of others (a ‘best response’ strategy). But we also assume that they have the ability to change their network connections and will do so in a manner that increases the utility of their current behaviour (a best response strategy applied to restructuring behaviours). Hence, we refer to these as ‘selfish’ changes to network structure. (In this paper we address continuous changes to network structure, i.e. changes in the strength and sign of connections, rather than a topological ‘re-wiring’ of the network [55,75,31,29,30]).

This is a very general assumption. For example, social agents that seek to minimise conflict/maximise consensus with friends may change their behaviour to be compatible with the behaviours of friends but logically they may also change friends to strengthen relationships with others who are already compatible (and weaken connections with friends who conflict) [69]. Similarly, evolving species in an ecosystem will evolve traits that maximise their fitness given the other species they interact with in their biotic context, but natural selection also demands that, where possible, they will evolve behaviours that alter the set of species they interact with, so as to maximise their fitness given their current traits. In an economic market, businesses may adopt pricing plans that maximise profit given the other businesses they currently trade with, but a profit-seeking business may also change the network of businesses they trade with to maximise profit from their current pricing plan [69]. In all cases, we do not need to adopt an arbitrary mapping between behaviours on the network and behaviours that change the network, instead we are simply assuming that agents will change network structure in a manner that is consistent with the same factors that govern their behaviours on the network – be it minimising conflict, or maximising fitness, profit or productivity. All such scenarios can be modelled abstractly as maximisation of utility.

We will assess how selfish changes to inter-agent relationships change the dynamics of the system, and specifically, whether they cause self-interested agents given these new relationships to subsequently adopt more cooperative behaviours, or conversely, result in more inter-agent conflict and lower levels of cooperation. The fundamental question of cooperation is whether agents that are acting to maximise individual utility have the result of increasing the utility of others. This is assessed by measuring total welfare or *total utility* which is simply the sum of individual utilities (see Methods). Using a sum of individual utilities in this manner in no way assumes that individual and global utilities are aligned nor that individual utilities are not conflicted – an increase in utility of one individual may cause a decrease in utility of another. Indeed, these constraints are what make total utility a quantity worth measuring. For example, the Prisoner’s dilemma [41,72] is interesting precisely because individuals seeking to maximise individual utility will arrive at an equilibrium that does not maximise the sum of individual utilities. Likewise, in systems where agents can benefit from a public good without contributing to it, a ‘tragedy of the commons’ can result where individual (relative) gains result in global (absolute) losses [41]. However, preoccupation with the Prisoner’s Dilemma – where the only Nash equilibrium of individual strategies *minimises* the sum of individual utilities – may lead one to assume that individual and global interests must be completely opposed in this sense in every two-player game in order for cooperation to be meaningful. But this is not the case. Even in two-player games cooperation can be meaningful without the Prisoner’s Dilemma [69,61,64], and more generally, even if the underlying two-player games present no obvious conflict of interest, games of multiple players can present unresolvable conflicts. In particular, multi-player non-zero-sum games may have many equilibria, and different equilibria may have different total utilities; Consequently, whereas all equilibria locally maximise the individual utility of all agents (by definition), different equilibria vary in the extent to which they are successful in maximising total utility. As the system approaches a given local equilibrium it may well be the case that total utility also increases, i.e. individual and total utility are not completely opposed. But a given local equilibrium will not in general globally maximise total utility, so neither are individual and total utility completely aligned. Such partial alignment between individual and total utility arises naturally given that individuals have conflicting or constrained interests (unless they are specifically contrived to have completely opposed interests as in the Prisoner’s Dilemma). Conflicts of interest between agents are created in our model by a network of random pairwise constraints (coordination and anti-coordination games [73]) that cannot be simultaneously satisfied. This representation

cannot represent all possible games, but it does represent tensions between individual and global utility, creating a system which exhibits multiple Nash equilibria some of which have higher total utility than others. We show that the problem of maximising total utility is equivalent to a combinatorial optimisation problem that is difficult in a formal sense.

In any model where the game is static, however, the degree to which individually selfish actions align with global welfare is simply an outcome intrinsic to the setup of the game employed. But when agents have the ability to alter who they play with, or the relative strength of their relationships with other agents, this is equivalent to changing the underlying game [72] (or changing how individuals perceive or experience the game [72,7]). Our question is not therefore, whether individual interests are aligned with global interests in the initial game we define, but rather whether the selfish restructuring of the interactions in the system leads, in effect, to agents that (because they are playing a modification of the underlying game) exhibit better alignment with total utility over time. We thus define global adaptation as an *increase* in the ability of the system to maximise total utility.

In summary, we examine systems under the following conditions: 1) Changes to agent behaviours and to connections between agents are carried out selfishly by the agents in the system to maximise individual utility; 2) A separation of timescales exists between fast changes to behaviours and slow changes to connections; 3) The initial dynamics of the agent behaviours (given the initial connections between them) exhibits multiple point attractors (Nash equilibria, [49]). 4) The agent behaviours are subject to external perturbations such that they sample many different attractors (on a timescale where connections change slowly).

The first two conditions are straightforward, and the third, although restrictive, is a useful simplifying assumption (we will show that if this condition is true of the original system it will remain true of the system as relationships change, and we will briefly discuss relaxation of this assumption later). But it is easy to see that these three assumptions alone do not necessarily produce any interesting consequences for global adaptation. That is, if the behavioural states change much more rapidly than connections then almost all changes to connections will occur after the system has reached whatever behavioural attractor it is going to reach. Selfish changes to connections (as we will show) have the effect of reinforcing the current state configuration (because they simply favour exploitation of local utility gradients) – making it more stable such that the agents involved are less inclined to change their behaviours. Thus, the long term dynamical consequences of such changes for global adaptation are degenerate [60]: Whichever local equilibrium happens to be found first becomes essentially canalised by changes to connections.

But we will show that together with the fourth assumption (the condition of external perturbations) selfish changes to connections have a much more interesting outcome. Under this condition the initial structure does not dictate a single attractor state but a distribution of attractor states, and this distribution of attractors changes as the connections of the system are modified. Specifically, we find that selfish agents change the structure of interactions in the system such that the equilibria of the new system are a special subset of those in the original system, specifically those that have high total utility. We can therefore view this restructuring as effecting equilibrium selection [14,5] that enhances the global level of adaptation (rather than merely finding an arbitrary local equilibrium). Accordingly, under these conditions, individually selfish changes to relationships have the effect of increasing global adaptation.

Selfish changes to connections and Hebbian learning

Our objective in this paper is not merely to show that selfish changes to connections can increase global adaptation under certain conditions but to provide a formal framework to understand why. Specifically, the key to understanding this result is that selfish agents necessarily modify connections in a manner consistent with Hebb's rule [16] – a simple learning rule familiar in computational neuroscience (Methods). This means that a system of selfish agents, each modifying its connections with other agents selfishly and in a completely distributed manner, will produce dynamical consequences for the system as a whole that are functionally identical to a learning neural network. This result seems initially surprising – there is no a priori reason why selfish agents should change connections with other agents in a Hebbian manner. In a cognitive system, in contrast, we suppose that natural selection is responsible for creating a Hebbian learning mechanism – but here we claim that in an ecosystem or multi-agent system of selfish agents, Hebbian learning at the system level arises spontaneously as a result of rational agent choices at the individual level. In fact, Hebb's rule is a very simple positive feedback principle and in light of this it is not so surprising that selfish components in a complex adaptive system will naturally adhere to such a principle, as we will discuss (Methods). Having shown this equivalence we can then employ concepts and theory from computational neuroscience to help us understand what consequences this will have for global behaviour and how global adaptation is enhanced [66,22,23,11,19,25,1,36,6].

In particular, recent work of our own [78,81] provides a foundation for the current paper by integrating two well-known behaviours of the Hopfield network [22] in a novel combination. Specifically, the ability of the Hopfield network with Hebbian learning to learn a set of training patterns and recall them from noisy or partial samples is well-known [22]. And in previously unrelated work, the Hopfield network can also be used as an optimisation method; here the weights of the network are not learned but pre-programmed to represent the constraints of an optimisation problem, and running the network then provides locally optimal solutions to the optimisation problem [23,24]. Our work shows that these two behaviours can be combined in the same network at different timescales to improve optimisation. Specifically, a network's initial weights are programmed to represent an optimisation problem as before, but the system is repeatedly perturbed, causing it to visit many different locally optimal solutions. Whilst this is happening the weights of the network are slowly updated by Hebbian learning. If the system spends most of its time at local optima (i.e. the system perturbations are infrequent and dynamical

transients are relatively short), the Hebbian learning will tend to enlarge the basin of attraction for existing local attractors in proportion to how often they are visited. In effect, the network is overlaying an associative memory [22] of its behavioural experience onto its own original dynamics, or ‘self-modelling’ [78].

Since the enlargement of one attractor must result in the decrease of another, attractors compete with one another during this process, and the system tends to enlarge the attractors that it visits most often at the expense of those that are visited less often. As its behaviour becomes dominated by learned connections rather than the original connections, positive feedback accelerates this competition between incompatible attractor patterns. The surprising finding is that this competition results in the enlargement of basins of attraction that lead to especially high total-utility attractors at the expense of low-utility attractors. This is explained in part by the fact that in systems built out of the superposition of many low-order constraints, low-energy (high-utility) attractors necessarily have large basins of attraction [11,38,37,39]. So, the better the attractor the more it is visited, thus the more it is enlarged by learning and the more it is visited in future, and so on. Thus, if the learning rate is sufficiently low, the system will with high likelihood enlarge the attractor of the globally optimal system configuration. More interestingly, generalisation inherent in associative memory [36,6] causes the system to generate new combinations of common features from past configurations, enlarging the basin of attraction for high-utility attractors that have not yet been experienced. Thus the system does not merely recall good configurations that have already been visited but can find good configurations more quickly even for the first time [78,81].

The novel contribution of the current paper is to show that these behaviours, shown straightforwardly in a neural network where we may assume Hebbian learning, also occur spontaneously in a complex adaptive system of self-interested components that alter connections selfishly (without assuming or mandating Hebbian learning). We argue that these behaviours are natural for any complex adaptive system with plastic connections that are altered by selfish agents, because selfish agents in such systems will necessarily alter connections in a Hebbian manner. We discuss the implications of this finding for various types of complex adaptive systems, thus transferring the optimisation result from neural networks into the domain of multi-agent systems of selfish components. Thus the principles of associative memory and distributed learning usually reserved for organismic neural networks, i.e. brains and nervous systems, apply directly to self-organised complex adaptive systems and multi-agent systems. This *associative induction* does not in any way require that the system as a whole has any motivations to adapt holistically nor does it require that the network has been holistically evolved to produce such adaptation. The behaviours that we observe are the spontaneous, self-organised result of many essentially uncoordinated agents acting selfishly. Nonetheless, we show that their independent motivation to optimise individual utility causes them to selectively strengthen or weaken connections with other agents such that the system as a whole becomes more likely to find combinations of behaviours that resolve conflicts and result in higher total utility in a non-trivial and quantifiable sense.

Methods

Behaviours, connections, and total utility.

Each agent in a network adopts one of two discrete behaviours or states, $s_i = \pm 1$ (e.g. tall/short, a/A, trade-x/trade-y, vote-A/vote-B), so as to maximise an individual pay-off, fitness or utility, u_i , which is a weighted sum of interactions between the state it adopts and the states of other agents in the network,

$$u_i = \sum_j^N \omega_{ij} s_i s_j . \quad (1)$$

where $\omega_{ij} = (-1,1)$ is the current interaction coefficient or connection strength between agent i and agent j . In principle, an agent may have interactions with all other agents, but non-connected agents may be represented by $\omega_{ij} = 0$. A behavioural dynamics of this form provides a general model for a dynamical system in many domains from neural networks, to spin-glass models, to ecosystems. Since an agent’s utility depends on many connections, not all connections can be satisfied simultaneously in all cases. An agent will thus adopt the behaviour that is the best compromise for it, given the states of other agents. For each connection we can say that the behavioural ‘constraint’ that the connection imposes is ‘satisfied’ if $\omega_{ij} s_i s_j > 0$, otherwise it remains unsatisfied or ‘in tension’.

In some scenarios, e.g. social networks, an agent may be free to modify interactions directly (i.e. make or break friendships). In other scenarios the network of interactions that are alterable by an agent will augment or supplement a network of interactions that are immutable or ‘external’ to the system. For example, epistasis between alleles is partly the result of external selective dependencies and partly the result of evolvable developmental details, and fitness dependencies between species in an ecosystem are likewise partly external and partly the result of evolvable characteristics of each species. We thus suppose that ω_{ij} is composed of two components, $\omega_{ij}(t) = \omega_{ij}^o + \omega_{ij}^l(t)$, where ω_{ij}^o represents the original dependencies between components of the system that are external and unalterable by the agents themselves, and ω_{ij}^l represents the component of agent interactions that are learnable, malleable or evolvable via agent adaptation [81]. For

simplicity we may suppose that the original behaviour of the system is determined entirely by ω_{ij}^O with $\omega_{ij}^L(t=0)=0$. As the agents in the network modify ω_{ij}^L this will alter the dynamics of the system.

As each agent chooses the state that maximises its own utility this may increase or decrease the utility of other agents in the network. We are interested in how this affects the *total system utility* or global pay-off of the system, which is simply the sum of individual utilities, measured using ω_{ij}^O :

$$U = \sum_i^N \sum_j^N \omega_{ij}^O s_i s_j. \quad (2).$$

This quantifies the degree to which this combination of behaviours satisfies the original behavioural constraints between all agents. In general, the collective effect of many independently-maximising agents will not produce combinations of states that find the maximum possible total utility, but merely locally optimal equilibria.

Selfish changes to connections.

When an agent modifies its connections with other agents selfishly this simply means that it does so in a manner that maximises its own utility. Formally, given Eq (1):

$$\begin{aligned} u'_i > u_i &\Rightarrow \sum_j^N \omega'_{ij} s_i s_j > \sum_j^N \omega_{ij} s_i s_j \\ &\Rightarrow \begin{cases} \omega'_{ij} > \omega_{ij} & \text{if } s_i s_j > 0, \\ \omega'_{ij} < \omega_{ij} & \text{if } s_i s_j < 0. \end{cases} \quad \text{or} \end{aligned} \quad (3)$$

Intuitively, this means that an agent strengthens connections that are currently satisfied and weakens connections that are currently in tension. Note that although the change in the magnitude or strength of a connection is dependent on the product $\omega_{ij} s_i s_j$, the sign of the absolute change in the connection, $\Delta\omega_{ij} = \omega'_{ij} - \omega_{ij}$, is not sensitive to the sign of ω_{ij} , but only to the signs of the agents' behaviours. This means that the effect of selfish changes to connections is to reinforce the current combination of behaviours, regardless of whether the current combination of behaviours satisfies the constraint between them. A general form of Hebb's rule [16] is often stated as $\Delta\omega_{ij} = \delta s_i s_j$ with $\delta > 0$ for Hebbian learning, and $\delta < 0$ for anti-

Hebbian learning [25,1], where δ is a constant controlling the learning rate. It is easy to see from the above that each and every change to connections that increases individual utility (for a given behavioural configuration) will necessarily change connections in a Hebbian direction regardless of the properties of the existing connection matrix. Accordingly, changes to connections will be determined by whether states are more often correlated or anti-correlated, not by the existing interaction coefficient between them. Although there is positive feedback between a connection and the correlation of states it connects, any one connection only partially controls the correlation of these states since each agent also has many other connections to satisfy. Importantly, this means that the sign of a connection may thus be inverted by selfish modifications if the agents' behaviours often disagree with the sign of the connection.

Connection Dynamics

For a deterministic model of connection dynamics as considered here (e.g. a rational agent model), we may suppose that both $\Delta\omega_{ij} = +r$ and $\Delta\omega_{ij} = -r$, where r is a constant controlling the rate with which weights are changed, are systematically considered and whichever case increases individual utility is applied. In other cases it may be natural to assume that the probability of accepting or rejecting a particular change in ω (generated from a distribution of $\Delta\omega$ values with mean zero), is proportional to the increase in utility, $\Delta u > 0$, that it confers, i.e. $P[\omega_{ij} \leftarrow \omega_{ij} + \Delta\omega_{ij}] = \delta\Delta u = \delta(\Delta\omega_{ij} s_i s_j)$. For example, if each agent is evolving by natural selection given random variation and non-deterministic selection in its connections to others then the probability of a beneficial change fixing in the population when it arises is proportional to the fitness benefit it confers [13]. Thus both the magnitude and sign of $\Delta\omega_{ij}$ under selfish changes will have an expected value proportional to $s_i s_j$, i.e. $E[\Delta\omega_{ij}] = \delta s_i s_j$, as per Hebb's rule.

The magnitude of a weight change under Hebbian learning, $|\delta s_i s_j|$, will be a constant in the discrete-state system, hence $|\Delta\omega_{ij}| = r$, as above, matches Hebbian magnitudes as well as signs in the deterministic system as well as probabilistic. In the following experiments all weights are also capped at a magnitude of 1 by a linear threshold function θ , i.e. if $x > 1$ then $\theta(x) = 1$, if $x < -1$ then $\theta(x) = -1$, else $\theta(x) = x$. Thus,

$$\omega_{ij}(t) = \theta[\omega_{ij}^O + \omega_{ij}^L(t)], \text{ where}$$

$$\omega_{ij}^L(t+1) = \omega_{ij}^L(t) + r s_i(t) s_j(t).$$

Note that for a given state configuration the pressures on changes to connections are symmetric, i.e. $\Delta\omega_{ij} = \Delta\omega_{ji}$, even if the original connections are asymmetric or different signs. This implies that over time, given the positive feedback on state configurations created by changes to connections, ω_{ij} and ω_{ji} will eventually take the same sign even if their initial values do not. This ensures that this type of modification to connections will tend to simplify the dynamics of systems over time (i.e. make fixed point attractors more likely) rather than cause them to become chaotic.

Perturbations and relaxations

The state of the system takes a random state configuration, $R = \{-1|1\}^N$, every τ time steps, where τ is sufficiently large to ensure that the system has spent most of its time at an attractor, but sufficiently small compared to r such that many attractors are visited in the time that significant modifications to connections takes place. We refer to each period of τ time steps as a

‘relaxation’ of the system. For the following experiments we use $\tau = \lceil (3/4)eN\ln(N) \rceil$ as a heuristic to allow sufficient time for the network to relax (derived from [48]). These perturbations show that no inheritance of system configurations is required – as one might have expected from either a stochastic local search process or simulated annealing [48,35], or a group selection model [92] – the only values inherited from one relaxation to the next are the connection values. Smaller perturbations will not fully explore the space of system configurations but will still provide opportunities to learn new associations so long as different attractors are visited. In evolutionary scenarios, changes in environmental conditions [57], for example, by migration to different spatial localities [83], may preclude the need for such perturbations by, in effect, enabling different attractors to be sampled in parallel rather than in series.

Given sufficiently short transients compared to the length of each relaxation, the cumulative effect of modifications to connections over a single relaxation of the network will be approximately equal to updating the connections (with a larger r) only at the attractor. We use this end-of-relaxation update method in the following experiments, but a continuous update model may be more natural in some systems, e.g. when attractors are not point attractors.

Initial connection matrices

We examine two simple test cases that are at opposite extremes of a structure dimension: 1) a network with initially unstructured connections, 2) a network with initially modular consistent connections. To ensure contractive dynamics, and to utilise formal treatments [78,81], we restrict our attention to symmetric connections, i.e. $\omega_{ij} = \omega_{ji}$ [22] (this ensures that the total system is a potential game [7]). Note that the results presented above, i.e. that $\Delta\omega_{ij}$ will be Hebbian, and $\Delta\omega_{ij} = \Delta\omega_{ji}$, do not depend on this symmetry. Symmetric interactions are not suitable to represent all types of interactions¹ (e.g. predator-prey interactions, which will generally create cyclic or sometimes chaotic attractors). But positive symmetric weights create a balanced coordination game (see also the stag hunt [69]), and negative weights define an anti-coordination game, and we address example systems where mixed constraints (S1) or differently weighted constraints (S2) make maximisation of total utility formally difficult; i.e. ω_{ij}^O creates numerous local optima in total utility such that selfish maximisation of individual utilities arrives at sub-optimal equilibria (see Results). Our set-up is thereby formally equivalent to scenarios where an energy function [71] is used to represent an optimisation problem [23,24], in which case the system necessarily has only fixed point attractors (corresponding to local minima of the energy function/local optima of the optimisation problem) but may nonetheless represent a difficult optimisation problem [81] (see the problem of finding the ground state of a spin-glass system [28]). We also know that (non-associative) local search processes or (non-associative) evolution cannot scale well on the modular problem (S2) even though it is consistent (i.e. in consistent problems, configurations exist that satisfy all constraints simultaneously, but many configurations that fail to satisfy all constraints are local equilibria) [80].

S1. Random sparse symmetric matrices. A connection matrix of N variables with Nk non-zero connections (i.e. density of non-zero connections is $Nk/N^2 = k/N$, or k dependencies per node on average). Non-zero connections take value -1 or 1 with equal probability. $\omega_{ij}^O = \{-1, 0, 1\}$. $N=120$, $k=8$.

S2. Modular consistent connectivity matrix: Intra-module connections=1, all inter-module connections= p ($p>0$). i.e.

¹ Making $\omega_{ij} = \omega_{ji}$ means that with respect to the reciprocal connection between two agents, a strategy change that increases the *contribution to utility from this one connection* to agent i also increases the *contribution to utility from this one connection* to agent j . But the significance of this should not be over-interpreted. In particular, since the balance of constraints on each agent depends on its connections with *all* its neighbours, it does not mean that ‘what is good for agent i is good for agent j ’ – i.e. a strategy that increases the utility of agent i may decrease the utility of agent j (because i forgoes mutual benefit with j in order to receive benefit from some other agent k). More specifically, circuits of constraints between multiple agents cannot be simultaneously satisfied when they contain an odd number of negative weights [28] (e.g. $\omega_{ij} = \omega_{ji} = 1$, $\omega_{jk} = \omega_{kj} = 1$, $\omega_{ki} = \omega_{ik} = -1$ creates a condition where j needs to coordinate with both i and k , but i needs to anti-coordinate with k).

$$\omega_{ij}^O = 1, \text{ if } \left\lfloor \frac{i}{k} \right\rfloor = \left\lfloor \frac{j}{k} \right\rfloor \quad \omega_{ij}^O = p, \text{ otherwise }^2.$$

where k is the size of modules. $N=100$, $k=5$. $p=0.01$.

In the modular function the connections within the block diagonal regions are stronger than the inter-module connections [68,79]. Intermediate p , as used here, where inter-module dependencies are weak but non-trivial, creates $2^{\frac{N}{k}}$ local equilibria (corresponding to local optima in total utility) only two of which are globally optimal [81].

Results

In overview, the experiments are as follows. A network of agents is created with either an initially random or initially modular connection structure. Both systems have a very large number of local equilibria which exhibit various total-utility values. The behaviours of the agents on the network are repeatedly perturbed to different initial conditions and allowed to relax again to an attractor. The total utility of the state configuration found at that attractor is recorded. We also allow agents to alter their connections with other agents (by a fixed-magnitude increment, r) in whichever direction increases their individual utility. We assess how these selfish changes to the connection structure, accumulated slowly as the system visits many local equilibria, alters the distribution of attractors that the system visits. Note that the total utility is measured using the original connection values so that we can assess how the new/augmented connections of the network affect the likelihood of finding state configurations that resolve the tensions imposed by the *original* constraints of the network.

Fig. 1 shows that selfish modification of connections results in enlargement of high-utility attractors at the expense of low-utility attractors, such that the total utility of equilibria found by the network with the modified connections is higher on average than the total utility of equilibria found by the network with the original connections. In the modular system the attractors of the system are known and we can estimate how the size of the basin of attraction for each attractor class (grouped by total utility) changes over time. There are many more low-utility attractors than high-utility attractors in this system and the global optimum is rarely visited initially (Fig. 2). With unmodified connections, the time to first hit of the global optimum is 1009.8 ± 92.00 relaxations (s.e.m.) on average (over 100 independent experiments). In contrast, with slowly modified connections the time to first hit is 75.9 ± 2.51 (s.e.m.) (mean time to first hit of the global optimum is different with confidence p-value $< 10^{-5}$). Indeed, after 200 relaxations with slowly modified connections the global optimum has become practically the only attractor of the system (Fig. 2). This means that selfish modifications to connections are not merely increasing the basin of high-utility attractors that have already been visited, but are enlarging the basin of attraction for these attractors *before* they are visited for the first time. In the modular system the original multi-modal ‘optimisation problem’ posed by the original connections has been transformed into an easy single-optimum ‘problem’ (See Discussion, Fig.3).

Naively, it might seem obvious that selfish agents will modify the malleable dependencies in the problem structure to make the problem they face easier for themselves. But recall that we are measuring total utility using only the original, unchanged connections/constraints. How does altering the malleable constraints improve satisfaction of the unaltered constraints? Although the total utility of any given behavioural configuration after connections are altered is the same as it was before connections were altered (i.e. alterations affect ω_{ij}^L but not ω_{ij}^O), the average total utility of configurations *that are visited* can change. This occurs not because the utility of any configuration has increased, but because the distribution of configurations that are visited has been altered such that higher-utility configurations are visited more often. Specifically, the easy-to-find local optimum of the new system dynamics is in fact the hard-to-find maximally-optimal local attractor of the original system dynamics. The malleable constraints of the system thus induce dynamical behaviour that solves the constraints imposed by the original and unmodifiable connections (Discussion).

² $\lfloor x \rfloor$ is the integer part of x ($x \geq 0$).

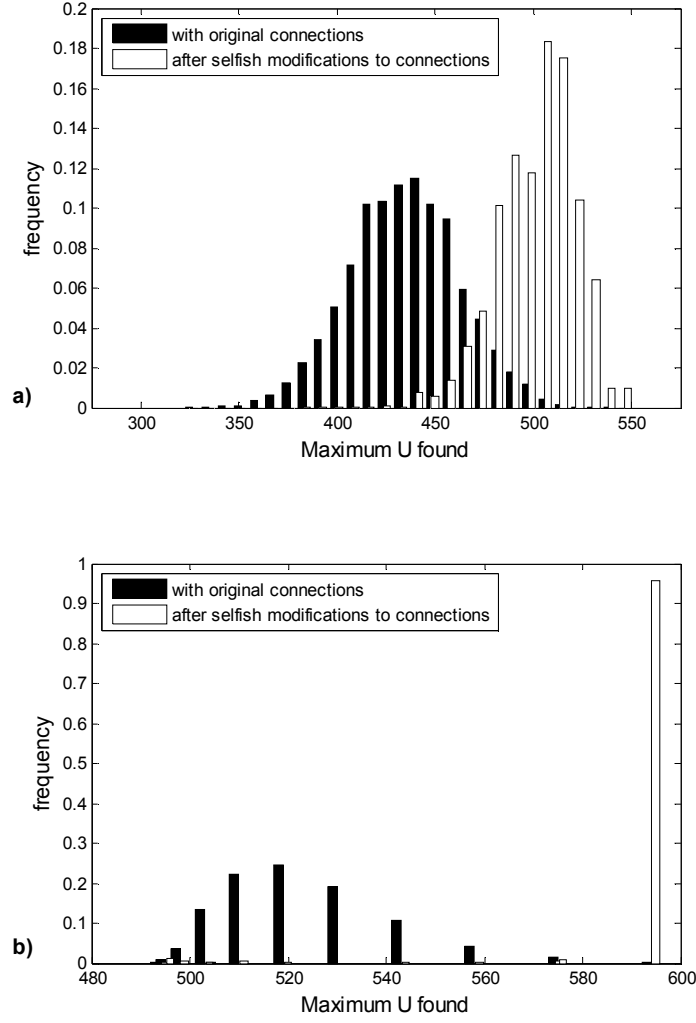


Fig. 1: Histograms of the total utility (U) at end-of-relaxation state configurations before and after modification of the network connections: a) initially random structure ($S1$, learning rate $r=0.0002$), b) initially modular structure ($S2$, $r=0.002$). ('after' = after 500 and 150 relaxations with restructuring for $S1$ and $S2$ respectively. Data are averaged from 100 independent runs in each case, and distributions are generated from 100 sample relaxations per run for $S1$ and 25 sample relaxations per run for $S2$.) These show that the probability of finding configurations that are high total utility is significantly increased by selfish modifications to the connections. In the unstructured ($S1$) case, 99.56% of the attractors found before restructuring have lower total utility than the *modal* attractor found after restructuring. The effect is even more pronounced in the modular system, $S2$, where the modal attractor after restructuring is the attractor with *globally maximal* total utility, and whereas only 0.08% of relaxations hit this configuration before restructuring, 95.8% of relaxations find the global optimum after restructuring.

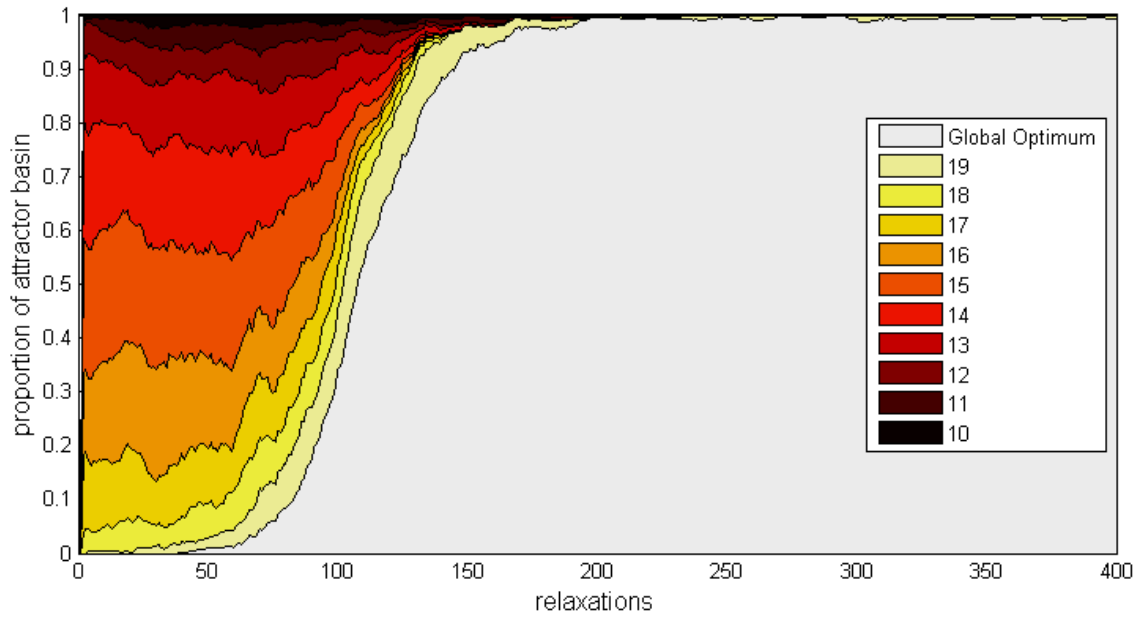


Fig. 2. The change in the size of the attractor basins in the (initially) modular system over time (stacked plot). Colour/shading indicates membership in the eleven possible U -value classes (*c.f.* 20 modules). The attractor of the globally optimal configuration (lightest shade) is initially very small but becomes the only significant attractor of the system. (The attractor basin size is averaged using a sliding window of 20 relaxations and averaged over 100 independent runs).

Discussion

Are selfish changes to connections necessarily Hebbian?

The suggestion that theory from neural networks might be directly relevant to other types of complex adaptive systems is not unique [60,53,9,10,43]. But the utility of this transfer depends on whether the manner in which selfish agents alter connections is homologous with the learning rules familiar in neural networks.

When an agent modifies connections to maximise individual utility it necessarily strengthens connections that support the current state of the agent and weakens those that oppose it (see Methods). It therefore strengthens connections that are currently satisfied and weakens connections that are currently in tension. Changes of this form are preferred because only these will increase the *immediate* pay-off that the agent receives - increasing its preference for the current combination of behavioural states. Since connections change slowly, the pair of behaviours experienced when the modification took place will also be preferred more *in future* (as well as immediately). The consequence for the dynamics of the system as a whole is to increase the basin of attraction for state configurations that have been previously experienced. Although it might not be immediately apparent, the fact that selfish changes to connections will necessarily be Hebbian should not be too much of a surprise given what we know about Hebbian learning. Specifically, Hebbian learning creates an associative memory by increasing the basin of attraction for the presented pattern, and to do that it must increase the preference (or decrease the energy) for the combinations of states in that pattern, and in order to do that it must (just as a selfish agent would) increase the strength of connections that agree with the pattern and weaken those that would disrupt it. Appendix 1 discusses how this result holds regardless of the form of utility function used.

To be clear, agents do not modify connections in a Hebbian manner *in order to* alter the dynamics of the system, nor to create an associative memory, nor to increase the utility of others or their own future utility (although the modifications do all these things), but purely for immediate individual benefit. Nonetheless, the result of these changes produces the same dynamical consequences that are well-known for Hebbian learning in neural dynamics such as the potential to form an associative memory, to ‘recall’ past patterns of activation from similar initial conditions, and in some cases to generalise patterns of activation to create new combinations of features observed in past state configurations [22].

Memory, optimisation and generalisation.

Although the ‘energy minimisation’ behaviour of the Hopfield network and its interpretation as an optimisation process (Fig. 3.a) are well known [23,24,71], and similarly, the ability of Hebbian learning to implement an associative memory of a set of training patterns and ‘recall’ them or ‘recognise’ them from noisy or partial examples (Fig. 3.b) is also well known [22], in contrast, the result of combining these two behaviours in the same network is novel (Fig. 3.c) [81,78]. Given that good attractors are large attractors [11,38,37,39], the results above can then be partly explained by the development of an

associative memory that simply ‘recalls’ the best configurations previously visited since these are the configurations that are most frequently visited.

However, the results we observe cannot be entirely explained this way because the network modification enlarges the basin of attraction for the global optimum *before* it is visited for the first time (see Results). In other words, the system is ‘predicting’ the location of high-utility attractors not merely recalling them. This behaviour can be understood via the inherent ability of associative learning to generalise training patterns (Fig. 4). Thus the enlargement of attractors depends not only on the frequency with which different behavioural patterns are visited but also on the commonality of features between those patterns.

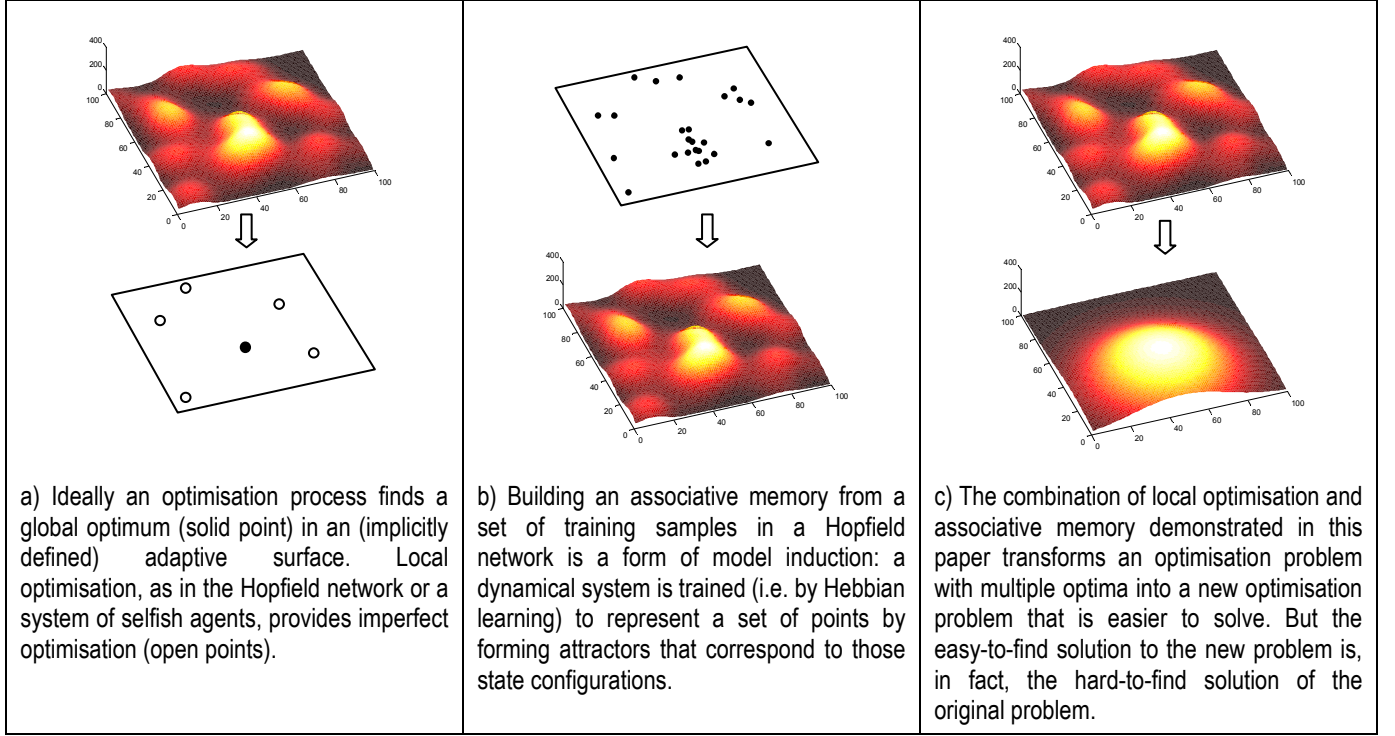


Fig. 3: Schematic of well-known behaviours of the Hopfield network (a) and (b), and comparison with their combination in our model where network connections are adapted in response to the systems own dynamical attractors (c) [78,81].

The potential for Hebbian learning to generate new combinations of features by generalising over observed patterns, including the phenomenon known as ‘spurious’ attractors [32], is particularly important in this respect. ‘Spurious’ attractors are attractors in an associative memory that are distant from any training pattern (Fig. 4.c). These are created when many training patterns have large components in common with each other and the new attractors correspond to novel combinations of these features [32]. In an associative memory this is naturally considered to be a nuisance (e.g. [12]) but in optimisation the ability to produce new combinations of successful features is highly-desirable [21,77,80,81]. Mills [46,47,44] shows that the automatic discovery and utilisation of modular structure in an optimisation problem, as facilitated by learned associations, can be used to provide significant optimisation performance (see also [26,27,84,93]). This result thereby shows that a distributed optimisation process, based on nothing more than repeated relaxation of state configurations plus local selfish reinforcement of connections has the effect not only of creating an associative memory of its past local optimisation behaviour but also generalising its past behaviour and enabling superior optimisation and, in the context of a multi-agent system, global adaptation.

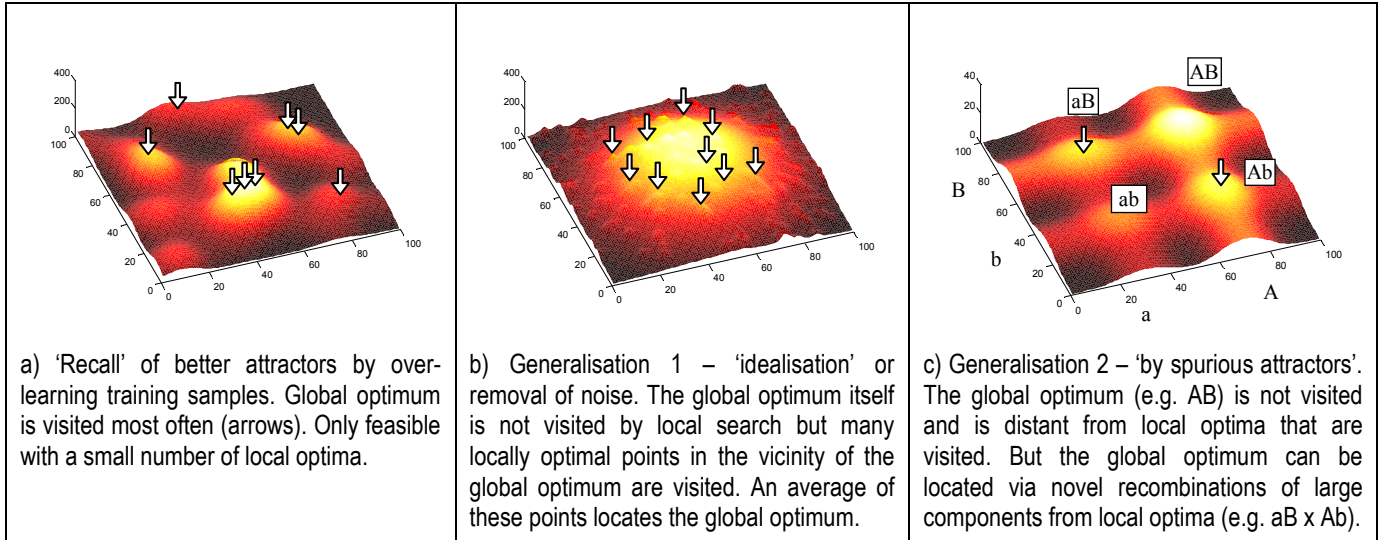


Fig. 4: Schematic of mechanisms involved in finding high-utility attractors: a) recall, b) generalisation in 'unstructured' problems, c) generalisation in modular problem structures. The latter is made possible by the fact that the landscape must be a superposition of many independent components. In high-dimensional structured spaces (unlike the two dimensional illustration shown here) the original system dynamics may be very unlikely to visit the global optimum but nonetheless very likely to visit a set of points that together contain all the components of the global optimum (hence providing the necessary data for successful generalisation).

The equivalence between an explicit global optimisation process and spontaneous global adaptation in multi-agent systems shown here is made exact by the use of an agent system whose dynamics can be described with an energy or potential function (increasing total utility) [71], and this is possible because we assumed symmetric constraints [22,24]. In more general systems, actions that increase individual utility will not necessarily increase total utility even locally (i.e. total utility will not provide a potential function for agent behaviours) and in this case the attractor states that the system spends most time in (and are thereby enlarged the most by changes to connections) will not necessarily correspond to local optima in total utility. In particular, the configuration that is globally optimal in total utility may not be an equilibrium and, in this case, it cannot be enlarged by selfish changes to connections. This does not mean that selfish changes to connections will not produce an increase in global adaptation, however. That is, from the subset of configurations that *are* local equilibria, some will have higher total utility than others. If it is still true that local equilibria with higher total utility tend to have larger basins of attraction, then those attractors (and generalisations thereof) will be enlarged at the expense of lower-utility attractors. We conjecture that, in systems built out of the superposition of many low-order interactions, this must be true *on average* [38,37,39]. But of course, this will not be true in every instance – and in systems that have only one equilibrium, as in a single, two-player Prisoner's Dilemma, for example, the enlargement of this attractor will have neither a positive nor a detrimental effect on total utility.

How can selfish behaviours be cooperative?

All the changes in our system, both the agents' basic behaviours and their network restructuring behaviours, are entirely explicable in terms of selfish motives. Finding the joint equilibrium of these two types of selfish behaviours thus explains the observations in our results [62,69,61]. Moreover, in a deep sense, this must be the case if we seek a mechanistic explanation of the results, avoiding mystical sources of spontaneous cooperation. So if individual selfishness motivates and explains the results, in what sense then is it necessary or even meaningful to suggest that the resulting system is holistically adapted?³ Indeed, the very idea that selfish behaviours can create increased cooperation seems contradictory. But the idea that selfish changes to *relationships* can create cooperative *behaviours*, is not a contradiction. That is, a perspective that maintains a conceptual separation of the ordinary behaviours and the restructuring behaviours resolves this paradox by viewing the structure of the system at any given time as a parameter of the ordinary behaviours. It will necessarily be the case that the ordinary behaviours *given the current structure at any point in time* are selfish, but selfish behaviours in one structure yield higher total utility than selfish behaviours in another structure. Without this conceptual separation between ordinary and restructuring behaviours, questions about how the system *adapts* over time have no meaning – we cannot say that the

³ We note that the same tension causes detailed group selection models to 'explain away' the phenomenon they sought to capture [91,89,61].

system’s behaviour is better or worse than it used to be, it is simply the same system in a new state⁴. This separation of population structure and social behaviour [69,61] is also exploited in related work examining the concurrent evolution of non-network-based population structuring traits and cooperative traits, both changing under natural selection [63,64].

Related work

In the social network domain we have also shown [7] that ‘habituating’ agents, acting as ‘creatures of habit’ due to a tendency to increase their preference for the ‘status quo’, are also formally equivalent to Hebbian learning at the system level and produce global adaptation in the same manner as that shown here. This work frames the utility function as a weighted poly-player coordination game [41,69,73], and the system dynamics as repeated episodes of attempting to reach consensus given constraints amongst players. This work separates the component of a constraint that can be changed and the component of a constraint that cannot be changed (here, α_{ij}^f and α_{ij}^o , respectively) into the ‘perception’ of a game and the actual game (or its ‘true’ utility). Since agent behaviours are selected based on their distorted perceptions they will sometimes select behaviours that decrease their true utility (which must be the case if the attractors of the system are to be different). But the different attractors that are consequently reached have higher total true utility (and accordingly, higher individual utility on average) than the attractors reached by immediate maximisation of true utility. This work also discusses the close relationship between actions which amplify the *immediate* preference for the status-quo *averaged over long timescales* (i.e. given a slow-changing preference) and actions that actually maximise the *long term* utility of an individual. This is closely related to a simple form of reinforcement learning where agents play the strategy that maximises their utility given the *expected* strategies adopted by other agents based on past behaviour [40]. But given the separation of timescales (i.e. rapid relaxations and slowly changing associations) adopted in the self-modelling framework [78], a simple ‘habituation’ of correlations observed at equilibria (i.e. preferring the *status quo*) achieves the same outcome as a preferential reinforcement of ‘good’ correlations.

In an evolutionary scenario, an individual-based simulation [85,83] supports the conclusion that individuals will evolve associations that reinforce locally stable equilibria as shown here. More exactly, this work shows that associations canalise common components of locally stable equilibria – in other words, they create groups that reflect the commonly occurring sub-patterns of configurations that are visited not the entire configuration patterns. Whereas the present work implements a subtle form of ‘social niche construction’ [59,62,63], altering interaction coefficients that affect the attractors of the behavioural dynamics [59,63], this related work shows how similar principles determine how evolving agents create new ‘vehicles’ [8] in the sense of strongly correlated co-dispersal probabilities (see also ‘macro variation’ [47,44]). The grouping of evolved species or symbioses can be understood as a decomposition of the problem variables into sub-sets that are approximately independent from one another but strongly interdependent internally. The scalability of an associative optimisation process based on these principles is also shown to be algorithmically superior to non-associative evolution in a formal sense [85,46,84]. The associative memory principles discussed here may therefore help us better understand the mechanisms of major evolutionary transitions [42] (developed elsewhere [83,84]).

Recent results on the evolution of evolvability suggest intriguing parallels between the evolution of gene expression networks and other complex adaptive systems discussed here. Specifically, Parter et al [57] have investigated how the evolvability of a population changes over time when it is subjected to a fluctuating environment (we achieve the same conditions using repeated perturbation in a static environment [82]). They find that organisms develop a “memory” of their evolutionary history and observe that evolved networks “generalise to future environments, exhibiting high adaptability to novel goals”. Wagner et al [76] explain part of the mechanism that might be involved by referring to genetic loci that affect the correlation of phenotypic traits [58] as follows: “natural selection can act on [variation in the strength of the pleiotropic effects between genes] to either increase the correlation among traits or decrease it depending on whether the traits are simultaneously under directional selection or not. ...[resulting in] a reinforcement of pleiotropic effects among co-selected traits and suppression of pleiotropic effects that are not selected together” [76]. This clearly describes a Hebbian modification of gene interactions. Recent work of our own [82] has unified Wagner and Parter’s observations with the findings of the current paper to provide a mechanism that explains the increase in evolvability that Parter et al observe. The interesting finding of these evolutionary scenarios is that changes to connections within a network will be Hebbian regardless of whether the unit being evolved is the individual node in the network or the whole network, thus explaining equally well both global adaptation (e.g. in an ecosystem) and increased evolvability (e.g. in a gene expression network).

We note that the interplay of rapid behavioural change and simultaneous but slower changes to connections offers a close analogy with the Baldwin effect [3,18,45,77] that warrants further investigation. In the Baldwin effect a relatively slow learning mechanism canalises the results of relatively fast phenotypic exploration. But by canalising associations between behaviours, rather than fixing individual behaviours, the effects of Hebbian learning on adaptation are more flexible, creating what we might term a ‘second-order canalisation’, i.e. canalisation that can bias which *combinations* of behaviours occur without precluding any individual behavioural values.

⁴ This can be described as a ‘meta-dynamical system’ [4]; i.e. a dynamical system of ordinary behaviours, S , with parameters that include the connection topology – but these parameters are actually slowly changing variables in a larger dynamical system, S' . We then ask how the dynamics of S' reorganise the structure of, and hence change the dynamics of, S .

Relaxation of the symmetry conditions (adopted here in the initial connections) and the cyclic or chaotic attractors that may result should also be investigated. The two extremes of initial structure studied here (i.e. totally random structure and neat modularity) suggest we should also explore the effect of semi-modular or ‘nearly decomposable’ systems [68,79,77]. The simplification we have employed by studying fully-connected weighted networks rather than sparsely connected topologies also warrants further investigation. When connections are all-or-nothing, as they often are in social network models, a connection either creates correlation or anti-correlation in the states that it connects and cannot represent an average of state correlations observed over time as modelled here. However, on average an ensemble of all-or-nothing connections may represent past state correlation in a probabilistic or aggregate sense.

Conclusions

We have shown that organisational principles familiar in organismic learning occur implicitly in distributed complex adaptive systems. System-level ‘learning’ or associative induction happens as a direct consequence of the fact that selfish modifications to relationships between components are equivalent to Hebb’s rule. Such changes are motivated by purely local and immediate benefits but have the side-effect of enlarging the basin of attraction for those behaviour combinations in future state dynamics. Thus networks of selfish agents self-organise the connection structure of a network in a manner that creates an associative memory of state configurations that the system experiences.

Given this, when the state of the system is sufficiently perturbed and the changes to connections are sufficiently slow, agents will modify connections in a manner that enhances total utility. More exactly, agents find connections that coordinate their behaviours with other agents more effectively such that the system as a whole finds better solutions to the original constraints between agents (even though these constraints cannot be altered). The improved global efficiency that results can be understood via the inherent ability of associative learning to generalise by idealising stored patterns and/or creating new combinations of learned sub-patterns. These findings suggest that distributed complex adaptive systems of self-interested components, such as individuals in a social network or species in an ecosystem, may exhibit organisational principles in common with those familiar in organismic learning, developing an associative memory of their past behaviour that enhances system-level efficiency in future. In the case of species in an ecosystem, although each species may adapt its relationships with other species via natural selection, the global adaptation of the network does not require that the ecosystem as a whole is reproduced with inheritance, variation and selection to achieve this. This work thereby demonstrates a completely distributed adaptive process that we view as a natural extension to the “emergent collective computational abilities” that come ‘for free’ in physical systems [22].

Acknowledgements. Thanks to Simon Powers, Alex Penn, Adam Davies, Chrisantha Fernando, Jason Noble, Seth Bullock, Hod Lipson, Archie Chapman and Nic Palmius.

References

1. Ackley, D.H., Hinton, G.E. & Sejnowski, T.J. (1985). A Learning Algorithm for Boltzmann Machines, *Cognitive Science*, 9, 147-169.
2. Aktipis, C.A., (2004). Know when to walk away: contingent movement and the evolution of cooperation. *Journal of Theoretical Biology*, 231(2), 249-260.
3. Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, 30:441–451.
4. Bourguine, P. & Varela, F.J. (1992). Toward a Practice of Autonomous Systems. In Varela, F.J. & Bourguine, P. (Eds.) *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, (pp. xi-xvii). Cambridge, Mass.: MIT Press
5. Boyd, R. & Richerson, P. J. (1990). Group Selection Among Alternative Evolutionarily Stable Strategies. *Journal of Theoretical Biology*, 145, 331-342.
6. Branchtein, M.C., & Arenzon, J.J. (1992). Categorization and Generalization in the Hopfield Model. *J. Phys. I (France)* 22019.
7. Davies, A.P., Watson, R.A., Mills, R., Buckley, C. L., & Noble, J. (in press). If you can't be with the one you love, love the one you're with: How individual habituation of agent interactions improves global utility. *Artificial Life*.
8. Dawkins, R. (1976). *The Selfish Gene*. New York: Oxford University Press.
9. Fernando, C. & Szathmari, S. (2010). Natural Selection in the Brain. *Chapter 20. Towards a Theory of Thinking Vol 2*.
10. Fernando, C., Liekens, A.M.L., Bingle, L.E.H., Beck, C., Lenser, T., Stekel, D.J., & Rowe, J.E. (2008) Molecular circuits for associative learning in single-celled organisms. *Journal of the Royal Society Interface*. 6(34):463-9.
11. Gross T. & Sayama, H. (eds.) (2009) *Adaptive Networks: Theory, Models and Applications*. Springer Verlag..
12. Gascuel, J.-D., Moobed, B. & Weinfeld, M. (1994). An Internal Mechanism for Detecting Parasite Attractors in a Hopfield Network, *Neural Computation* 6, 902-915.

13. Haldane, J. B. S. (1927). A mathematical theory of natural and artificial selection. V. Selection and mutation. *Proceedings of the Cambridge Philosophical Society* 23, 838–844.
14. Harsanyi, J.C., & Selten, R. (1988). *A General Theory of Equilibrium Selection in Games* (MIT Press, 1988).
15. Hauert, C., De Monte, S., Hofbauer, J., & Sigmund, K. (2002). Volunteering as Red Queen Mechanism for Cooperators in Public Goods Games. *Science* 296, 1129–1132.
16. Hebb, D.O. (1949). *The organization of behaviour*. New York: Wiley.
17. Heylighen, F., Gershenson, C., Staab, S., Flake, G.W., Pennock, D.M., Fain, D.C., De Roure, D., Aberer, K., Wei-Min Shen, Dousse, O. & Thiran, P. (2003). Neurons, viscose fluids, freshwater polyp hydra-and self-organizing information systems. *Intelligent Systems*, 18(4) 72 – 86.
18. Hinton, G. E. & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.
19. Hinton, G.E., & Sejnowski, T.J. (1985). Learning in Boltzmann Machines, *Cognitive*, 85, Paris, France.
20. Hogeweg, P. (2007). From population dynamics to ecoinformatics: Ecosystems as multilevel information processing systems. *Ecological Informatics*, 2, 103–111.
21. Holland, J.H. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.
22. Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities, *PNAS USA*, 79 (8) 2554–2558.
23. Hopfield, J.J., & Tank, D.W. (1985). 'Neural' computation of decisions in optimization problems. *Biol Cybern* 52:141–152.
24. Hopfield, J.J., & Tank, D.W. (1986). Computing with neural circuits: A model. *Science* 233: 625–633
25. Hopfield, J.J., Feinstein, D. & Palmer, R. (1983). 'Unlearning' has a Stabilizing Effect in Collective Memories. *Nature*, 304. 158–159.
26. Iclanzan, D., & Dumitrescu, D. (2007). Overcoming hierarchical difficulty by hill-climbing the building block structure. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2007)*. 1256–1263.
27. Iclanzan, D., & Dumitrescu, D. (2008). How can Artificial Neural Networks help making the intractable search spaces tractable. In *IEEE Congress on Evolutionary Computation (CEC 2008)*. 4015–4022.
28. Istrail, S. (2000) Statistical Mechanics, Three-Dimensionality and NP-Completeness: I. Universality of Intractability of the Partition Functions of the Ising Model Across Non-Planar Lattices, In *Proceedings of the 32nd ACM Symposium on the Theory of Computing (STOC00)*, ACM Press, p.87–96.
29. Jackson, M.O. & Watts, A. (2002a). On the Formation of Interaction Networks in Social Coordination Games, *Games and Economic Behavior*, 41(2), 265–291.
30. Jackson, M.O. & Watts, A. (2002b). The Evolution of Social and Economic Networks, *Journal of Economic Theory*, 106(2), 265–295.
31. Jackson, M.O. & Wolinsky, A. (1996). A Strategic Model of Social and Economic Networks, *Journal of Economic Theory*, 71, 44–74.
32. Jang, J.-S., Kim, M.W., & Lee, Y. (1992) A Conceptual Interpretation of Spurious Memories in the Hopfield-type Neural Network. *Neural Networks, 1992. IJCNN., International Joint Conference on*. 1:21–26.
33. Kauffman, S. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford: Oxford Univ. Press.
34. Keeling, M. J. (2005). The implications of network structure for epidemic dynamics. *Theoretical Population Biology* 67, 1–8.
35. Kirkpatrick, S., Gelatt, C.D. & Vecchi M.P. (1983). Optimization by Simulated Annealing. *Science*. New Series 220 (4598): 671–680.
36. Krebs, P.R., & Theumann, W.K. (1993). Generalization in a Hopfield network with noise. *J. Phys. A*. 26 - 3983.
37. Kryzhanovsky, B., Magomedov, B.M. & Fonarev, A.B. (2006). On the Probability of Finding Local Minima in Optimization Problems. *IJCNN 2006*: 3243–3248.
38. Kryzhanovsky, B., & Kryzhanovsky, V. (2008). Binary Optimization: On the Probability of a Local Minimum Detection in Random Search. *ICAISC 2008*: 89–100.
39. Kryzhanovsky, B., Kryzhanovsky, V. & Mikaelian, A. L. (2007). Binary optimization: A relation between the depth of a local minimum and the probability of its detection. *ICINCO-ICSO 2007*: 5–10
40. Leslie, D.S. & Collins, E.J. (2006). Generalised weakened fictitious play, *Games and Economic Behavior*, 56: 2, pp 285–298.
41. Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge.
42. Maynard Smith, J., & Szathmary, E. (1995). *The major transitions in evolution*. Oxford.
43. Mikhailova, A.S., Mit'kov, I.V. & Sveshnikov, N.A. (1990). Molecular associative memory. *Biosystems*. 23(4):291–5.
44. Mills, R. (2010). *How Micro-Evolution Can Guide Macro-Evolution: Multi-Scale Search via Evolved Modular Variation*. Doctoral dissertation, University of Southampton, Southampton.
45. Mills, R. & Watson, R. A. (2006). On crossing fitness valleys with the Baldwin effect. *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pp. 493–499, MIT Press
46. Mills, R., & Watson, R.A. (2009). Symbiosis Enables the Evolution of Rare Complexes in Structured Environments. *Proceedings of European Conference on Artificial Life 2009*.

47. Mills, R., & Watson, R.A. (in prep.) "Multi-scale search, modular variation, and adaptive neighbourhoods". Submitted.
48. Muehlenbein, H. (1992). How genetic algorithms really work: mutation and hill climbing. *Proceedings of Parallel Problem Solving From Nature 2*, pp. 15-26, Elsevier.
49. Nash, J. (1951). Non-Cooperative Games. *The Annals of Mathematics* 54(2):286-295.
50. Newman, M. E. J. (2002). Spread of epidemic disease on networks. *Physical Review E* 66, 016128.
51. Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review* 45 (2), 167-256.
52. Newman, M. E. J., Barabasi, A. L. & Watts, D. J. (Eds.) (2006). *The Structure and Dynamics of Networks*. Princeton University Press.
53. Noonburg, V.W. (1989). A Neural Network Modeled by an Adaptive Lotka-Volterra System. *SIAM Journal on Applied Mathematics*, 49:6. pp. 1779-1792.
54. Odling-Smee, F.J., Laland, K.N., & Feldman, M.W. (2003). *Niche Construction. The Neglected Process in Evolution*. Monographs in Population Biology. 37. Princeton University Press.
55. Pacheco, J. M., Lenaerts, T. & Santos, F.C. (2007). Evolution of Cooperation in a Population of Selfish Adaptive Agents. *Advances in Artificial Life 2007. LNCS 4648*, 535-544.
56. Pacheco, J. M., Traulsen, A. & Nowak, M. A. (2006). Coevolution of Strategy and Structure in Complex Networks with Dynamical Linking. *Phys. Rev. Lett.* 97, 258103.
57. Parter, M., Kashtan, N., & Alon, U. (2008). Facilitated Variation: How Evolution Learns from Past Environments To Generalize to New Environments. *PLoS Comput Biol* 4(11): e1000206.
58. Pavlicev, M., Kenney-Hunt, J.P., Norgard, E.A., Roseman, C.C., Wolf, J.B. & Cheverud, J.M. (2008). Genetic variation in pleiotropy: differential epistasis as a source of variation in the allometric relationship between long bone lengths and body weight. *Evolution*. 2008 Jan;62(1):199-213.
59. Penn, A. S. (2006). *Ecosystem Selection: Simulation, Experiment and Theory*. Doctoral dissertation, University of Sussex, Falmer.
60. Poderoso, F.C., & Fontanari, J.F. (2007). Model ecosystem with variable interspecies interactions, *J. Phys. A: Math. Theor.* 40 8723-8738.
61. Powers, S. T. (2010) *Social Niche Construction: Evolutionary Explanations for Cooperative Group Formation*, Doctoral dissertation, ECS Southampton University, Southampton.
62. Powers, S.T., Mills, R. Penn, A.S., & Watson, R.A. (2009) Social niche construction provides an adaptive explanation for new levels of individuality (ABSTRACT), *Proceedings of Workshop on Levels of Selection and Individuality in Evolution*, In European Conference on Artificial Life.
63. Powers, S.T., Penn, A.S. & Watson, R.A. (2007) Individual Selection for Cooperative Group Formation. *Proceedings of European Conference on Artificial Life 2007*. pp. 585-594.
64. Powers, S.T., Penn, A.S., & Watson, R.A. (in press) The concurrent evolution of social traits and the population structures that support them. *Evolution*.
65. Roughgarden, J., (1979). *Theory of Population Genetics and Evolutionary Ecology: An Introduction*. New York: Macmillan Publishing Company.
66. Rumelhart, D.E., & McClelland, J.M. (1986). *Parallel distributed processing- Explorations in the Microstructure of Cognition Volume 1: Foundations*, MIT Press.
67. Sachtjen, M. L., Carreras, B. A. & Lynch, V. E. (2000). Disturbances in power transmission system. *Physical Review E* 61, 4877.
68. Simon, H.A. (1969). *The Sciences of the Artificial*. Cambridge, Mass.: MIT Press.
69. Skyrms, B. (2004). *The Stag Hunt and Evolution of Social Structure*. Cambridge: Cambridge University Press.
70. Solé, R. V.; Alonso, D. & McKane, A. (2000). Scaling in a network model of a multispecies ecosystem *Physica A*, 286, 337-344.
71. Strogatz, S.H. (1994). *Nonlinear Dynamics and Chaos*. Reading, Mass.: Addison-Wesley.
72. Taylor, C., & Nowak, M.A. (2007). Transforming the dilemma. *Evolution*. 61(10):2281-92.
73. Tomassini M., & Pestelacci E. (2010). Coordination Games on Dynamical Networks. *Games*. 1(3):242-261.
74. Traulsen, A., Santos, F.C. & Pacheco, J.M. (2008). Evolutionary games in self-organizing populations, in *Adaptive Networks: Theory, Models and Applications*, Gross, T. & Sayama, H. (Eds.) New England Complex Systems Institute Book Series on Complexity. N.Y.:Springer.
75. Van Segbroeck S., Santos F.C., Pacheco J.M., Lenaerts T. (2010). Coevolution of Cooperation, Response to Adverse Social Ties and Network Structure. *Games*. 1(3):317-337.
76. Wagner, G.P., Pavlicev, M., & Cheverud, J.M. (2007). The road to modularity. *Nature Reviews Genetics* 8: 921-931.
77. Watson, R. A. (2006). *Compositional Evolution: The impact of Sex, Symbiosis and Modularity on the Gradualist Framework of Evolution*. (Vienna series in theoretical biology). Mass.:MIT Press.
78. Watson, R. A., Buckley, C. L. & Mills, R. (2010). Optimisation in 'Self-modelling' Complex Adaptive Systems. *Complexity*.
79. Watson, R.A. & Pollack, J.B. (2005). Modular Interdependency in Complex Dynamical Systems. *Artificial Life*, 11 (4). pp. 445-457.

80. Watson, R.A., & Jansen, T. (2007). A Building-Block Royal Road Where Crossover is Provably Essential. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2007)*. pp. 1452-1459.
81. Watson, R.A., Buckley, C.L. & Mills, R. (2009). *The Effect of Hebbian Learning on Optimisation in Hopfield Networks*. Technical Report, ECS, University of Southampton, Southampton.
82. Watson, R.A., Buckley, C.L., Mills, R. & Davies, A. (2010) Associative memory in gene regulation networks. *In procs. Artificial Life 2010*. 194-202.
83. Watson, R.A., Jackson, A., Palmius, N., Mills, R., Powers, S.T. & Penn, A.S. (in prep) “The Origin of Group Selection via Individual Selection”.
84. Watson, R.A., Mills, R., & Buckley, C. L. (in press) “Transformations in the Scale of Behaviour and the Global Optimisation of Constraints in Adaptive Networks”, *Adaptive Behaviour*.
85. Watson, R.A., Palmius, N., Mills, R., Powers, S.T., & Penn, A.S. (2009). Can Selfish Symbioses Effect Higher-level Selection? *European Conference on Artificial Life 2009*.
86. Watts, D. J. & Dodds, P. S. (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research* 34, 441.
87. Watts, D. J. & Strogatz, S. H. (1998). Collective dynamics of small-world networks. *Nature* 393, 440-442.
88. Werfel, J. & Bar-Yam, Y. (2004). The evolution of reproductive restraint through social communication. *Proceedings of the National Academy of Science* 101 (30), 11019-11024.
89. Wild, G., Gardner, A. & West, S.A. (2009). Adaptation and the evolution of parasite virulence in a connected world. *Nature* 459, 983-986.
90. Wilmers, C.C., (2007). Understanding ecosystem robustness. *Trends in Ecology & Evolution*, Volume 22, Issue 10, 504-506.
91. Wilson, D.S. & Wilson, E. O. (2007). Rethinking the theoretical foundation of socio-biology. *The Quarterly review of biology*, Vol. 82, No. 4. (December 2007), pp. 327-348.
92. Wilson, D.S. (1980). *The Natural Selection of Populations and Communities*. Benjamin/Cummings.
93. Yu, T.-L., Goldberg, D. E., Yassine, A. & Chen, Y.-p. (2003). Genetic algorithm design inspired by organizational theory: Pilot study of a dependency structure matrix driven genetic algorithm, *Proceedings of Artificial Neural Networks in Engineering 2003 (ANNIE 2003)*, 327–332, St. Louis, Missouri, USA, Nov. 2003.

Appendix 1

Syntactically, the result that selfish changes to connections are Hebbian is sensitive to the form of the original utility function we assumed⁵ (Eq. 1). For example, if instead we had assumed maximisation of Eq. 4, then selfish changes to connections would (syntactically) be anti-Hebbian (i.e. $\delta < 0$), not Hebbian.

$$u_i = - \sum_j^N \omega_{ij} S_i S_j \quad (4).$$

However, this is merely a labelling error with respect to the semantics of a connection. Specifically, Hebb’s rule assumes that positive weights increase the correlation between connected states, rather than decrease it (i.e. an ‘excitatory’ connection causes the activation of one neuron to excite another not depress it). Under Eq. 4 a connection’s sign has the opposite meaning and likewise the syntax of Hebb’s rule must be reversed in order to maintain the intended effect of Hebbian learning. Specifically, the general rule is that Hebbian learning must increase the basin of attraction for the current configuration – otherwise it would not create a ‘memory’ for the current configuration in the system’s dynamics. Under Eq.4, Hebbian learning must be reversed to achieve this (i.e. $\delta < 0$). The syntactic definition of Hebbian learning is therefore unfortunately misleading in this respect. But the important observation for our purposes is that selfish changes to connections will *always* increase the utility of the current state configuration (by definition of a selfish change), and will therefore increase the basin of attraction for this configuration, regardless of the form of utility function used. This is the dynamical consequence that matters for the intended semantics of Hebbian learning. That is, by definition, a selfish change to a connection amplifies the benefit of the current states – and this will always increase the likelihood that this configuration of behaviours is visited again in future. Intuitively, this just means that selfish changes to connections ‘exploit’ the current configuration of behaviours by amplifying their benefit rather than ‘exploring’ or inducing new/different state configurations.

⁵ We thank Hywell Williams for this observation.