Industrial and Corporate Change, Volume 12, Number 4, pp. 843–877

Economic production as chemistry

John F. Padgett, Doowan Lee and Nick Collier

Inspired by the hypercycle model of the origins of chemical life on earth, this paper develops an autocatalytic model of the co-evolution of economic production and economic firms, represented as skills. The production and distribution of goods by firms are only half of what is accomplished in markets. Firms are produced and transformed, via learning, through goods passing through them. Through means of both agent-based and analytic modeling, this article establishes three principles of social organization that provide sufficient foundations for the unconscious evolution of technological complexity: structured topology, altruistic learning and stigmergy.

1. Introduction

The production and distribution of goods by firms are only half of what is accomplished in markets. Firms also are produced and transformed through goods passing through them. This transformation is not just a matter of profits. Skills and the core competencies that define firms are developed and maintained through 'learning by doing' and other learning processes that are triggered by exchange among firms. In periods of decentralization and outsourcing, like today, it is more evident than ever that linked chains of skills are distributed across firms. In this context especially, evolution in and learning of distributed skill sets reverberates directly into the reconstitution of firms. Evolving links among firms, in turn, guide and shape the recombinant new-product possibilities latent in distributed skill sets.

The duality of this co-evolution between product and organization is often ignored, as analysts assume away one side of the dynamics in order to focus attention on the other. A number of economists, including many who publish in this journal, are aware of the issue of co-evolution (e.g. March and Simon, 1958; Nelson and Winter, 1982; Hughes, 1983, 1987; Dosi *et al.*, 1992, 2000; Malerba and Orsinigo, 1993; Nelson, 1994, 1995; Warlien, 1995; Landesmann and Scazzieri, 1996; Powell, 1996; McKelvey, 1997; Antonelli and Marchionatti, 1998; Coriat and Dosi, 1998; Rosenkopf and Tushman, 1998). However, more tools are needed to help to analyze the nonlinear and path-dependent dynamics of feedback among evolving networks that such processes entail.

One place to turn for analytic inspiration is chemistry. From the chemical perspective, life is an interacting ensemble of chemicals that reproduces itself through time, in the face of turnover of its parts.[1] Biological organisms are not fixed entities; they are autocatalytic networks of chemical transformations, which continually reconstruct both themselves and their physical containers. The origin-of-life problem, under this view, is how such an ensemble can self-organize, from a ‘soup’ of random chemicals in interaction and flux.

This chemical perspective can be applied to the analysis of co-evolution of products and firms through the following analogy: skills, like chemical reactions, are rules that transform products into other products. Products, like chemicals, are transformed by skills. Firms, like organisms, are containers of skills that transform products. Trade, like food, passes transformed products around through exchange networks, renewing skills and thereby firms in the process. In the macroeconomic aggregate, product inputs flow into, and outputs flow out of, this trading network of firms and skills.

Firms in this view are sites through which a distributed ‘chemical reaction’ production process flows. At minimum, firms can be considered to be mere collection bins for diverse skills. Trading among firms regulates both the activation and the evolution of skill sets distributed across firms. Composition of skills within firms evolves, among other methods, through learning-by-doing: the more a skill is used, the more the skill is reinforced. Skills not used are forgotten. These two processes of learning and forgetting impose selection pressure on an evolving network-of-skills-through-firms production system. The ‘origin-of-life’ problem for markets is to discover how a randomly distributed set of skills across firms can self-organize, through exchange, into a coherent product-transformation network,[2] which then reproduces itself through time and ‘grows’ a set of firms to sustain it.

Inspired by a specific literature in chemistry, that on hypercycles, in this paper we develop one family of economic production models that operationalizes this co-evolutionary perspective on markets. Extensions beyond the hypercycle framework will be discussed at the end of this paper.

The ‘hypercycle’ is a specific model of the chemical origin of life pioneered by Eigen (1971) and Eigen and Schuster (1979), and extended by others (e.g. Kauffman, 1986, 1993; Hofbauer and Sigmund, 1988; Fontana and Buss, 1994; full literature reviewed in Stadler and Stadler, 2002). From random distributions of chemicals, the hypercycle model seeks to find and to grow sets of chemical transformations that include self-reinforcing loops: $\{1\rightarrow2, 2\rightarrow3, 3\rightarrow4, \ldots, n\rightarrow1\}$. Chemical cycles are crucial to the issue of life because they are the motors behind the self-reproduction of metabolic networks, in the face of continuous turnover in component chemicals. Without cycles, there is no positive feedback for growth; without them, any chemical reaction left to

[1]From the physics and biological points of view, additional criteria to the definition of life are sometimes added. Physicists (e.g. Prigogine and Glansdorff, 1971) sometimes add the criterion of far-from-equilibrium throughput of energy. Biologists (e.g. Maturana and Varela, 1980) sometimes add the criterion of permeable encapsulation.

[2]Such a network could be called a ‘metabolism’ or a ‘technology’, depending upon the application context.

itself will stop or 'die'. Eigen and Schuster, Hofbauer and Sigmund, and others have explored how variation in reaction rates, in density, and in number of components affected the dynamic stability or 'survivability' of various classes of hypercyclic chemical reactions, within a well-stirred liquid reaction tank. Boerlijst and Hogeweg (1991) and Padgett (1997) extended the investigation beyond the original liquid context to a spatial topology of interaction.

Viewing economics as chemistry entails extraordinarily minimalist assumptions about economic production: firms become nothing more than bins of transformation rules; products randomly flow in and through these bins, without purpose; rules reproduce or die only as functions of use. There is no guiding intelligence, either at the level of the market or at the level of the firm.[3] In such a minimalist setup, the analytic question is: can any coherent and self-reproducing systems of production (that is, coevolved sets of products and firms) emerge? And if they can, what mechanisms affect the likelihood of such emergence? *A priori* one might not expect much complex economic organization to be possible from randomly iterated rules. Yet the history of chemical and biological life on earth suggests that minimalist systems can generate astounding complexity under the right circumstances. Intelligence, we speculate, may not have been necessary for markets to emerge.[4] We are not arguing thereby that humans are no more complicated than chemicals. We are arguing that a surprising amount of social and economic organization does not depend on humans being complicated.

2. The hypercycle model of economic production

We shall describe our hypercycle model of economic production in pseudo-algorithmic fashion, because we have implemented it in the form of an agent-based simulation.[5] First we shall describe our core models of production and learning. These will give the logic of our basic 'dependent variable': hypercycle emergence. Then we shall describe experimental variations of our core model—number of products, interaction topology, mode of learning, input environment, and input search. These are the 'independent variables' that may affect the likelihood of hypercycle emergence. The simplest versions of our spatial hypercycle model can be solved analytically. We

[3]This is not only bounded rationality, this is the absence of consciousness altogether.

[4]Hayek (1948) made a similar 'self-organizing' argument about the operation of markets that we are making about the emergence of markets.

[5]Our agent-based model is publicly available for both demonstration and modification, and can be found on the web site http://repast.sourceforge.net under the application module HYPERCYCLE. Repast is a comprehensive software framework and library for creating agent-based simulations, built in the Java language. It was developed at the Social Science Research Computing Center at the University of Chicago.

present such solutions below, both to verify and to aid interpretation of the simulation results.[6]

2.1 Core model of production

1. There are three components in the model: rules ('skills'), balls ('products'), and bins ('firms').
2. Rules/skills transform balls/products into other balls/products. For example, if balls/products are indexed by $i = 1, 2, 3, \ldots, n$, then the set of transformation rules obeying a cyclic structure[7] would be represented as $\{1\to2, 2\to3, 3\to4, \ldots, n\to1\}$. The set of transformation rules we call a 'technology'; n indexes the 'complexity' of the technology.
3. Rules/skills are contained in bins/firms. At the beginning of each simulation run, skills are just randomly distributed across available firms, without any logic. The number of firms initially is large.
4. Bins/firms are arrayed on a spatial grid, with wrap-around boundaries. Each firm has eight possible nearest-neighbor trading partners.
5. At each asynchronous iteration of the model, a random rule is chosen 'looking for action'. The firm containing that rule/skill reaches into the input environment (modeled as an urn) and draws an input ball/product. If the input ball/product selected is compatible with that rule, then the ball/product is transformed according to that rule. (For example, if a firm possessed an activated '1→2' skill, and it drew a '1' as input from the urn environment, then it would transform the input '1' into the output '2'.) If the ball/product selected could not be processed by the activated rule, then input ball/product passes through the firm into the output environment (also modeled as an urn) unchanged.
6. Products successfully transformed within the firm are passed randomly to one of the firm's eight possible trading partners. If that trading partner possesses a compatible skill, then it transforms the product further, and passes that along in a random direction. (For example, if the second firm possessed a '2→3', then after receiving the output '2' from the first firm, it would transform the '2' into a '3', and then pass that

[6]In economics, though not in physics, there is frequently a fruitless methodological debate about agent-based modeling versus analytic modeling. Our position is that one can and should do both—namely, solve simple settings analytically and then scale up through computer modeling. Analytic solutions are more transparent than computer simulations, but frequently require the imposition of highly restrictive and unrealistic homogeneity assumptions. Computers can numerically solve highly non-linear models with heterogeneous agents in non-homogeneous topologies, and there is no reason not to let them do so as long as one can understand the results.

[7]Rule sets more general than the cyclic structure are of course quite possible to set up and explore. In this paper we restrict ourselves to the cyclic structure in order to root ourselves in the pre-existing literature on hypercycles. In future papers, we intend to explore other rule sets. The HYPERCYCLE code has already been written with this extension in mind.

on to a third firm or possibly back to the first.) In this way, transformed products pass through sequences or chains of skills.

7. Bins/firms continue passing around transformed products among themselves until the product lands on a firm that does not possess a compatible skill to transform it further. At that point the product is ejected into the output environment. And a new input ball is selected to begin the iterative process all over again.

Overall, the production process looks like this: input balls/products come in from an input environment, then pass in random directions through randomly distributed production chains of skills, being transformed *en route*, until they pass back out into an output environment. For this random production process to self-organize into coherence, there must be some sort of a feedback mechanism. For us, this is learning through trade.

2.2 Core model of learning

1. 'Learning by doing' is modeled in chemical fashion as follows: if one skill transforms a product and then passes it on to another transforming skill, then a skill is reproduced. We call such a sequence a 'successful transaction', since both sides transform products.[8] Which of the two skills is reproduced in a successful transaction is an experimental variation within the model, to be discussed below.
2. 'Forgetting' is modeled in chemical fashion as follows: whenever one skill reproduces anywhere in the system, another skill, randomly chosen from the overall population of skills, is killed off. The total population volume of skills in the population thereby is held constant.[9]
3. Once a firm loses all its skills, it 'goes bankrupt' or 'dies', never to recover any skills.

Learning by firms is equivalent here to reproduction of their skills. Learning by firms and reproduction of skills, we argue, are the same process, just described at different levels of analysis. This is like a germ's eye view of disease: instead of focusing on the organism getting sick, we focus instead on the reproduction and spread of germs. Firms learn and adapt in our model, but the underlying mechanism is not conscious reasoning. Rather it is the reproduction of their inherited skills through use.[10] Firms are kept alive or are killed off solely through the 'chemical' reactions of skills that operate through them.

This combination of learning, forgetting and dying imposes selection pressure on the production system of skills. In the face of inexorable forgetting, skills must

[8]Final consumption is the output urn.

[9]This conservation-of-skills assumption mimics the conservation-of-mass assumption in chemistry. While perhaps too harsh an assumption for many human populations, this constraint is one chemistry-style way to model competition among firms.

[10]In future extensions of this model, we intend to add diffusion of skills among trading firms, in order to mimic 'collaborative dialogue'. But that extension is not developed in this paper.

reproduce in order to survive. In the harsh conservation-of-skills setup employed here, indeed, the very success of rules in one place in the system imposes sharply competitive selection pressure on rules elsewhere in the system. Heavily used subsets of the distributed skill set reproduce, and rarely used subsets of the distributed skill set disappear. The death of a firm is an absorbing state that permanently eliminates its unsuccessful skills.[11] As the skill composition of rules within firms thereby evolves, surviving firms cluster into mutually reinforcing trading groups, reminiscent of Marshallian industrial districts. And production chains of compatibly sequenced rules self-organize their way through these spatially contiguous groups of firms.

A conscious desire to cooperate, indeed consciousness at all, is not necessary for mutually reinforcing clusters of trading firms to emerge and to survive. In this model, the minimal requirement for long-term survival, both of firms and of clusters, is to participate in at least one spatially distributed production chain that closes in on itself, to form a loop. Not all production chains within a trading cluster need be closed into loops. And more than one loop within a cluster is possible, in which case we may have a dense hypercyclic network of spatially distributed production rules. But loops within distributed chains of skill are crucial, not for production itself, but for the competitive reproduction of skills. Loops set up positive feedbacks of growth in skills that give firms that participate in them the reproductive ability to out-produce firms that do not. Put another way, clusters of firms can produce products with or without hypercycles, but firms whose skill sets participate in production chains that include loops have the further capacity to keep renewing each other through time. This is the chemical definition of life.

From our chemical perspective, therefore, the secret to understanding competitive success, both of firms and of industrial districts, is to find the conditions that foster the spontaneous self-organization of skills into self-reinforcing hypercyclic production chains, which wend their way through firms, knitting them together in trade and helping them to reproduce each other through continuous learning.

2.3 Experimental variations

There are five 'independent variables'—that is, experimental treatments in the simulation model—whose effect on the likelihood of finding and sustaining self-organized hypercycles of skills will be explored in this paper.

1. Complexity. A parametrically fixed volume of rules or skills is scattered randomly around the space of firms at the beginning of each run. In this paper, there are 200 rules being scattered. We vary the composition or 'complexity' of the rule set so scattered. In a cyclic structure of rules, complexity is indexed by *n*. We shall vary *n* from 2 to 9: that is, we shall explore 2-skill hypercycles, 3-skill hypercycles, and so forth, up to 9-skill

[11]Allowing the entry of new firms is another obvious extension to our model that we do not explore here.

hypercycles. Presumably the more complex the rule set, the more difficult it will be to find and to sustain hypercyclic production chains.

2. Interaction topology. The basic spatial topology for trading to be explored in this paper is the 10×10 wraparound grid. That is, at the beginning of each run, there are 100 firms, one firm per cell in the grid, each of which can trade products with their eight nearest neighbors. This is the so-called Moore-neighborhood topology.[12]

As experimental variation, we shall compare the hypercycle behavior of this spatial topology to that of the non-spatial 'well-stirred liquid reactor' topology, more traditional in chemistry. In non-spatial or random topology every rule is equally likely to pass a product to any other surviving rule, irrespective of spatial/firm location.

A major finding in the existing hypercycle literature (Hofbauer and Sigmund, 1988: 96) is that non-spatial hypercycles are dynamically stable up to 4-elements, but not beyond that. In other words, in non-spatial interaction when hypercyclic sets are 5-elements and up, one or more of the component chemicals is always driven to zero during the reaction process, thereby breaking the reproductive loop and causing the hypercycle to 'crash'. This is a 'complexity barrier' that self-organizing hypercycles, and hence 'life', cannot penetrate when chemical interaction is non-spatial or random. Padgett (1997) has shown that in spatial interaction topologies, dynamically stable hypercycles of complexity 5-elements and above can be grown, albeit at increasingly lower frequencies at higher levels of complexity. Spatial interaction, in other words, can break the complexity barrier. Presumably this is one reason why complicated chemical life is embodied. We shall reconfirm both the Hofbauer and Sigmund (1988) non-spatial findings and the Padgett (1997) spatial findings here, in a new context.[13]

3. Learning/reproduction. In the spatial topology setting, there are two variants of 'learning by doing' that can and will be explored:[14]

1. 'Source reproduction' is where the originating rule in a successful transaction is reproduced.
2. 'Target reproduction' is where the receiving rule in a successful transaction is reproduced.

[12]In future work we plan to investigate additional topologies as well. Padgett (1997) used 4-neighbor (von Neumann) neighborhoods. The impact of social networks of various kinds, such as cliques and small worlds, is an especially important avenue to explore.

[13]The models in this paper are extensions of the model presented in Padgett (1997). The main extension is to add explicit products that are being transformed. In the earlier paper, there were action–reaction chains of 'play', but nothing was actually produced or accomplished. We believe that the setup in Padgett (1997) was appropriate to the emergence of informal organization among people within a firm, whereas this setup here is more appropriate to trading among firms in an economy.

[14]In non-spatial interaction, these two reproduction modes behave identically (see the appendix). Space is what separates target from source. In Padgett (1997), a third mode was also explored: 'joint reproduction', where both rules in a successful transaction reproduce. Because two rules are reproduced in this hybrid, two offsetting skills need to be killed off to preserve conservation-of-mass.

For example, if (1→2) receives a 1 from the input environment, transforms it into a 2, and then successfully passes that 2 onto a neighboring (2→3), who transforms it again, then 'source reproduction' is where the initiating (1→2) reproduces, and 'target reproduction' is where the recipient (2→3) reproduces.[15] Variation in mode of reproduction thus defines who benefits from the transaction.

We think of source reproduction as 'selfish learning', because the initiator of the successful transaction reaps the reward (like a teacher). And we think of target reproduction as 'altruistic learning', because the recipient of the successful transaction reaps the reward (like a student). 'Selfish' and 'altruistic' are verbal labels that accurately characterize who benefits. In using these suggestive labels, however, one should avoid importing motivational connotations. In the minimalist models developed here, there are no motivations—just actions and reactions, like in chemistry.

Padgett (1997) demonstrated that, in comparison with source reproduction, target reproduction dramatically increases the likelihood of growing stable hypercycles. And it also increases the spatial extensiveness and complexity of the firm cluster that hypercycles produce. Both of these findings will be reconfirmed here.

In addition to these three experimental manipulations, two more experiments will be performed here, which vary the input environment in which hypercycles grow. Such additional experiments were not possible in Padgett (1997), because previously there was no explicit modeling of products or of product environments.

4. Input environment. Input environments of resources or products can be conceived as fixed or as variable, and they can be conceived as rich or as poor.

Among fixed resource environments,

1. 'Rich' input environments will be modeled by letting the input urn of resources contain all possible inputs, never to be depleted even as products/resources are withdrawn.
2. 'Poor' input environments will be modeled by letting the input urn of resources contain only one possible input (by convention, we call that '1'), not depleted even as products/resources are withdrawn.

Among variable resource environments,

3. 'Endogenous' input environments will be modeled by letting the input urn be constructed over time by the outputs of the production system. Under the endogenous-environment variant, in other words, our model will withdraw one input product, transform it into other products through distributed production chains, and then place the final output back into the original input urn.

At initialization, we set the endogenous input environment to be equal to the poor resource environment (that is, 'all 1s'), in order see if and when the economy can by

[15] Of course the recipient (2→3) could easily turn into an initiator in the next tick, if a neighboring (3→4) is subsequently found.

itself transform an initially poor environment into one more supportive of its own life. Structurally, therefore, the endogenous resource environment migrates between the fixed-poor resource environment and the fixed-rich resource environment.

Presumably, rich input environments are more congenial to hypercycle emergence than are poor environments. What is less clear *a priori* is the relative ranking of endogenous environments. Given that we have defined ‘rich’ virtually as nirvana (namely as ‘all possible inputs available all the time, never to be depleted’), our expectation is that nothing can outperform that. However, modelers of social insect behavior (e.g. Camazine *et al.*, 2001) have discovered that ‘stigmergy’—the ability of social insects to transform their physical environments into nests, mounds, paths, and the like—can sometimes exert surprisingly powerful feedback onto the development of social organization itself. The open question, therefore, is whether the social organization achievable in the endogenous environment is superior in any way to that achievable in the rich environment.

5. Input search. The final experimental manipulation we shall perform varies the precision of search through the environment:

1. ‘Random search’ is when an activated rule reaches into the input environmental urn and chooses inputs randomly, in proportion to what is there.
2. ‘Selective search’ is when an activated rule reaches into the input environmental urn and selects the exact input it needs to transform, if it is there.

Random search is like literal chemistry.[16] Selective search is more like animal behavior.[17] This is the only place in the model where we vary degree of intelligence. We expect the more intelligent selective-search procedure to outperform the stupid random-search procedure in finding and nurturing production hypercycles.

In the appendix, we lay out mathematically the operationalization of all these experimental variations, for the very simple special case of two skills and two firms.

To sum up the logic of our modeling enterprise: chemistry teaches us that life is an ensemble of products and transformation rules that reproduces itself through time. Distributed economic production activity qualifies under this minimal definition, especially if the fragility and malleability of both firms and their skills are recognized. We take the analogy between economic and chemical exchange to its extreme by assuming away all human rationality and even consciousness, holding on only to the features of blind adaptive learning and selection. We hope to demonstrate that firms, production chains, and industrial districts can emerge and reproduce even under these minimalist assumptions. And we hope to discover the structural and interactive imperatives that help to foster economic self-organization.

[16] Metaphorically we think of this as ‘the intelligence of an atom, bouncing around’.

[17] Metaphorically we think of this as ‘the intelligence of a cow, looking for grass’.

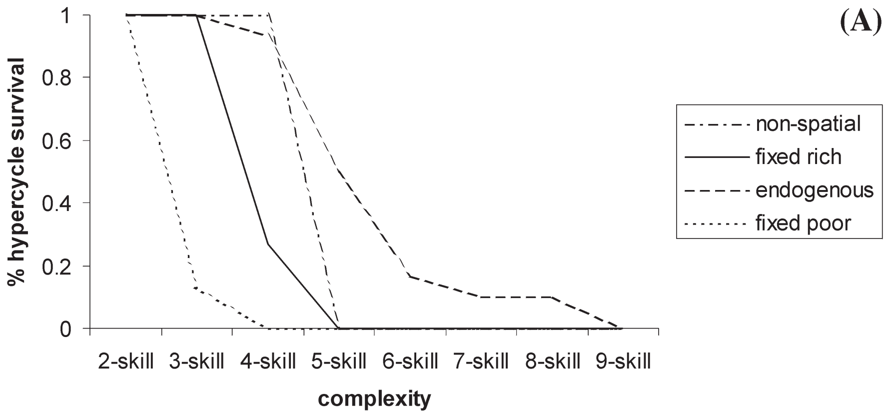

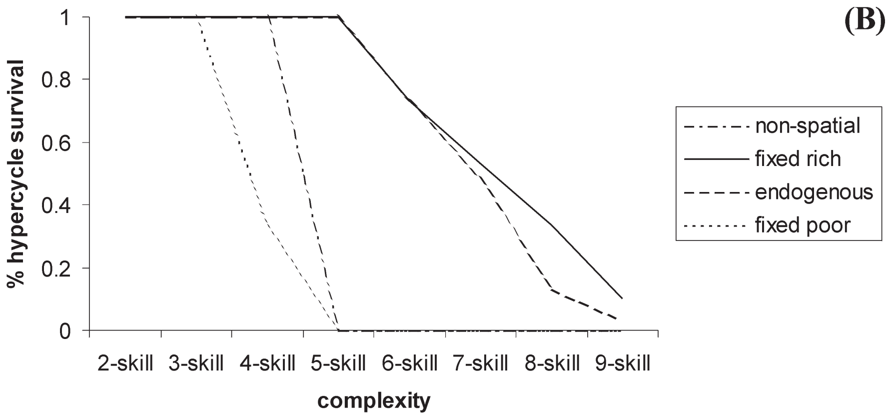

Figure 1 (a) Source reproduction and selective search. (b) Target reproduction and selective search. Each point is an average of 30 simulation runs.

3. Results: hypercycle emergence

The basic findings from our agent-based simulation model of hypercycle economies are presented in Figures 1 and 2. These figures present on their y-axes our dependent variable: long-term[18] probability[19] of hypercycle survival. They present on their x-axes

[18]Our operational definition of 'long-term' came inductively from observing many, many runs, and how long even the slowest among them took to converge to equilibrium. We finally settled on the following as liberal stopping points for our simulations: 30 000 ticks for 2-element hypercycles; 40 000 ticks for 3-element hypercycles; 60 000 ticks for 4-element hypercycles; 80,000 ticks for 5-element hypercycles; 120 000 ticks for 6-element hypercycles; 180 000 ticks for 7-element hypercycles; 250 000 ticks for 8-element hypercycles; and 300 000 ticks for 9-element hypercycles. For the more complex of these hypercycles, much computing time was required.

[19]Each of the points in these graphs represents the average of 30 simulation runs.

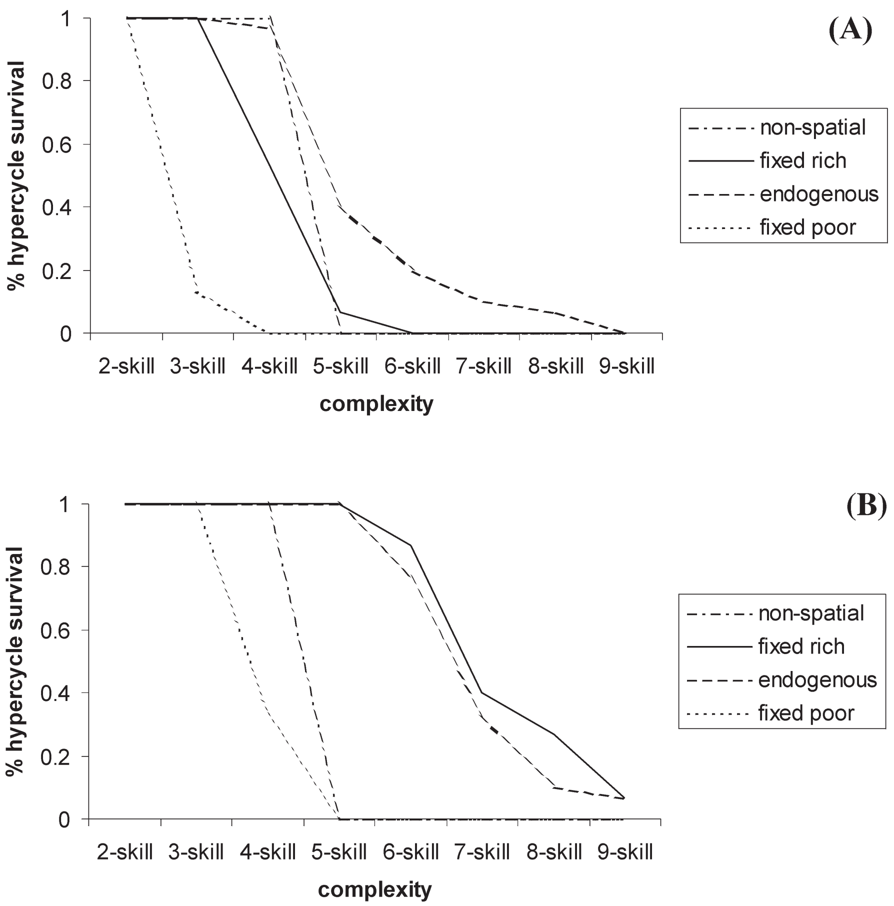

Figure 2 (a) Source reproduction and random search. (b) Target reproduction and random search. Each point is an average of 30 simulation runs.

varying degrees of complexity in the simulated economies' technologies: simple 2-skill technologies, slightly more complicated 3-skill technologies, and so forth, up to our most complex 9-skill technologies. Different lines within these graphs present the results of our various experimental manipulations: interaction topology, mode of reproduction/learning, and input environments. Figure 1 presents these comparisons for selective search; Figure 2 presents these comparisons for random search.

Much information is contained in these graphs, so we shall unpack the findings one independent variable at a time.

3.1 The effect of spatial topology

As Hofbauer and Sigmund (1988) have shown analytically, and as we have already mentioned, non-spatial hypercycles face a dynamic 'complexity barrier' at the level of 5-elements and above. In the non-spatial or 'liquid' topology of random interaction, where there are no firms, the volumes of the various reproducing skills undergo accelerating oscillations under a hypercycle regime with complex rule sets, until eventually one skill is driven to zero, thereby breaking the reproductive loop and causing the overall hypercycle to 'crash'. This finding is reconfirmed in our simulations, it being displayed graphically by the fact that hypercycle survival rates abruptly plummet from 100% to 0% in the non-spatial portion of all of our figures, as complexity passes the threshold from 4-skills to 5-skills.

In sharp contrast to this dynamic instability among 5+ skills, once spatial constraints on interaction are introduced—that is, once firms with delimited trading patterns are permitted—then higher complexity in skill sets becomes dynamically possible (albeit not 100% of the time). This finding is illustrated graphically by the fact that, for complexity five skills and above, survival rates of spatial hypercycles are equal to or are superior to survival rates of non-spatial hypercycles.[20]

Another way of expressing these findings is this: non-spatial 'freedom of trade' of every skill with every other, with no firms to channel that trade, generates so much volatility in skill reproduction that complexity becomes dynamically unsustainable. The opposite extreme—complete internalization of all skills within a single firm—eliminates entirely the trade that renews learning. Skills spatially dispersed through clusters of firms are necessary (but not sufficient) in our model for economic production markets to be sustainably complex. Simple economies, with four or fewer products, do not need firms or spatial clusters of firms to reproduce. But complicated economies, with five or more products, do.

To understand why this is the case, let us first show representative pictures of equilibrium outcomes for successful spatially embedded hypercycles, and then give a verbal description of how a typical run develops.[21] Figure 3 presents some 5-skill hypercycles produced by target reproduction, under a variety of environmental and search conditions. Ellipses are firms; number pairs are skills; arrows are routes through which products can successfully be transformed. When arrows are solid, they are part of one or more cyclic loops of arrows—that is, part of a hypercycle. The total number of distinct cycles[22] and the volume of rules per firm are also reported in these diagrams. In all successful 'no crash' runs, final convergence was to a single extended cluster of firms, through which ran one or usually more hypercycles of sequenced skills.

As is evident in the pictures by the dotted arrows, even in long-term equilibrium not all surviving firms or skills participate in hypercycles. 'Parasites'—namely, firms and

[20]We exclude the expected terrible performance of the fixed-poor environment from this statement.

[21]Readers can visually see for themselves at http://repast.sourceforge.net, under the application module HYPERCYCLE.

[22]One cycle is distinct from another if at least one node in the loop is different. Usually, as is illustrated in Figure 3, multiple hypercycles are produced by our model and overlap into a dense hypercyclic network. Such redundancy of 'metabolic pathways' gives resilience to a cluster in the face of selection pressure, which is why hypercycles overlap in the first place. Redundancy is sometimes viewed as a mystery by those who equate natural selection with maximal efficiency.

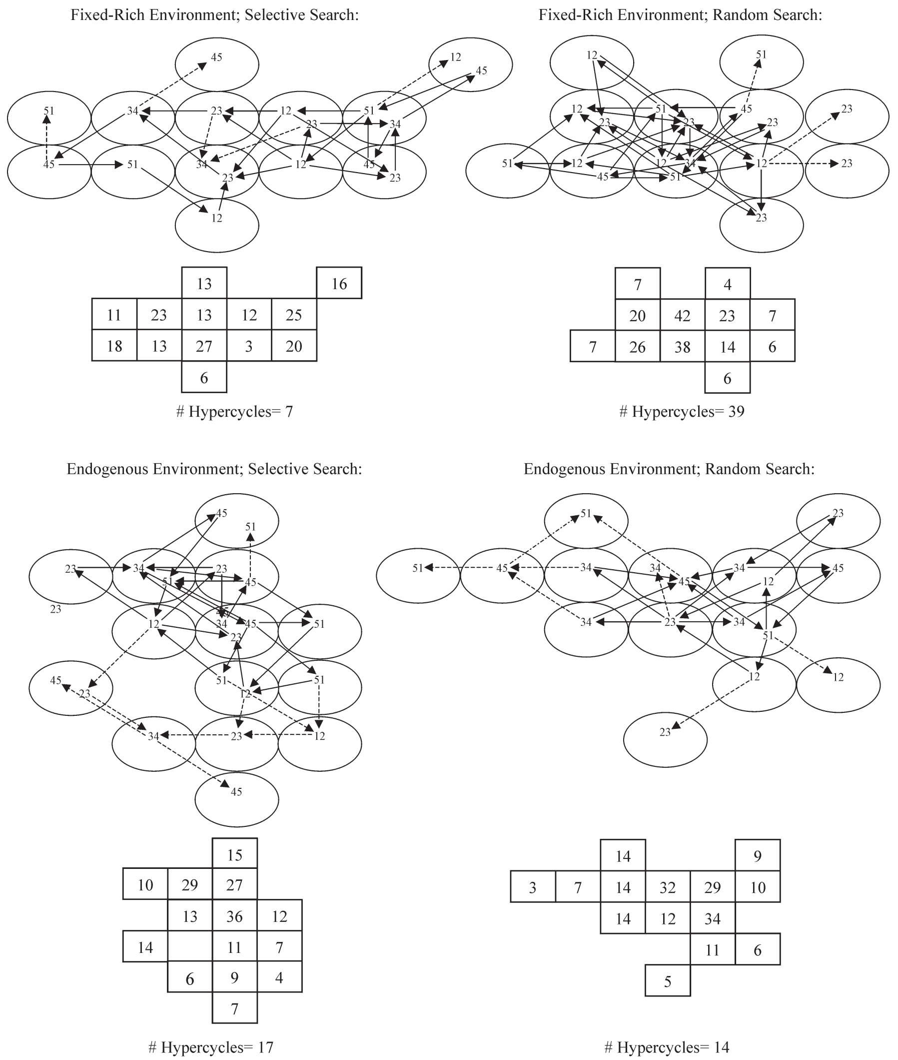

Figure 3 Representative 5-skill hypercycles at equilibrium: target reproduction. Ellipses are firms; within ellipses, number pairs are skills; products flow along arrows. Solid arrows participate in hypercycles. Dotted arrows link to parasite rules. Boxes give the total volume of rules contained within corresponding ellipses.

skills who do not themselves help to keep others alive, but simply 'free ride' on the reproductive work of others—are both possible and common in our model. In chemistry, unlike in rational choice theory, free riders do not necessarily threaten successful cooperation in the core, if the core is strong enough.[23] In a subsequent

[23]'Strong enough core' means overlapping hypercycles. A single hypercycle, while possibly stable in and of itself, is more vulnerable to parasites than are overlapping hypercycles that reinforce each other.

section on hypercycle structure near the end of this paper, we shall report the percentage parasites produced under various conditions.

The way these hypercyclic spatial clusters emerge is usually as follows: initially, as explained above, skills are just randomly scattered over the grid of firms. Given the parameter values used here—200 rules and 100 firms—this random distribution is also a fairly thin distribution, albeit one with slight inhomogeneities. At the random outset, there are hundreds or thousands of discrete spatially distributed cycles, some quite convoluted indeed, just because of the combinatorics of random trading among 100 firms. As reproduction and selection processes kick in, however, skills and firms start to die, leaving surviving firms, which trade heavily amongst each other, to congeal throughout the grid.[24] These multiple spatial clusters then begin to compete with one another in the reproduction of their skills: each continuing to grow as long as a hypercycle is contained within it, but then decaying once interior hypercycles have crashed. Eventually only one cluster remains.

This surviving hypercyclic cluster may or may not itself be dynamically stable over the long term. Even after 'victory', in other words, a hypercycle can (and often does) crash; 'survival of the fittest' may or may not be good enough for 'life'.

The observed process reminds one generically of the origin of the universe. As physical space expanded, an initially uniform distribution of particles developed very slight inhomogeneities, which then positively fed back on one other through gravity until galaxies and stars were produced.[25] Expanding space 'broke the symmetry' of random 'liquid' interaction, creating localized inhomogeneities that grew.

In our case, the primary feedback reason for observed spatial clustering is 'memory'—namely, the inscription of the past into the structure of the firm. Successful transactions reproduce compatible skills that are located in neighboring firms. This creates the following positive feedback among compatible-skilled neighboring firms: the more the skills are activated, the more they reproduce, the more they are activated, the more they reproduce. The volume of skills in a firm at a given point in time thereby becomes the cumulative history of past interaction with its compatible neighbors. Meanwhile, the counterbalancing forgetting process kills off skills randomly in the population. Firms not participating in positive feedbacks among neighbors eventually get cleaned out of their initial skills. Long-term success or failure of a given firm is the path-dependent result not only of that firm's own history, but also of that firm's neighbors' histories.

None of this would have been true without physical space, or some social-network functional equivalent to space (cf. Cohen *et al.*, 2001). Spatial or other constraint on interaction 'breaks the symmetry' of firms potentially trading with all other firms, and allows localized skill inhomogeneities to form. Positive feedback through continued

[24]In early stages when all clusters are thin and extended, these multiple clusters overlap to such an extent that to call them 'multiple' is poetic license.

[25]Of course, we do not wish to imply by this simple statement that such feedback was linear or smooth.

trading then inscribes the memory of past interactive success into the structure of each co-adapting firm.

A second mechanism behind spatial clustering, not present in Padgett (1997), is chaining. The physical act of passing products around orchestrates sequences of learning. Not only do compatible neighbors have positive feedback loops in their own growths, but also compatible neighbors trigger other compatible neighbors. Thus feedback loops are evoked more efficiently once hypercyclic clusters begin to emerge. Perhaps this is one evolutionary reason for why artifacts, either physical or symbolic, are helpful for humans learning in groups (cf. Hutchins, 1995). The mere act of passing around transformed products, even when purposeless, coordinates learning sequences of humans through chaining.

3.2 The effect of reproduction/learning mode

Embedding production and trading in physical (or social) space has a second non-obvious consequence: it induces an asymmetry between target and source reproduction.[26] There is no difference between 'selfish' and 'altruistic' when there are no firms, or other types of organisms, to begin with.

In the production and nurture of hypercycles, target reproduction is superior to source reproduction. This is shown in Figures 1 and 2 by the facts that in all cases where target and source differ in the first place the survival plots of target reproduction are displaced to the right of the corresponding survival plots of source reproduction. Rephrasing this finding at a different level of analysis, spatial hypercycles of whatever complexity are easier to grow when learning by firms is altruistic than when it is selfish.

As explained in Padgett (1997), the basic reason for this superiority of target reproduction, or altruistic learning, is repair. Target reproduction combats dynamic instability in a way that source reproduction does not. The basic process of dynamic instability, causing hypercycles to crash, is that as one skill reproduces rapidly, under competition other skills are driven to zero, thereby breaking the reproductive loop of skills. Spatial topology distributes this dynamic into an overlapping series of neighborhoods, thereby inducing local heterogeneity. This opens the door for localized co-adaptive feedbacks to operate.[27] But source reproduction, or selfish learning, does not really attack the basic dynamic instability itself. Source reproduction is this: an initial activated rule passes on its transformed product to a neighboring compatible rule, causing the original activated rule to reproduce. Under source reproduction, frequently activated rules reproduce more frequently, often eventually driving out of business even compatible neighbors on whom they depend for their own survival. As we shall see in the next section, other factors like endogenous environment can

[26]As shown in the appendix, in non-spatial topology target and source reproduction become identical processes.

[27]To repeat: spatial topology is necessary but not sufficient for complex hypercycle emergence.

sometimes moderate this destructive dynamic, but source reproduction in and of itself does not eliminate the non-spatial instability problem.

In contrast, target reproduction is this: an initial activated rule passes on its transformed product to a neighboring compatible rule, causing the recipient rule to reproduce. Here the more frequently the initial rule is activated the more frequently the second recipient rule reproduces. In this way, a hypercycle can repair itself: as the volume of one skill in a loop gets low, high volumes of compatible skills in neighboring firms reach in to that low-volume skill to build it back up. Peaks and valleys along loops are smoothed.

This simulates altruistic behavior, although here of course no skill or firm is 'trying to preserve' the hypercycle. Target-reproduction repair does not guarantee that a hypercycle will survive, but it does directly alleviate the dynamic instability problem that afflicts both the non-spatial and the spatial-source settings.

This repair mechanism can also be shown analytically for one special case of our model. In the appendix we derive differential equations of skill growth for the extremely simple 2-skill-hypercycle setting of a single dyad: two adjacent firms trading only with each other. Box 1 collates these differential-equation results for ease of inspection. In our agent-based simulations interlinked dyads are proliferated across the entire grid, generating interaction effects not captured in the stripped-down dyadic setting. Simplification, however, permits analytic solutions not otherwise possible. Such solutions are useful both to increase transparency and to double-check computer code.

The analytic contrast between target reproduction and source reproduction is most sharply and cleanly illustrated in the setting of fixed-rich environment. There, in both of the target-reproduction equations, $E(n_{12,t+1})$ always goes up when $n_{12,\mathrm{t}} < n_{21,t}$, and $E(n_{12,t+1})$ always goes down when $n_{12,t} > n_{21,t}$. The converses are true for $E(n_{21,t+1})$. In other words, target reproduction generates a consistent tendency toward homeostatic stability, over the entire range of n_{12}. In sharp contrast, in both of the corresponding source-reproduction equations, both $E(n_{12,t+1})$ and $E(n_{21,t+1})$ equal zero. Source reproduction exhibits no built-in tendency toward homeostatic stability: n_{12} drifts in random-walk fashion until eventually it crashes into the absorbing states of either $n_{12} = 0$ or $n_{12} = 1$.

For hypercycles more complex than the 2-element dyad, we can no longer derive solutions analytically. But the simulations show this basic dyadic finding to be true more generally. Target reproduction generates higher rates of hypercycle survival than does source reproduction for all corresponding spatial settings. To repeat, the mechanism generating this sizable superiority is direct 'altruistic' repair of complementary rules by each other. Target reproduction repairs hypercycles without intending to do so, once given the precondition of spatial 'symmetry breaking', which induces the distinction between altruistic and selfish in the first place.

3.3 The effect of input environment

Figures 1 and 2 also point to the existence of a second repair mechanism, more relevant to source ('selfish') reproduction than to target ('altruistic') reproduction.

In the learning mode of target reproduction, hypercycles survived at almost as high rates under the endogenous environment as they did under the fixed-rich environment. This result fit with our expectations, which were that, although no environment could outperform 'nirvana', production hypercycles might be able effectively to build their way out of a poor environment into one more supportive of themselves. The slight inferiority of endogenous was due only to the fact that sometimes they were not quick enough to do so.[28] But usually, they were impressively quick.

The same comparison for source reproduction, however, generated quite a surprise: hypercycle survival rates under endogenous environments actually were superior to fixed-rich ('nirvana') environments, and by a substantial amount. It was not true that the boost given to source reproduction by endogenous environment lifted it to the level of target reproduction. But it came close.

After some detective work,[29] we discovered the reason behind this surprising phenomenon. As explained above, source reproduction left to itself generates self-destructive ('cancerous') growth. Namely, the more skills are activated, the more they reproduce, the more they are activated, etc., until the neighboring partner upon whom that skill depends is destroyed. For complex hypercycles, this is fatal. Endogenous environments do not eliminate this cancerous growth, but they help to control it. The mechanism is this: The more a skill reproduces and hence is activated, the more it consumes its compatible input in the environment, and transforms it into something else. The environment of compatible resources for low-volume skills thereby is enriched, while the environment of compatible resources for high-volume skills is starved. This does not completely smooth the peaks and valleys around the hypercycle, as does the more direct skills-to-skills intervention of target reproduction.[30] But the indirect skills-to-environment-to-skills method of regulation, induced by environmental endogeneity, does function to keep skill-volume peaks within bounds.

This reminds one of 'stigmergy' among social insects (e.g. Camazine *et al.*, 2001). Ants coordinate their behavior with one another not directly but indirectly through modifying their environment (for example through pheromones), in ways that feedback to their own behaviors. This leads not to static equilibrium behavior, but

[28]Appropriate 2-skill equations in Table 1 support this interpretation: differential equations of skill growth for the fixed-rich and the endogenous environments are identical, both for target and for source reproduction, except for the boundary condition of input balls being present in the endogenous urn.

[29]The detective strategy that finally worked was comparing the sequentially updated distribution of rules/skills to the sequentially updated distribution of output balls/products, for a large number of sample runs.

[30]Hence target reproduction remains superior to source reproduction plus endogenous.

rather to flexible physical structures (like 'roads') that have the capacity to adapt, both to exogenous shocks and to what the ant colony itself does.

Supporting this analogy further is our observation that product outputs in our model do not converge to a fixed composition under endogenous environments. Peaks of modal product production stochastically change through time, like waves, even in hypercycle skill equilibrium. Such moving peaks are observed to be very sharp under selective search ('specialized production'), whereas they are more gentle under random search ('diversified production').

This flexible-output capacity can be understood through equation (7) in Box 1. Output '1' increases in volume the more '2→1' skills there are, and it decreases in volume the more '1→2' rules there are. As explained above, this protects enough against cancerous growth to permit the skill equations (2)–(6) to kick in.[31] But these skill equations operate homeostatically to bring n_{12} and n_{21} back into line. Eventually $n_{21} > n_{12}$ tips over to $n_{21} < n_{12}$, producing a switch in modal output. And so the output cycling continues.

No particular production output is favored over any other in this model.[32] So there is no basis for defining 'optimality' in output. The flexibility of output produced by this model, however, demonstrates that endogenous hypercyclic production systems have considerable capacity for collective adaptation, were such adaptation to be rewarded. In the short-run, output can be shifted easily without any radical reorganization of hypercyclic production chains. In the long-term, output *has* to shift around, or else the natural repair mechanism of endogenous environment will be disabled.[33]

We close this section with the following evolutionary speculation: assuming that greater complexity is good for some reason not specified in this paper, then endogenous environments permit selfish learning to thrive and to compete, even though altruistic learning is 'naturally' better. Target reproduction and endogenous environments are redundant: if target reproduction is present, then endogenous environments are not needed. But if selfish learning is presumed for whatever reason, then endogenous environments are crucial for progress in technological complexity.

[31]We have observed in our simulations that the skill differential equations operate on a slower time scale than the product equations, especially for selective search, even though they were not explicitly designed to do so.

[32]In other words, there are no consumers with particular tastes.

[33]This sentence applies more to selective search than to random search. In random search the distribution of output is peaked, but somewhat gently so ('diversified output'). In specialized search, the distribution of output is extremely peaked ('specialized output'). Unconstrained this specialized output poses no problems to the reproducibility of hypercycles, since output rotates around, thereby renewing all skills along the chain. But if this rotating output stopped, then eventually some skills would be starved into death, especially those skills located at the end of long 'food chains' of transformed products. Such long chains, of course, are characteristic of complex technologies

Box 1 Growth-of-skill equations for the 2-element hypercycle in a trading dyad

Setup: a) Transformation Structure:

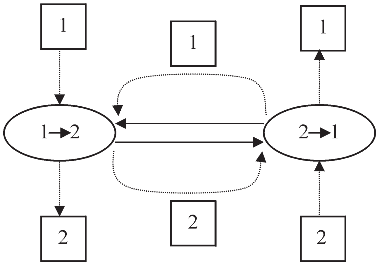

b) Spatial Layout:

	n_{12}	n_{21}	

Volume of rules/skills: $n_{12} + n_{21} = N$
Volume of balls/products: $b_1 + b_2 = B$
Growth-of-skill equations:
1. Non-spatial topology: unchained

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12}) = \left(\frac{n_{12}}{N}\right)\left(\frac{n_{21}}{N}\right)\left[\left(\frac{n_{21}}{N}\right) - \left(\frac{n_{12}}{N}\right)\right]$$

2. Spatial topology: fixed-rich environment, with selective search
(a) Source reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12}) = 0$$

(b) Target reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12}) = \left(\frac{9}{64}\right)\left[\left(\frac{n_{21}}{N}\right) - \left(\frac{n_{12}}{N}\right)\right]$$

3. Spatial topology: fixed-rich environment, with random search
(a) Source reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12}) = 0$$

(b) Target reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12}) = \left(\frac{17}{256}\right)\left[\left(\frac{n_{21}}{N}\right) - \left(\frac{n_{12}}{N}\right)\right]$$

4. Spatial topology: fixed poor environment, with either search
(a) Source reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12}) = \left(\frac{1}{64}\right)\left(\frac{n_{12}}{N}\right)\left(\frac{n_{21}}{N}\right)\left[8 + \left(\frac{n_{12}}{N}\right)\right]$$

Box 1 Continued

4(b) Target reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12})=\left(\frac{1}{64}\right)\left(\frac{n_{12}}{N}\right)\left[8+\left(\frac{n_{12}}{N}\right)\right]$$

5. Spatial topology: endogenous environment, with selective search
(a) Source reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12})=0$$

as long as $b_1 > 0$ and $b_2 > 0$.
(b) Target reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12})=\left(\frac{9}{64}\right)\left[\left(\frac{n_{21}}{N}\right)-\left(\frac{n_{12}}{N}\right)\right]$$

as long as $b_1 > 0$ and $b_2 > 0$.
6. Spatial topology: endogenous environment, with random search
(a) Source reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12})=\left(\frac{1}{64}\right)\left(\frac{n_{12}}{N}\right)\left(\frac{n_{21}}{N}\right)\left[\left(\frac{b_1}{B}\right)-\left(\frac{b_2}{N}\right)\right]\left[9-\left(\frac{b_1}{B}\right)\left(\frac{n_{21}}{N}\right)-\left(\frac{b_2}{N}\right)\left(\frac{n_{12}}{N}\right)\right]$$

(b) Target reproduction of rules

$$\frac{\mathrm{d}}{\mathrm{d}t}E(n_{12})=\left(\frac{1}{64}\right)\left[\left(\frac{b_2}{B}\right)\left(\frac{n_{21}}{N}\right)^2-\left(\frac{b_1}{N}\right)\left(\frac{n_{12}}{N}\right)^2\right]\left[8+\left(\frac{n_{12}}{N}\right)-\left(\frac{n_{12}}{N}\right)^2\right]$$

7. Spatial topology: endogenous environment, with either search
(a) Either method of reproduction of balls

$$\frac{\mathrm{d}}{\mathrm{d}t}E(b_1)=\left\{\left(\frac{7}{8}\right)\Big/\left[1-\left(\frac{1}{8}\right)^2\right]\right\}\left[\left(\frac{n_{21}}{N}\right)-\left(\frac{n_{12}}{N}\right)\right]$$

3.4 The effect of search intelligence

The main finding about degree of search intelligence is a negative one. Contrary to our expectations, selective search did not improve the chances of hypercycle emergence over random search. There is no significant difference[34] between the data reported in Figure 1 and the data reported in Figure 2.

[34]Perhaps one might make a case for a slight difference between the two search settings for source reproduction under the fixed-rich environment. But that difference goes in the opposite direction from

Selective search does increase the speed of convergence to equilibrium, compared to random search. It just does not change the equilibrium outcome. The decision tree representations of the 2-skill hypercycle, in the appendix, reveal why this is true. Search procedure affects the second branch of those trees. Selective search triggers the sequences specified in the appendix. Random search only adds the probability of one half of branching into 'no change' to these selective-search sequences. Other than this, nothing is altered.

In this context, therefore, search efficiency is not all it is sometimes cracked up to be. Search efficiency may be beneficial for a particular agent. But search efficiency through a given structure does not itself alter the evolution of that structure.

More generally, intelligence is not necessary for complexity to emerge—a point we knew already from observing evolution. Rather the evolutionary sequence is the opposite: complexity is necessary for intelligence to emerge.

4. Results: hypercycle structure

Finally, in addition to reporting on emergence, we report on structure: what sorts of 'industrial districts'—that is, what sorts of firms and firm clusters—are generated by our hypercycle model of economic production? Table 1 presents data, for successful non-crash hypercycles only, on three features of equilibrium spatial clusters: average number of firms supported by the hypercycles, average number of discrete skills[35] per firm, and average percentage of discrete skills that are parasites.

There are a few subtleties in these data, but the primary finding is quite simple: reproduction/learning mode has a more substantial effect on the shape and nature of firm clusters than does either resource environment or search procedure. Target reproduction ('altruistic learning') generates more surviving firms, who are more multi-skilled in character, than does source reproduction ('selfish learning'). And target reproduction tolerates more parasitic free-riding than does source reproduction, with no ill effect on hypercycle survival. The quantitative impact of target reproduction on the number-of-skills profile of firms is modest, but it is noticeable and consistent. The quantitative impact of target reproduction on the number of surviving firms in trading clusters is substantial. And altruistic learning's high tolerance of free-riding may be less surprising than the fact that this tolerance is not weeded out, even in a sharply competitive, zero-sum context.

The mechanism that produced these morphological features is the same as that

our expectations: random search there is superior to selective search. Rather than try to make something out of this isolated instance, we prefer to stick to our basic finding of 'no significant difference' between the search procedures.

[35]We use the term 'discrete skills' to refer to non-redundant skills. The total volume of skills in this conservation-of-mass simulation was fixed at 200, as we explained above. But most of these are replications. 'Discrete skills' are the number of unique types of skills located in each firm, summed across all surviving firms.

which produced the corresponding metabolic effect: direct repair. In target reproduction, high-volume skills reach out to low-volume (hence threatened) compatible neighbors, whether those neighbors are contributing to the hypercycle that sustains themselves or not. In contrast, consistency of effect across morphological and metabolic levels does not apply to our second repair mechanism: endogenous environment. The stigmergy mechanism of endogenous environment controls cancerous growth, but it does not induce selfish skills pro-actively to reach out to restructure and thereby to save their neighbors. Because neighbors are not changed, firm and cluster morphologies are not changed.

An intriguing second-order effect contained in Table 1 is that increasing techno-

Table 1 Structure of non-crash hypercyclic spatial clusters, at equilibrium

	Selective				Random			
	Source		Target		Source		Target	
	Rich	Endog.	Rich	Endog.	Rich	Endog.	Rich	Endog.
(A) Average number of surviving firms								
2-skill	7.2	7.6	12.1	12.7	8.1	7.8	13.3	12.2
3-skill	5.5	6.0	10.9	12.0	6.1	6.7	12.4	10.6
4-skill	5.0	5.0	10.1	9.7	6.2	6.0	10.5	11.2
5-skill		5.1	7.6	7.6	6.0	6.8	8.9	8.2
6-skill		5.4	7.0	6.2		6.0	7.1	7.3
7-skill		6.7	7.3	7.1		6.3	7.5	7.2
8-skill		5.7	7.4	8.5		6.0	7.0	6.0
9-skill			6.7	9.0			11.0	7.5
(B) Average number of discrete skills per firm								
2-skill	1.15	1.12	1.24	1.20	1.17	1.12	1.21	1.31
3-skill	1.22	1.03	1.35	1.42	1.30	1.27	1.38	1.49
4-skill	1.25	1.21	1.40	1.52	1.24	1.29	1.50	1.48
5-skill		1.29	1.57	1.57	1.17	1.35	1.36	1.50
6-skill		1.37	1.56	1.58		1.42	1.46	1.41
7-skill		1.15	1.44	1.51		1.42	1.50	1.53
8-skill		1.53	1.45	1.59		1.33	1.46	1.67
9-skill			1.50	1.67			1.55	2.0
(C) Percentage of discrete skills that are parasites								
2-skill	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3-skill	0.00	0.05	0.09	0.06	0.05	0.04	0.07	0.10
4-skill	0.00	0.04	0.07	0.06	0.03	0.05	0.12	0.05
5-skill		0.06	0.16	0.16	0.07	0.08	0.29	0.17
6-skill		0.00	0.22	0.24		0.07	0.33	0.35
7-skill		0.00	0.14	0.18		0.04	0.26	0.29
8-skill		0.00	0.20	0.17		0.00	0.27	0.17
9-skill			0.27	0.06			0.26	0.06

logical complexity is packed into increasing spatial concentration, both of firms in a cluster and of skills in a firm. In our very simple setup, there are limits to how far this compression can proceed:[36] for target reproduction, up to 9-skill hypercycles can compressed into 7-firm clusters with 1½ skills per firm, on average. Yet the qualitative trend produced by our model—the more complex the technology, the more dense the spatial container—has biological verisimilitude.[37]

5. Summary and discussion

Even within our chemical perspective, this article has not fully addressed the issue of the co-evolution of technology and industry, because the evolution of products has not been modeled explicitly. That is the next step.[38] What this article has done, however, is to establish three principles of social organization that provide sufficient foundations for the unconscious evolution of technological complexity: structured topology, altruistic learning, and stigmergy.

1. Unstructured interaction topologies are not conducive to the emergence of complex technologies. Without help through embodiment, long sequences of skills cannot dynamically regulate their own stable reproduction. 'Structured topology' does not have to mean spatial, as it does here (cf. Cohen *et al.*, 2001). But constraints on interaction are necessary, firstly, in order to break the symmetry of full mixing and induce localized heterogeneity, and secondly, in order to allow positive reproductive feedback to turn that raw heterogeneity into path-dependent memory of past successes. This is the chemistry answer to why firms exist:[39] dynamic barriers of

[36]With higher initial total-volume density that the parameter setting (200 rules spread over 100 firms) explored here, these exact numbers probably could be pushed up some.

[37]It is hard to imagine, for example, low-density creatures such as algae or sponges ever becoming very complex. The surprise in this model is that spatial compression happens 'automatically' (at least up to a point), without any special mechanism.

[38]In our HYPERCYCLE Repast code, we have already prepared for this next step by enabling the specification of an arbitrary set of initial ($i \rightarrow j$) rules, not just a cyclic set. This extension, which moves beyond the hypercycle framework of Eigen and Schuster, will permit exploration of the emergence of arbitrary networks of skills. More generalized network structures will have to contain cycles within them in order to reproduce. But apart from this, the broader class of network structures that is sustainable by our learning dynamics is currently unknown. Our proposed extension bears a family resemblance to the work of Jain and Krishna (1998).

[39]Padgett (1997) discusses why the traditional explanations for the firm given in neo-classical economics—namely, transaction-cost economics and principal-agent theory—are inadequate from a biological perspective. 'Such a transposition of "the firm" down into a series of dyadic negotiations overlooks the institutionalized autonomy of all stable organizations. In organisms, social or biological, rules of action and patterns of interaction persist and reproduce even in the face of constant turnover in component parts, be these cells, molecules, principals, or agents. In the constant flow of people through

technological complexity can be transcended once global is transformed into the concatenation of locals.

Classic Marshallian industrial districts receive the benefits of physical space naturally. In an era of globalization, densely interconnected firms may or may not be so fortunate. What our model implies is that trading within these new 'virtual industrial districts' will have to become interactionally constrained for technological progress, not instability, to be the consequence of increased connectivity (cf. May, 1974; Davidow, 2000).

2. The potential benefits of localized embodiment are more easily reaped through altruistic learning than through selfish learning. When recipients, not initiators, of transactions receive the reproductive rewards, complex technologies are more readily nurtured and repaired.[40] Free-riding occurs, but that does not threaten system stability.

This conclusion is consistent with anthropological emphases on gift-giving in primitive economies (Mauss, 1967; Sahlins, 1972). It is also consistent with sociological observations about the 'strange' persistence of generous behavior in modern economies (Macauley, 1963; Granovetter, 1985; Uzzi, 1997; Padgett and McLean, 2002). Ours may not be the only explanation of generosity. But repair is one evolutionary reason for the natural selection of this behavior in competitive economies of all sorts. Altruistic learning stabilizes the reproduction of distributed technological skills, on which all depend.[41]

3. When altruistic learning is not present for whatever reason, then stigmergy—the endogenous construction of resource environments—is second best. Entomologists (e.g. Bonabeau *et al.*, 1999; Camazine *et al.*, 2001) have shown that stigmergy flexibly can coordinate sophisticated collective behavior among myopic social insects. We have shown that stigmergy also can regulate the cancerous growth of selfish learners, keeping even long chains of distributed skills alive.

Adams (1966, 1996) has long argued that cities are crucial to the history of technology. His exemplar case is Mesopotamia, where spatial feedbacks between settlements and rivers guided the joint emergence of urban concentration, irrigation technology,

organizations, the collectivity typically is not renegotiated anew. Rather, within constraints, component parts are transformed and molded into the ongoing flow of action' (Padgett, 1997: 199–200).

[40]Sabel (1994) recommends squeezing the temporal distance between the two sides of an iterated transaction until this distinction is effaced. Such relational constraints are consistent with our first conclusion. Regarding our second conclusion, however, Padgett (1997) demonstrated that joint reproduction, the closest analogue in chemistry to this recommendation, does not succeed in breaking complexity barriers.

[41]This may be news to some rational choice theorists, but it will not come as a surprise to parents.

and the shapes of the rivers themselves.[42] Of course our model is far too minimalist for real history, but it may illustrate one reason why the spatial reorganization of land into cities and the development of complex technologies proceeded hand in hand. Technology causes cities, as we all know; less obviously, the spatial products of technology channel and regulate the social forces that produce it. To put it simply: cities stabilize selfishness.[43]

In this article, we have developed a few simple tools, imported from chemistry, that enable us to investigate systematically the co-evolution of distributed technology and social organization. Extreme assumptions about the absence of consciousness are implied by our specification. The payoff of such extreme simplification is the discovery of three social-organizational principles enabling technological evolution. How robust such principles are to alternate specifications is an important issue to explore in the future. Regardless of the answer to that question, however, we hope to have demonstrated at minimum that complex cognition is not necessary for the emergence and functioning of complex economies. Just as March and Simon (1958) argued long ago.[44]

Acknowledgements

We have the honor of writing this paper for James G. March, on the occasion of his 75th birthday. In our opinion, he is the greatest organization theorist ever (for justification of this claim, see Padgett, 1992). The first author would like to acknowledge the financial support of the Santa Fe Institute and the Hewlett Foundation. We would also like to express our gratitude to David Sallach, director of Social Science Research Computing at the University of Chicago. David has spearheaded the creation and development of Repast, our agent-based simulation platform, and he has been very supportive of our particular project within Repast. Nick Collier is the chief designer and programmer of Repast.

[42]It would be stretching it, to say the least, to expect the mechanisms of our model to explain the invention of writing. But it is worth noting that this too was implicated in these co-evolutionary developments.

[43]Given the absence of cities in the prehistory of mankind, perhaps one is as justified in speculating about 'the evolution of selfishness' as one is justified in speculating about 'the evolution of altruism'. More reasonably, we should all support the efforts of the evolutionary economists cited at the beginning of this article, and more historically oriented others (e.g. Herrigel, 1996; Sabel and Zeitlin, 1997), to try to specify alternate social-organizational paths to technological complexity.

[44]March and Simon went beyond our point to also argue that social structures enable human cognition, by mapping the world down to levels we can comprehend.

Address for correspondence

J. F. Padgett, University of Chicago & Santa Fe Institute, Chicago, IL, USA. Email: jpadgett@midway.uchicago.edu.

References

Adams, R. M. (1966), *The Evolution of Urban Society: Early Mesopotamia and Prehispanic Mexico*. Aldine: Chicago, IL.

Adams, R. M. (1996), *Paths of Fire: An Anthropologist's Inquiry into Western Technology*. Princeton University Press: Princeton, NJ.

Antonelli, C. and R. Marchionatti (1998), 'Technological and organizational change in the process of industrial rejuvenation: the case of the Italian cotton textile industry,' *Cambridge Journal of Economics*, **22**, 1–18.

Boerlijst, M. C. and P. Hogeweg (1991), 'Spiral wave structures in pre-biotic evolution: hypercycles stable against parasites,' *Physica B*, **48**, 17–28.

Bonabeau, E. , M. Dorigo and G. Theraulaz (1999), *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press: Oxford.

Camazine, S., J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz and E. Bonabeau (2001), *Self-organization in Biological Systems*. Princeton University Press: Princeton, NJ.

Cohen, M. D., R. L. Riolo and R. Axelrod (2001), 'The role of social structure in the maintenance of cooperative regimes,' *Rationality and Society*, **13**, 5–32.

Coriat, B. and G. Dosi (1998), 'Learning how to govern and learning how to solve problems: on the co-evolution of competencies, conflicts and organizational routines,' in A. Chandler, P. Hagstrom and O. Solvell (eds), *The Dynamic Firm: The Role of Technology, Strategy, Organization and Regime*. Oxford University Press: Oxford.

Davidow, W. (2000), 'Social tornadoes: interconnection, instability, and social policy,' unpublished manuscript.

Dosi, G., R. Giannetti and P. A. Toninelli (1992), *Technology and Enterprise in a Historical Perspective*. Clarendon Press: Oxford.

Dosi, G., R. Nelson and S. Winter (2000), *The Nature and Dynamics of Organizational Capabilities*. Oxford University Press: Oxford.

Eigen, M. (1971), 'Self-organization of matter and the evolution of biological macromolecules,' *Naturwissenschaften*, **58**, 465–523.

Eigen, M. and P. Schuster (1979), *The Hypercycle: A Principle of Natural Self-organization*. Springer-Verlag: New York.

Fontana, W. F. and L. W. Buss (1994), 'The arrival of the fittest: toward a theory of biological organization,' *Bulletin of Mathematical Biology*, **56**, 1–64.

Granovetter, M. (1985), 'Economic action and social structure: the problem of embeddedness,' *American Journal of Sociology*, **91**, 481–510.

Hayek, F. (1948), *Individualism and Economic Order*. University of Chicago Press: Chicago, IL.

Herrigel, G. (1996), *Industrial Constructions: The Sources of German Industrial Power*. Cambridge University Press: Cambridge.

Hofbauer, J. and K. Sigmund (1988), *Dynamical Systems and the Theory of Evolution*. Cambridge University Press: Cambridge.

Hughes, T. P. (1983), *Networks of Power: Electrification in Western Society, 1880–1930*. The Johns Hopkins University Press: Baltimore, MD.

Hughes, T. P. (1987), 'The evolution of large technological systems,' in W. Bijker, T. Hughes and T. Pinch (eds), *The Social Construction of Technological Systems*. MIT Press: Cambridge, MA.

Hutchins, E. (1995), *Cognition in the Wild*. MIT Press: Cambridge, MA.

Jain, S. and S. Krishna (1998), 'Autocatalytic sets and the growth of complexity in an evolutionary model,' *Physical Review Letters*, **81**, 5684–5687.

Kauffman, S. A. (1986), 'Autocatalytic sets of proteins,' *Journal of Theoretical Biology*, **119**, 1–24.

Kauffman, S. A. (1993), *The Origins of Order: Self-organization and Selection in Evolution*. Oxford University Press: New York.

Landesmann, M. and R. Scazzieri (1996), *Production and Economic Dynamics*. Cambridge University Press: Cambridge.

Macauley, S. (1963), 'Non-contractual relations in business: a preliminary study,' *American Sociological Review*, **28**, 55–67.

Malerba, F. and L. Orsinigo (1993), 'Technological regimes and firm behavior,' *Industrial and Corporate Change*, **2**, 45–71.

March, J. G. and H. A. Simon (1958), *Organizations*. Wiley: New York.

Maturana, H. and F. Varela (1980), 'Autopoesis: the organization of the living,' in *Autopoesis and Cognition*. D. Reidel: Dordrecht.

Mauss, M. (1967), *The Gift: Forms and Functions of Exchange in Archaic Societies* [1899]. Norton: New York.

May, R. M. (1974), *Stability and Complexity in Model Ecosystems*. Princeton University Press: Princeton, NJ.

McKelvey, M. (1997), 'Co-evolution in commercial genetic engineering,' *Industrial and Corporate Change*, **6**, 503–532.

Nelson, R. R. (1994), 'The co-evolution of technology, industrial structure, and supporting institutions,' *Industrial and Corporate Change*, **3**, 47–64.

Nelson, R. R. (1995), 'Recent evolutionary theorizing about economic change,' *Journal of Economic Literature*, **33**, 48–90.

Nelson, R. R. and S. G. Winter (1982), *An Evolutionary Theory of Economic Change*. Belnap Press of Harvard University Press: Cambridge, MA.

Padgett, J. F. (1992), 'Learning from (and about) March,' *Contemporary Sociology*, **21**, 744–749.

Padgett, J. F. (1997), 'The emergence of simple ecologies of skill: a hypercycle approach to economic organization,' in W. B. Arthur, S. N. Durlauf and D. A. Lane (eds), *The Economy as an Evolving Complex System II*. Addison-Wesley: Reading, MA.

Padgett, J. F. and P. McLean (2002), 'Economic and social exchange in Renaissance Florence,' unpublished manuscript.

Powell, W. W. (1996), 'Inter-organizational collaboration in the biotechnology industry,' *Journal of Institutional and Theoretical Economics*, **152**, 197–215.

Prigogine, I. and P. Glansdorff (1971), *Thermodynamic Theory of Structure: Stability and Fluctuations*. Wiley: New York.

Rosenkopf, L. and M. L. Tushman (1998), 'The co-evolution of community networks and technology: lessons from the flight simulation industry,' *Industrial and Corporate Change*, **7**, 311–346.

Sabel, C. F. (1994), 'Learning by monitoring: the institutions of economic development,' in N. J. Smelser and R. Swedberg (eds), *The Handbook of Economic Sociology*. Princeton University Press: Princeton, NJ.

Sabel, C. F. and J. Zeitlin (eds) (1997), *World of Possibilities: Flexibility and Mass Production in Western Industrialization*. Cambridge University Press: Cambridge.

Sahlins, M. (1972), *Stone Age Economics*. Aldine: New York.

Stadler, B. M. R. and P. F. Stadler (2002), 'Molecular replicator dynamics,' working paper 02-09-049, Santa Fe Institute.

Uzzi, B. (1997), 'The sources and consequences of embeddedness for the economic performance of organizations,' *American Sociological Review*, **42**, 35–67.

Warlien, M. (1995), 'Hierarchical selection and organizational adaptation,' *Industrial and Corporate Change*, **4**, 161–186.

Appendices: derivation of growth-of-skill equations for 2-element hypercycles within one trading dyad

1. *Non-spatial setup: unchained (= one-step) (same as in Hofbauer and Sigmund, 1988)*

1a. SOURCE Reproduction of Rules:

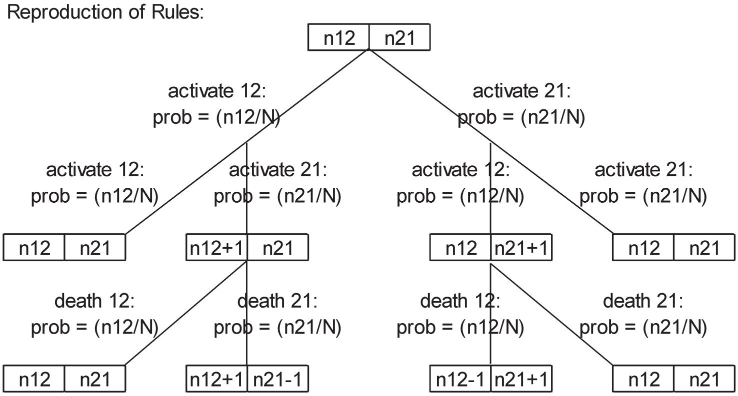

Therefore, transition probabilities are:

prob {[n12 | n21] → [n12+1 | n21-1]} = (n12/N) (n21/N) (n21/N)

prob {[n12 | n21] → [n12-1 | n21+1]} = (n21/N) (n12/N) (n12/N)

Thus, d/dt E(n12) = (n12/N) (n21/N) [(n21/N) - (n12/N)]

1b. TARGET Reproduction of Rules:

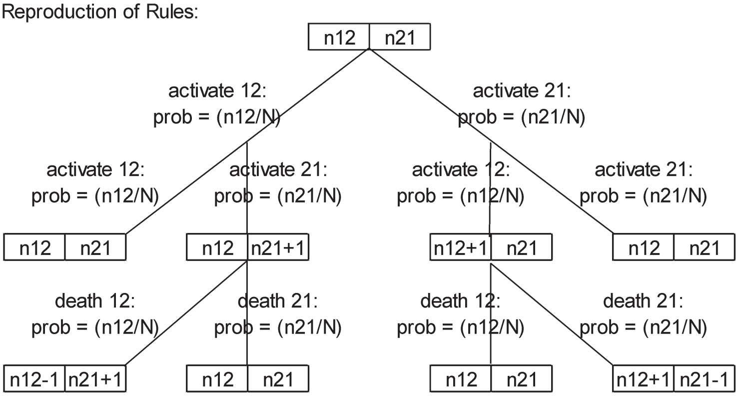

Therefore, transition probabilities are:

prob {[n12 | n21] → [n12+1 | n21-1]} = (n21/N) (n12/N) (n21/N)

prob {[n12 | n21] → [n12-1 | n21+1]} = (n12/N) (n21/N) (n12/N)

Thus, d/dt E(n12) = (n12/N) (n21/N) [(n21/N) - (n12/N)] (i.e., identical to SOURCE)

2. *Spatial topology: fixed-rich environment, selective search, product chained*

2a. SOURCE Reproduction of Rules:

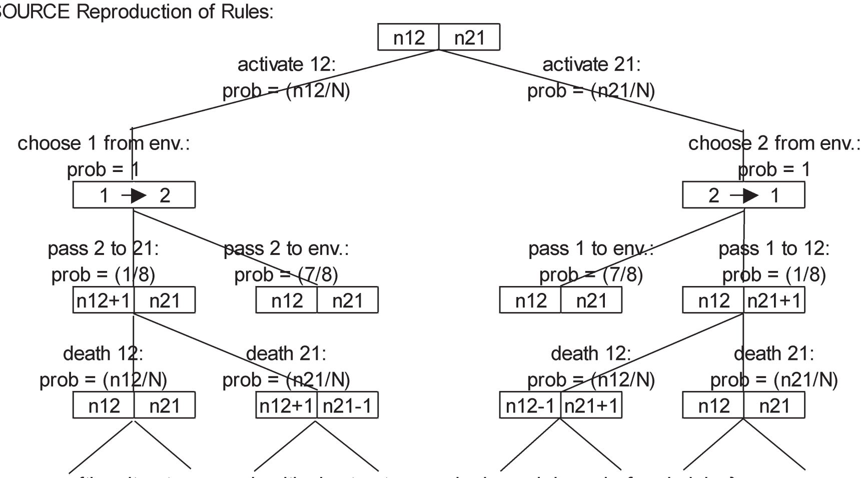

{then iterate same algorithmic structure again, in each branch, for chaining}

After two cycles, this yields:

prob { n12 | n21 → n12+1 | n21-1 } = (1/8)**2 (n12/N)(n21/N) [2(n12/N)**2 + 2(n21/N)**2 + 7]

prob { n12 | n21 → n12-1 | n21+1 } = (1/8)**2 (n12/N)(n21/N) [2(n12/N)**2 + 2(n21/N)**2 + 7]

prob { n12 | n21 → n12+2 | n21-2 } = (1/8)**2 (n12/N)**2 (n21/N)**2

prob { n12 | n21 → n12-2 | n21+2 } = (1/8)**2 (n12/N)**2 (n21/N)**2

Thus, d/dt E(n12) = 0

2b. TARGET Reproduction of Rules:

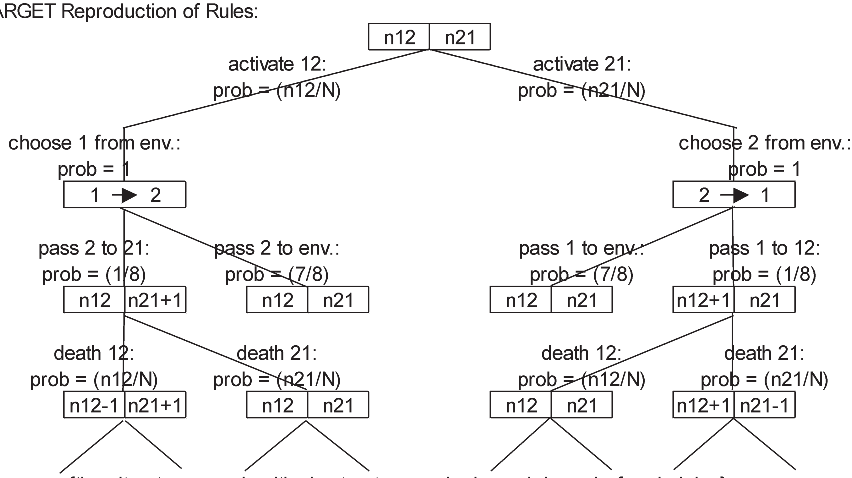

{then iterate same algorithmic structure again, in each branch, for chaining}

prob { n12 | n21 → n12+1 | n21-1 } = (1/8)**2 (n21/N)**2 [4(n12/N)(n21/N) + 7]

prob { n12 | n21 → n12-1 | n21+1 } = (1/8)**2 (n12/N)**2 [4(n12/N)(n21/N) + 7]

prob { n12 | n21 → n12+2 | n21-2 } = (1/8)**2 (n21/N)**4

prob { n12 | n21 → n12-2 | n21+2 } = (1/8)**2 (n12/N)**4

Thus, d/dt E(n12) = (9/64) [(n21/N) - (n12/N)]

3. *Spatial topology: fixed-rich environment, random search, product chained*

3a. SOURCE Reproduction of Rules:

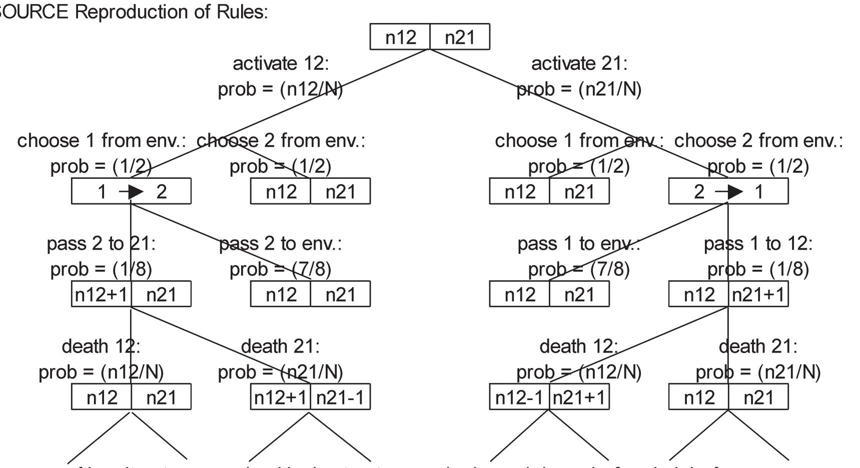

{then iterate same algorithmic structure again, in each branch, for chaining}

After two cycles, this yields:

prob { n21 → n12+1 n21-1 } = (1/2)**2 (1/8)**2 (n12/N)(n21/N) [2(n12/N)**2 + 2(n21/N)**2 + 15]
prob { n21 → n12-1 n21+1 } = (1/2)**2 (1/8)**2 (n12/N)(n21/N) [2(n12/N)**2 + 2(n21/N)**2 + 15]
prob { n21 → n12+2 n21-2 } = (1/2)**2 (1/8)**2 (n12/N)**2 (n21/N)**2
prob { n21 → n12-2 n21+2 } = (1/2)**2 (1/8)**2 (n12/N)**2 (n21/N)**2

Thus, d/dt E(n12) = 0

3b. TARGET Reproduction of Rules:

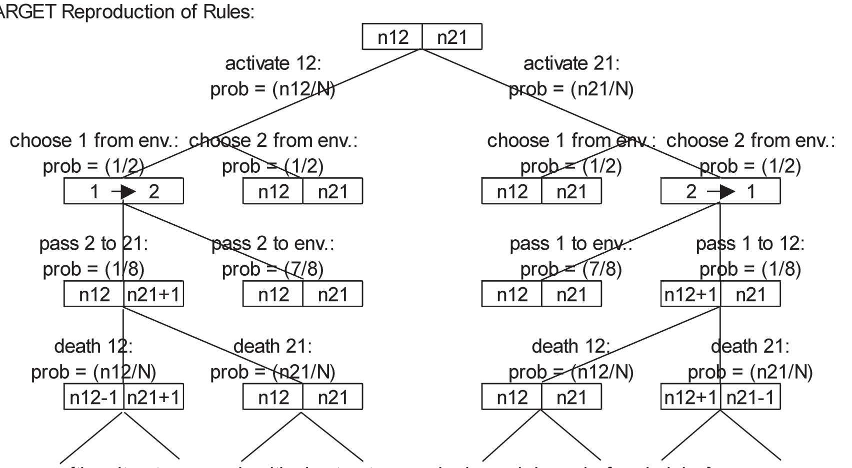

{then iterate same algorithmic structure again, in each branch, for chaining}

prob { n12 n21 → n12+1 n21-1 } = (1/2)**2 (1/8)**2 (n21/N)**2 [4(n12/N)(n21/N) + 15]
prob { n12 n21 → n12-1 n21+1 } = (1/2)**2 (1/8)**2 (n12/N)**2 [4(n12/N)(n21/N) + 15]
prob { n12 n21 → n12+2 n21-2 } = (1/2)**2 (1/8)**2 (n21/N)**4
prob { n12 n21 → n12-2 n21+2 } = (1/2)**2 (1/8)**2 (n12/N)**4

Thus, d/dt E(n12) = (17/256) [(n21/N) - (n12/N)]

4. *Spatial topology: fixed-poor environment, either search, product chained*

4a. SOURCE Reproduction of Rules:

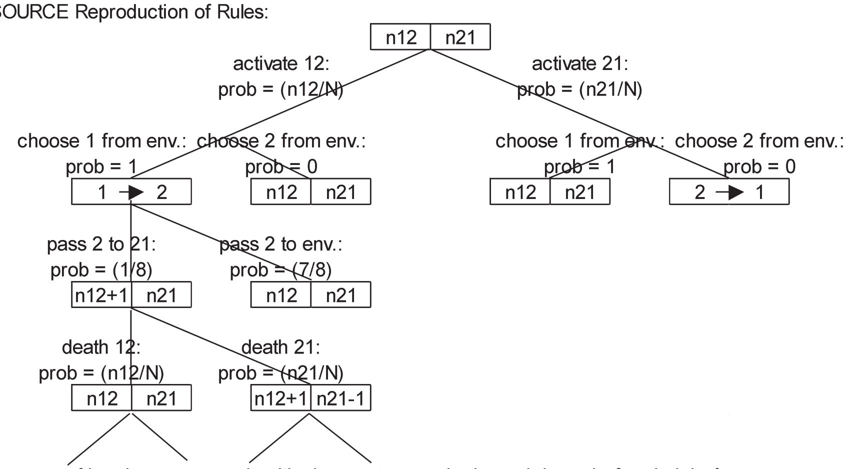

{then iterate same algorithmic structure again, in each branch, for chaining}

After two cycles, this yields:

prob { n12 | n21 → n12+1 | n21-1 } = (1/8)**2 (n12/N)(n21/N) [2(n12/N)**2 + (n21/N) + 7]
prob { n12 | n21 → n12-1 | n21+1 } = 0
prob { n12 | n21 → n12+2 | n21-2 } = (1/8)**2 (n12/N)**2 (n21/N)**2
prob { n12 | n21 → n12-2 | n21+2 } = 0
Thus, d/dt E(n12) = (1/64) (n12/N)(n21/N) [8 + (n12/N)]

4b. TARGET Reproduction of Rules:

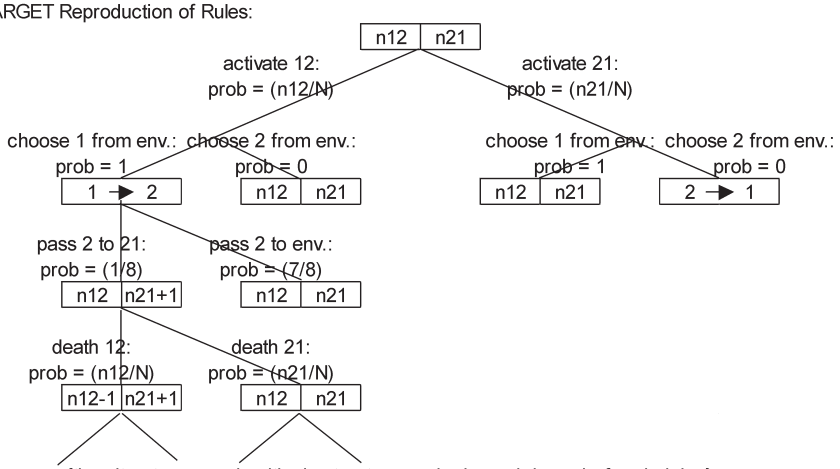

{then iterate same algorithmic structure again, in each branch, for chaining}

prob { n12 | n21 → n12+1 | n21-1 } = 0
prob { n12 | n21 → n12-1 | n21+1 } = (1/8)**2 (n12/N)**2 [2(n12/N)(n21/N) + (n21/N) + 7]
prob { n12 | n21 → n12+2 | n21-2 } = 0
prob { n12 | n21 → n12-2 | n21+2 } = (1/8)**2 (n12/N)**4
Thus, d/dt E(n12) = - (1/64) (n12/N)**2 [8 + (n12/N)]

5. *Spatial topology: endogenous environment, selective search, product chained*

5a. SOURCE Reproduction of Rules:

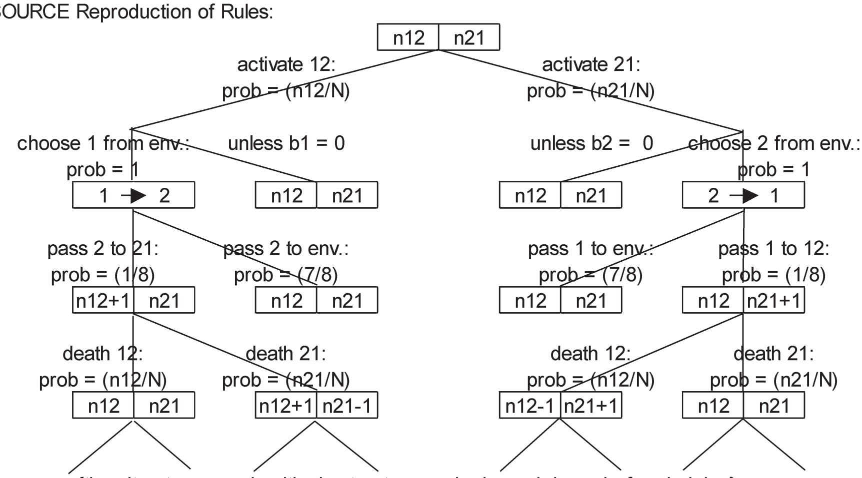

{then iterate same algorithmic structure again, in each branch, for chaining}

After two cycles, this yields:

prob { n12 | n21 → n12+1 | n21-1 } = (1/8)**2 (n12/N)(n21/N) [2(n12/N)**2 + 2(n21/N)**2 + 7]

prob { n12 | n21 → n12-1 | n21+1 } = (1/8)**2 (n12/N)(n21/N) [2(n12/N)**2 + 2(n21/N)**2 + 7]

prob { n12 | n21 → n12+2 | n21-2 } = (1/8)**2 (n12/N)**2 (n21/N)**2

prob { n12 | n21 → n12-2 | n21+2 } = (1/8)**2 (n12/N)**2 (n21/N)**2

Thus, d/dt E(n12) = 0, as long as b1>0 and b2>0

5b. TARGET Reproduction of Rules:

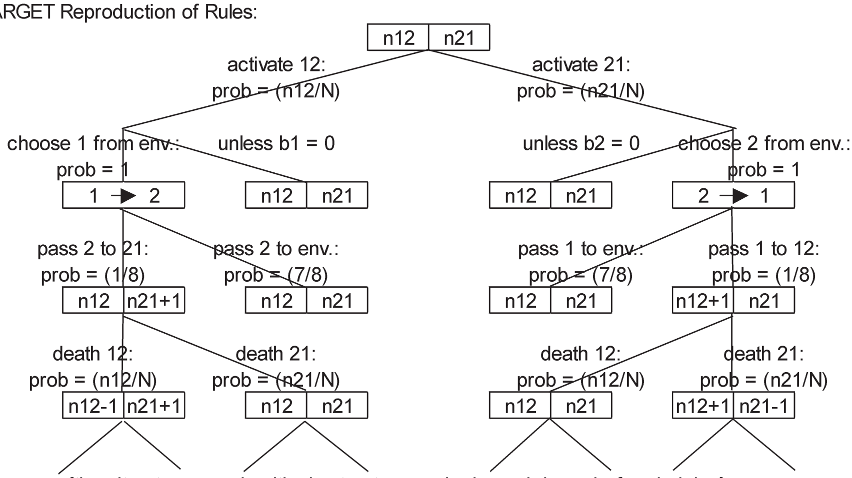

{then iterate same algorithmic structure again, in each branch, for chaining}

prob { n12 | n21 → n12+1 | n21-1 } = (1/8)**2 (n21/N)**2 [4(n12/N)(n21/N) + 7]

prob { n12 | n21 → n12-1 | n21+1 } = (1/8)**2 (n12/N)**2 [4(n12/N)(n21/N) + 7]

prob { n12 | n21 → n12+2 | n21-2 } = (1/8)**2 (n21/N)**4

prob { n12 | n21 → n12-2 | n21+2 } = (1/8)**2 (n12/N)**4

Thus, d/dt E(n12) = (9/64) [(n21/N) - (n12/N)], as long as b1>0 and b2>0

6. *Spatial topology: endogenous environment, random search, product chained*

6a. SOURCE Reproduction of Rules:

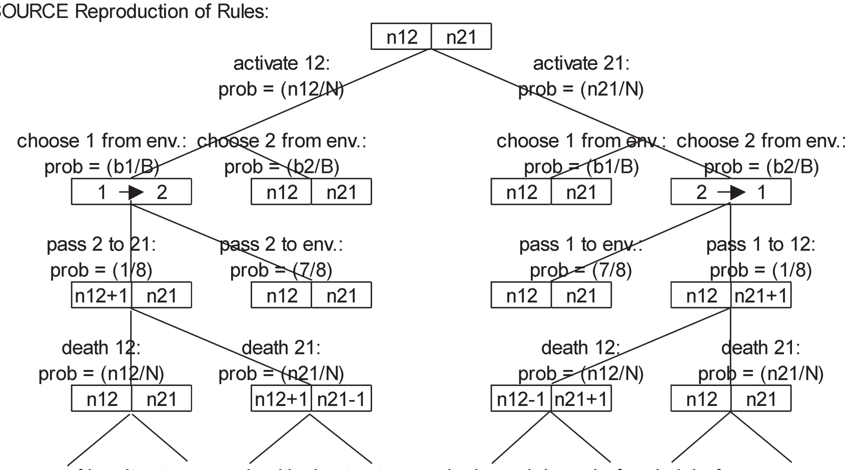

{then iterate same algorithmic structure again, in each branch, for chaining}

After two cycles, this yields:

prob { n12 | n21 → n12+1 | n21-1 } = [answer too long to write here]

prob { n12 | n21 → n12-1 | n21+1 } = [answer too long to write here]

prob { n12 | n21 → n12+2 | n21-2 } = (1/8)**2 (b1/B)**2 (n12/N)**2 (n21/N)**2

prob { n12 | n21 → n12-2 | n21+2 } = (1/8)**2 (b2/B)**2 (n12/N)**2 (n21/N)**2

Thus, d/dt E(n12) = (1/64) (n12/N)(n21/N) [(b1/B) - (b2/B)] [9 - (b1/B)(n21/N) - (b2/B)(n12/N)]

6b. TARGET Reproduction of Rules:

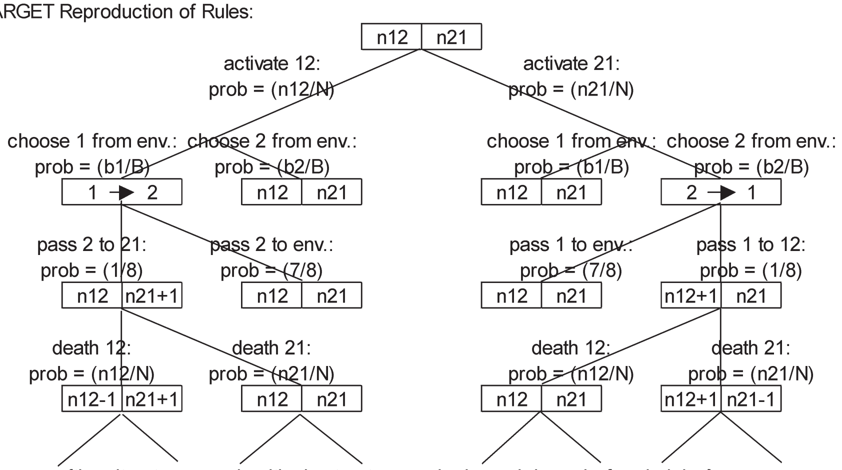

{then iterate same algorithmic structure again, in each branch, for chaining}

prob { n12 | n21 → n12+1 | n21-1 } = [answer too long to write here]

prob { n12 | n21 → n12-1 | n21+1 } = [answer too long to write here]

prob { n12 | n21 → n12+2 | n21-2 } = (1/8)**2 (b2/B)**2 (n21/N)**4

prob { n12 | n21 → n12-2 | n21+2 } = (1/8)**2 (b1/B)**2 (n12/N)**4

Thus, d/dt E(n12) = (1/64) [(b2/B)(n21/N)**2 - (b1/B)(n12/N)**2] [8 + (n12/N) - (n12/N)**2]

7. *Spatial topology: endogenous environment, either search, product chained*

7a. EITHER Reproduction of BALLS:

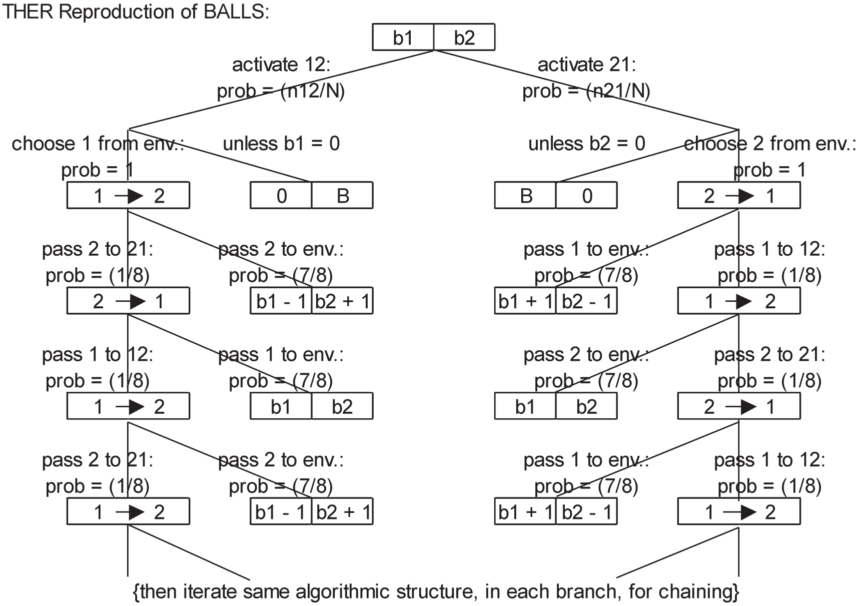

After infinite cycles, this yields:

prob { | b1 | b2 | ⟶ | b1+1 | b2-1 | } = {(7/8)/[1-(1/8)**2]} (n21/N)

prob { | b1 | b2 | ⟶ | b1-1 | b2+1 | } = {(7/8)/[1-(1/8)**2]} (n12/N)

Thus, d/dt E(b1) = {(7/8)/[1-(1/8)**2]} [(n21/N) - (n12/N)]