

ECONOMIC PRODUCTION AS CHEMISTRY II*

John F. Padgett, Peter McMahan and Xing Zhong

University of Chicago & Santa Fe Institute

April 2008

Forthcoming in John Padgett and Woody Powell (eds.),
Organizational and Market Emergence

*The 'II' in the title refers to fact that this chapter is a major extension, to other chemistries, of 'Economic Production as Chemistry' by Padgett, Lee and Collier, published in Industrial and Corporate Change (2003) 12: 843-877. We gratefully acknowledge the financial support of the National Science Foundation (program on Human and Social Dynamics) and the Hewlett Foundation, which funded the development of both the first and this second round of these agent-based models.

ECONOMIC PRODUCTION AS CHEMISTRY II

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 and social scientists are aware of the issue of co-evolution (for example, 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, Powell 1996, McKelvey 1997, Coriat and Dosi 1998, Rosenkopf and Tushman 1998, Murman 2003, Padgett and McLean 2006). 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, in this view, is how such an ensemble can self-organize and sustain itself, 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. Economic ‘life’ exists if an autocatalytic network of interlinked skills and products can emerge and renew itself, in the face of continual turnover and ‘death’ in its component skills and products.

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

Firms in this view are sites through which distributed and living ‘chemical reaction’ production processes flow. 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 and learning, into a coherent product-transformation network,² which then reproduces itself through time and ‘grows’ a set of firms to sustain itself.

Inspired by a specific literature in chemistry, that on hypercycles, Padgett, Lee and Collier (2003) developed one family of economic production models that operationalized this co-evolutionary perspective on markets. This chapter extends that modeling platform to cover a wider range of technological chemistries.

The ‘hypercycle’ is a specific model of the chemical origin of life pioneered by Eigen and Schuster (1971, 1979) and extended by others (for example, Hofbauer and Sigmund 1988, Kauffman 1986, 1993, 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: $\{1 \rightarrow 2, 2 \rightarrow 3, 3 \rightarrow 4, \dots, n \rightarrow 1\}$. Chemical cycles are crucial to the emergence and maintenance of life because these are the motors that sustain 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 itself will stop or ‘die’. Eigen and Schuster, Hofbauer and Sigmund, and others have explored how variation in reaction rates, in chemical density, and in number of components affect the dynamic stability or ‘survivability’ of various classes of hypercyclic chemical reactions, within a well-stirred liquid reaction tank. Boerlijst and Hogeweg (1991), Padgett (1997), and Padgett, Lee and Collier (2003) extended the investigation beyond the original liquid context to a spatial topology of interaction.

The hypercycle-modeling literature started with random distributions of interacting rules within linear or sequential rule sets, such as the one described above, and then searched for experimental conditions that would dynamically reproduce that rule set in the face of death. Here we re-label such linear-rule-set chemistries as SOLO H, which stands for ‘single hypercycle.’ In this chapter, we reproduce previous hypercycle (a.k.a. SOLO H) findings as benchmarks, but we also investigate a more general class of chemistries – namely, complete sets of interacting transformation rules $\{(i \rightarrow j)\}$, within experimentally varied maxima or n . Such complete families of chemistries we label ALL, standing for ‘all permutations.’³ For example, the ALL chemistry within a maximum of two products is $\{(1 \rightarrow 2), (2 \rightarrow 1)\}$. The ALL chemistry within a maximum of three products is $\{(1 \rightarrow 2), (1 \rightarrow 3), (2 \rightarrow 1), (2 \rightarrow 3), (3 \rightarrow 1), (3 \rightarrow 2)\}$. The ALL chemistry within

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

³ Only redundant and dynamically degenerate rules of the form $\{(i \rightarrow i)\}$ are excluded.

a maximum of four products is $\{(1 \rightarrow 2), (1 \rightarrow 3), (1 \rightarrow 4), (2 \rightarrow 1), (2 \rightarrow 3), (2 \rightarrow 4), (3 \rightarrow 1), (3 \rightarrow 2), (3 \rightarrow 4), (4 \rightarrow 1), (4 \rightarrow 2), (4 \rightarrow 3)\}$. And so forth, up to experimentally controlled maxima of n products. The rule-set size of the ALL chemistry grows rapidly as $(n^2 - n)$, whereas the rule-set size of SOLO H grows only as n .

The discovery and nurturance of cycles within randomized rule sets are as important to the creation of life within the ALL chemistry as they are in the more restricted SOLO H chemistry. In either chemistry, without emergent cycles of rule reproduction there are no positive feedbacks for growth, to combat endemic turnover or death. The difference is that in the ALL chemistry a great many distinct rule cycles are available to be discovered and reinforced, whereas in the SOLO H chemistry, there is only a single target type of cycle to try to grow and reinforce. This much larger space for exploration in ALL opens the door for the possible emergence of multiple production networks of linked rules, which may intertwine with each other. Given the emergence of multiple production networks, or ‘technologies’, the structure of how such multiple networks overlay and dynamically feedback into each other, positively or negatively, becomes a topological topic from which to derive the consequences for system robustness and evolvability (cf. Wagner and Altenberg 1996, Wagner 2005, and Fontana 200x, this volume).

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.⁴ 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 or organizations to emerge.⁵ 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.

CHEMISTRY MODELS OF ECONOMIC PRODUCTION: HYPERCYCLE AND ALL

We shall describe both variants of our chemical model of economic production in pseudo-algorithmic fashion, because we have implemented it in the form of an agent-

⁴ This is not only bounded rationality, this is the absence of consciousness altogether.

⁵ Hayek (1948) made a ‘self-organizing’ argument about the operation of markets similar to the one we are making about the emergence of markets.

based simulation.⁶ First we shall describe our core models of production and learning. These will give the logic of our core ‘dependent variable’: autocatalytic-network emergence. Then we shall describe experimental variations of our core model – type of chemistry, number of products, interaction topology, mode of learning, input environment, and input search method. These are the ‘independent variables’ that may affect the likelihoods of autocatalytic-network emergence and of the structures of any production networks and firms that do emerge.

The simplest version of our spatial hypercycle model has been solved analytically, with closed-form solutions already published (Padgett et al. 2003, pp. 861-62, 871-77). For this reason, we do not reproduce those mathematical results here.⁷

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, according to one of two families of chemistries: SOLO H and ALL.
3. Balls/products are indexed by $i = 1, 2, 3, \dots, n$. The parameter n serves two functions: it indexes each member of the two chemistry families, and it characterizes the relative ‘complexity’ of the particular rule set under investigation.
4. Rules/skills are contained in bins/firms. At the beginning of each simulation, skills are just randomly distributed across available firms, without any logic. The number of firms initially is large.
5. Bins/firms are arrayed on a spatial grid, with wrap-around boundaries. Each firm has eight possible nearest-neighbor trading partners – the so-called Moore-neighborhood structure of physical space.
6. 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

⁶ Our agent-based model, publicly available for both demonstration and open-source modification, 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. Peter McMahan has reprogrammed this Repast model in R; this version, available on the web site xxxx, was the version used to generate results reported in this chapter.

⁷ In economics, though not in physics, there exists 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. In particular the move from non-spatial homogeneity to spatial heterogeneity frequently causes problems for analytic tractability. 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.

- 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.
7. 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 on to a third firm or possibly back to the first.) In this way, transformed products pass through sequences or chains of skills.
 8. 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.

Core model of learning:

1. 'Learning by doing' is modeled in chemical fashion as follows: If one skill transforms a product and passes it on to a second skill that transforms it further ('can use it'), 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 – sender or receiver – 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

⁸ Final consumption is the output urn.

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

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 technological 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 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 – more than one loop of the same sequence in the case of SOLO H, or more than one loop of various sequences in the case of ALL. In either case, 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 cycles, but firms whose skill sets participate in production chains that include loops have the further capacity to keep renewing each other endogenously 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 cyclic production chains, which wend their way through firms, knitting them together in trade and helping them to reproduce each other through continuous learning.

Experimental variations:

There are six ‘independent variables’ – that is, experimental treatments in the simulation model – whose effect on the likelihood of finding and sustaining self-organized autocatalytic networks of skills will be explored in this chapter.

¹⁰ In future extensions of this model, we intend to add diffusion of skills among trading firms, in order to mimic ‘teaching.’ But that extension is not developed in this paper.

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

1. Type of Chemistry:

The SOLO H and the ALL sets of initial rules have already been described.¹²

2. 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 chapter there are always 200 specific instantiations of rules being scattered, chosen uniformly from the relevant chemistry's rule set. We vary the composition or 'complexity' of the rule set so scattered, 'complexity' being indexed by n . We shall vary n from 2 to 9: that is, within each of our chemistries, we shall explore 2-ball/product rule sets, 3-ball/product rule sets, and so forth, up to a maximum of 9-ball/product rule sets. Presumably the more complex and thin the resulting rule sets, the more difficult it will be to find and to sustain living production chains.¹³

3. Interaction topology:

The basic spatial topology for trading to be explored in this paper is the 10x10 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.¹⁴

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 or 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 or SOLO H rule 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) and Padgett et al (2003) have 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 complex chemical life is embodied. We shall reconfirm here both the Hofbauer and Sigmund (1988) non-spatial findings and the Padgett (1997, 2003) spatial findings, in a new context.

¹² We also explored a variant of ALL, called ALL (no reciprocity), explained in footnote 26 below.

¹³ In the extreme case of ALL 9-ball chemistry, for example, there are only $200 / (9^2 - 9) = 2.778$ initial instances of each specific rule being randomly scattered over the 100 firms. Rounding of numbers of specific rule instantiations (to 3 or 2 in this example), necessary to maintain a total of 200, was done randomly.

¹⁴ Padgett (1997) investigated 4-neighbor (von Neumann) neighborhoods. In future work, described in the next chapter, we shall investigate endogenous topologies. That extension will allow investigation of the emergence and impact of social, not just spatial, networks.

4. Learning/reproduction:

In the spatial topology setting, there are two variants of ‘learning by doing’ that can and will be explored:¹⁵

- (a) ‘Source reproduction’ is where the originating rule in a successful transaction is reproduced.
- (b) ‘Target reproduction’ is where the receiving rule in a successful transaction is reproduced.

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.¹⁶ Variation in mode of reproduction thus defines who benefits from the transaction.

We think of source reproduction as like ‘selfish learning,’ because the initiator of the successful transaction reaps the reward (like a teacher). And we think of target reproduction as like ‘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) and Padgett et al. (2003) 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, but extended with quantitative variation to more general chemistries.

In addition to these model variations, two more experimental manipulations will be performed here, which vary the input or resource environment in which autocatalytic networks grow. Such experiments were not possible in Padgett (1997), because in the original framework there was no explicit modeling of products or of product environments.

5. 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,

- (a) ‘rich’ input environments will be modeled by letting the input urn of

¹⁵ In non-spatial interaction, these two reproduction modes behave identically (see appendix in Padgett et al. (2003)). 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.

¹⁶ Of course the recipient (2→3) could easily turn into an initiator in the next tick, if a neighboring (3→4) is subsequently found.

resources contain all n possible inputs, never to be depleted even as products/resources are withdrawn; and

- (b) ‘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,

- (c) ‘endogenous’ input environments will be modeled by letting the input urn be reconstructed 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; and
- (d) endogenous environments can be rich or poor, depending upon initialization of the input urn, which endogenous gradually modifies.

Presumably, rich input environments are more congenial to autocatalytic-network 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. Consistent with arguments in this social-insect tradition, Padgett et al. (2003) discovered in hypercycles that stigmergy performed almost as well as fixed-rich in target reproduction, and that stigmergy even outperformed fixed-rich in source reproduction.

6. Input search:

The final experimental manipulation possible within our model varies the precision of search through the input environment:

- (a) ‘Random search’ is when an activated rule reaches into the input environmental urn and chooses inputs randomly, in proportion to what is there.
- (b) ‘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.¹⁷ Selective search is more like animal behavior.¹⁸ This is the only place in the model where we vary degree of intelligence. We originally expected the more intelligent selective-search procedure to outperform the stupid random-search procedure in finding and nurturing production hypercycles. Padgett et al. (2003) discovered to their surprise that selective search affected speed to equilibrium, but not any properties of the equilibrium itself. Search intelligence, in other words, speeded up evolution, but it did not alter the outcome of evolution. Because of this

¹⁷ Metaphorically we think of this as ‘the intelligence of an atom, bouncing around.’

¹⁸ Metaphorically we think of this as ‘the intelligence of a cow, looking for grass.’

earlier finding, we do not focus on this independent variable in this chapter. Instead we fix the search method at random, in order to focus our attention elsewhere.

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 even industrial districts can emerge and reproduce even under these minimalist assumptions. And we hope thereby to discover structural and interactive imperatives that help to foster economic self-organization.

RESULTS: HYPERCYCLE EMERGENCE IN SOLO H

Figure 1 presents a small sample of successful equilibria of our agent-based hypercycle model, as they would appear on the computer screen.¹⁹ ‘Success’ means that at least one autocatalytic network was found and survived the selection process to reliably reproduce itself.²⁰ To interpret the figure: (a) ellipses are surviving bins or firms; (b) numbers within the ellipses are surviving rules (a ‘51’ in this figure stands for ‘5→1’); (c) lines are products being passed between firms and their rules; (d) solid lines are trades that participate in cycles; (e) dotted lines are trades that do not participate in cycles (so-called ‘parasites’ or ‘free riders’). The entire ensemble resembles an industrial district of spatially contiguous firms, trading among each other through distributed production rules.²¹

Were one to observe on the computer screen the dynamics leading up to such equilibria, they would look like this: At first random rules pass around random products in all directions, like popcorn in a vat. Soon, however, where production-rule cycles randomly exist underneath of the apparent chaos, clusters of more fervent production-cum-trading activity tend to congeal, with firms outside these clusters starting to die off. Gradually, a period of inter-cluster competition ensues, with the final outcome being either no surviving cluster or only a single surviving cluster. If any cluster survives, usually it is the largest cluster that does. The outcome of no survivor is due to dynamic instability in the underlying distributed technology network – such a system just did not reliably reproduce its component rules, in the face of inevitable turnover and death. The outcome of no more than one surviving ‘cooperative’ cluster is due to competitive

¹⁹ Once again, to download and observe the models themselves, see the websites listed in footnote 6.

²⁰ [Peter: could you re-write this footnote for me, on how long models ran before ‘equilibrium’ was declared.]

²¹ Padgett (1997) interpreted pictures like these to be people in a workgroup or children in a playground. The principles of self-organization in our chemistry model are quite general, potentially applicable to many evolutionary scales, from chemistry to ecology. We do not mean to imply the patently absurdity that no important differences exist across scales, but in this line of modeling we choose to emphasize instead potentially universal principles common to life at any scale.

exclusion – if production-network clusters do not reproduce better than their neighbors, those neighboring networks extract more input resources, gradually starving them out of existence. Competitive exclusion in this model ultimately is the selection consequence of all rules facing the same input environment/urn.

This generic dynamics is observed in both SOLO H and ALL variants of the model, with varying probabilities. The experimental questions to explore thus become: What conditions affect the likelihood of any autocatalytic network of production rules growing and surviving? And given survival, how are production network structures shaped by experimental conditions? We shall answer these questions first for the old SOLO H chemistry (cf. Padgett et al. 2003), and then for the new chemistry ALL.

Figure 2 presents our results for SOLO H. These results reproduce, for longer run times,²² the findings in Padgett, Lee and Collier (2003). Figure 2 presents on its y-axis our first-order dependent variable: long-term probability²³ of hypercycle survival. Figures 3 and 4, below, will present comparable information for the ALL chemistry. All these present on their x-axes varying degrees of complexity in the simulated economies: simple 2-ball/product technologies, somewhat more complicated 3-product technologies, and so forth, up to our most complex 9-product rule set. Different lines within these graphs present the results of our various experimental manipulations: interaction topology, learning mode of reproduction, and input environment. As explained above, all these runs used the mindless random-search method of choosing inputs.

We shall unpack the findings in figure 2 one independent variable at a time.

A. 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 5-skills and above, survival rates of spatial hypercycles can become

²² Compare footnote 18 in Padgett et al (2003) to footnote 19 above. We have more powerful computers at our disposal this time, compared to 2003.

²³ Each of the points in these graphs represents the average of 30 simulation runs.

superior to survival rates of non-spatial hypercycles. This statement is especially true for target reproduction operating under rich environments. But it is also true for ‘stigmergy’ – that is, for source reproduction operating under endogenous environments (of either rich or poor initial configuration). This statement is not true for either mode of reproduction operating under fixed poor environments; those provide too few nutrients for any but the simplest 2-product hypercycles to survive.

Another way of expressing these findings is this: Non-spatial ‘freedom of trade’ of every skill with every other, with no firms to restrict and channel that trade, generates so much volatility in skill reproduction that high 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 complex economic production networks to be sustainable. 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. This is the chemistry reason for why spatially embedded firms exist in the first place – to break the complexity barrier.

The mechanistic reason for this result is that physical space ‘breaks the symmetry’ of homogeneous interaction, and thereby allows heterogeneous ‘memory’ of past interaction history to accrue within bins/firms. Through learning-by-doing, accretions of successes and failures in their past interaction become inscribed into the rule competencies of the two trading firms. The way this works is this: Successful transactions reproduce compatible skills that are located in neighboring firms. This sets up positive feedback between compatible-skilled neighbors: the more skills are activated, the more they reproduce, the more they reproduce, the more they are activated. The volume of skills in a firm at any given point in time thereby becomes the cumulative history of past interaction with its neighbors. Conversely, non-reproducing rules are ‘forgotten’ because all skills are killed off randomly in the population. Rules not participating in positive feedbacks with spatial neighbors thereby go extinct. The long-term success or failure of any given firm is the path-dependent consequence not only of that firm’s own history, but also of that firm’s neighbors’ histories.

None of this heterogeneous path-dependence would have occurred without physical space, or without some social-network functional equivalent to space (cf. Cohen et al. 2001). Spatial or social 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 trading then inscribes the memory of past interactive success into the structure of each co-adapting firm, thereby permitting the non-spatial complexity barrier to be breached.

A secondary mechanism behind the effectiveness of spatial clustering is chaining. The physical act of passing products around orchestrates sequences of learning. Not only do compatible neighbors generate positive feedbacks in their own growths, but also compatible neighbors trigger other compatible neighbors. Once hypercyclic clusters begin to emerge, micro feedback loops are evoked and orchestrated more efficiently.

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.

B. 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.²⁴ Without bins/firms, there is no difference between 'selfish' and 'altruistic' because without bins/firms there are no phenotypic actors to begin with, to be selfish or not.

In the production and nurture of spatial hypercycles, target reproduction is superior to source reproduction. This is shown in figure 2 by the fact that the survival plots of target reproduction are displaced to the right of the corresponding survival plots of source reproduction. In fixed poor environments, this difference is trivial because there almost everything dies. But in fixed rich environments, the difference is very dramatic. 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 mechanism that produces this superiority is repair. Target reproduction combats dynamic instability in a way that source reproduction does not. To repeat, the process of dynamic instability, causing hypercycles to crash, is this: if one skill reproduces too rapidly, competition drives other skills to zero, thereby breaking the reproductive loop of skills upon which all depend. Spatial topology distributes this dynamic into overlapping series of neighborhoods, thereby inducing local heterogeneity which may provide a partial buffer. But source reproduction, or selfish learning, does not really attack the basic dynamic instability itself. In source reproduction, an initial activated rule passes on its transformed product to a neighboring compatible rule, which causes the original activated rule to reproduce. Frequently activated rules thereby reproduce more frequently, often driving out of existence even compatible neighbors on whom they depend for their own survival. Like a cancer. As we shall see in the next subsection, endogenous environments sometimes can ameliorate the negative effects of this uncontrolled growth, but source reproduction in and of itself does not eliminate the basic instability problem.

In sharp contrast, an initial activated rule in target reproduction passes on its transformed product to a neighboring compatible rule, thereby causing the recipient (not the sender) rule to reproduce. Here the more frequently the initial sender rule is activated the more frequently the second recipient rule reproduces. This difference induces the following soothing homeostatic feedback: increased activation in the sender rule leads to increased reproduction in the recipient rule, which leads to increased probability of death and hence activation next time in the sender rule. In this way, hypercycles repair themselves. As the volumes of skills in a loop gets low, high volumes of compatible

²⁴ Padgett et al (2003, p. 871) demonstrated that in non-spatial topology, target and source reproduction mathematically become identical processes.

skills in neighboring firms reach in to those low-volume skills to build them back up. Threatening peaks and valleys along loops in the hypercycle are smoothed.

This simulates altruistic behavior, even though no skill or firm is trying to aid the public good. Target-reproduction repair does not guarantee that a hypercycle will survive, but it does directly alleviate through repair the dynamic instability problem that afflicts both the non-spatial and the spatial-source settings.

Padgett et al. (2003, p. 858) demonstrates analytically how this repair mechanism works, in a special simplified case of the model.

C. The Effect of Input Environment:

Figure 2 also points to the existence of a second repair mechanism, more relevant to source ('selfish') reproduction than to target ('altruistic') reproduction. The label 'stigmergy' in figure 2 and in all subsequent figures refers to the combination of source reproduction with endogenous input environment. Figure 2 demonstrates for SOLO H that, while stigmergy never reaches the nirvana of target reproduction operating in a rich environment, the addition of endogenous environment to source is capable of regulating source reproduction out of its otherwise self-destructive behavior.

The mechanism behind this surprising result is not direct repair of neighbor by neighbor, as in target reproduction, but rather indirect repair, via the intermediary of the input environment. As explained above, the basic problem is that source reproduction left to itself generates self-destructive ('cancerous') growth. The more skills are activated, the more they reproduce, the more they are activated, et cetera, until the neighboring partner upon whom that skill depends is destroyed. Endogenous environments do not eliminate this cancerous growth, but they help to control it in this way:²⁵ The more a skill is activated, the more it consumes its own input from the environment, and transforms it into something possibly useful for its compatible neighbor. If this product is not given directly to a compatible neighbor, then it is tossed back into the endogenous urn. 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 eliminate peaks and valleys around the hypercycle, as almost does the more direct skills-to-skills intervention of target reproduction – thus target reproduction remains superior to stigmergy. But the indirect skills-to-environment-to-skills method of regulation, induced by environmental endogeneity, can function to keep skill-volume peaks within bounds.

This is like what social insects do (e.g., Camazine et al. 2001). Bees communicate directly, but ants coordinate their behavior with one another indirectly through modifying their environment (for example through pheromones), in ways that feedback into their own behaviors. This leads not to static equilibrium behavior, but 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.

²⁵ This limiting of growth through nutrient starvation resembles anti-angiogenesis in real cancer treatment.

Supporting further this analogy 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 equilibrium. Such moving peaks are observed to be very sharp under selective search, whereas they are more gentle under random search.

No particular production output is favored over any other in the self-organizing model of this chapter. But were purposeful production to be introduced, the varying waves of production that hypercyclic production chains produce naturally in endogenous environments reveal that output could be adapted easily to changing circumstances in the short run. Hypercycles of production, in other words, generate flexible arrays of products, not just a single product. A potential downside might be that flexibility in output is a necessity: In the long run, output *has* to shift around, or else the natural repair mechanism of endogenous environment will be disabled.

D. The Effect of Search Intelligence:

The main finding about degree of search intelligence in Padgett et al. (2003) was a negative one. Contrary to our expectations, selective search did not improve the chances of hypercycle emergence over random search. Equilibria were reached more speedily under selective than under random search, but both the probabilities of survival and the structures of the final networks were the same. For this reason, we have chosen not to focus on this independent variable in this chapter.

Instead we just repeat our earlier conclusion about intelligence in search: “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 might have been the opposite: complexity is necessary for intelligence to emerge.” (Padgett et al. 2003, p. 863)

RESULTS: AUTO-CATALYTIC NETWORK EMERGENCE IN ALL

Figures 3 and 4 present the survival probabilities for the ALL chemistry, comparable to those just presented in figure 2 for SOLO H. In general, the relative effects of our independent variables in ALL are reassuringly similar to the rank ordering observed in SOLO H, but the absolute probabilities are changed dramatically. It is much easier for *something* to survive in ALL than it is for specifically targeted hypercycles to emerge in SOLO H. While initially surprising to us, in hindsight there is not much of a puzzle in the higher absolute values. Even in high-n ALL chemistries, where the potential for complexity is great, what actually survives is often quite simple. To take an extreme example for illustration, in a 9-product ALL chemistry, the equilibrium could turn out to be only the simpleton 2-product hypercycle: $\{(5 \rightarrow 8), (8 \rightarrow 5)\}$. Fortunately autocatalytic networks within ALL are not usually this boring, but if possible equilibria are

unconstrained then the likelihood of finding *something* is high. The real question for the ALL chemistry, therefore, is not the likelihood of producing something, but rather the likelihood of producing something complex and interesting.

In figure 4, we therefore re-present the same set of ALL-chemistry runs, but this time measuring only the probability of finding and nurturing to equilibrium ‘complex’ autocatalytic networks, where ‘complex’ for this purpose means “cycles involving three or more distinct rules.” Using this measure of survival, target reproduction in a fixed rich environment produces complex networks about 80% of the time, regardless of the number of balls/products. In contrast, source reproduction operating in a fixed rich environment produces complex networks only 10-20% of the time. Stigmegy, in either rich or poor initial environments, produces complex networks 30-50% of the time. Source reproduction in poor environments never produces complex networks, yet target reproduction operating in poor networks can do so, from 10% to 40% of the time. The symmetry-breaking and repair mechanisms lying beneath these comparative results have already been explained.²⁶

What are the autocatalytic networks that emerge in the ALL chemistry like? Figures 5 through 8 present structural properties of the autocatalytic networks that survive in ALL, given that they survived in the first place.

Figure 5 plots the average number of surviving firms that comprise emergent autocatalytic networks in ALL. Target reproduction generates, on average, quite robust ‘industrial districts’ of eight or so firms, regardless of chemistry complexity level, if the resource environment is rich. Even in poor resource environments, where the constraints on life are very severe, target reproduction can foster the emergence of small ‘industrial districts’ of eight, declining to four firms. Source reproduction, in contrast, produces mainly boring two-firm dyads, although at low levels of chemical complexity the number of surviving simple firms is higher. ‘Industrial districts’ is hardly an apt metaphor for such impoverished networks. Stigmergy – which combines source reproduction with endogenous environment – helps raw source reproduction, as expected, but not as dramatically as it did in SOLO H. In the ALL rule set there are many more rules than in the SOLO H rule set; therefore, the benefits from indirect repair are spread more thinly.

Figure 6 plots the number of distinct rules in ALL’s surviving autocatalytic networks, which we call the ‘rule complexity’ of the distributed production network. ‘Distinct’ means unique types of rules in the entire population. Many of these rules/skills

²⁶ We also ran a different procedure for evaluating survival rates above and beyond the lowest level of 2-product cycles. Namely, we implemented a modification of the ALL chemistry, which we label ALL (no reciprocity). This procedure kept the ALL rule set intact, but did not reward with reproduction any reciprocal dyad in which firms just passed the same two products back and forth. Turning rubber into tires, and then tires back into rubber, happens all the time in real chemistry, but seems rather pointless in human analogues. Somewhat surprising to us was the result that disabling low-level reciprocity did not much affect the ability of the ALL chemistry to generate more complicated trading patterns, one way or the other. The presence or absence of ‘pointless’ production behavior does not seem to matter for the emergence of more complex production behavior. Graphs proving these statements are available upon request.

are duplicated in different firms; hence the total volume of rules in the production networks is larger than is the number of distinct rules in the networks.

The rank ordering of experimental effects is similar for rule complexity as it was for population of firms, with target reproduction dominating this measure of complexity success, with stigmergy second, and with raw source reproduction coming up the rear – living perhaps, but at very low levels of complexity. Even more so than in SOLO H, target reproduction in ALL is the main route for generating complex autocatalytic networks. One interesting wrinkle emerges from examining the ratio of rule complexity to population (and even more so from examining the ratio of total volume of rules to population). For target reproduction in both rich and poor environments, this ratio rises as number of products increase. This means that more distinct rules on average become packed into surviving firms as the chemistry as a whole diversifies. That is, firms themselves become more complex: firms with more than one skill sit at the center of autocatalytic networks. This also was discovered in SOLO H chemistry (Padgett et al, 2003, pp. 864-5). Evolution therefore requires embodiment in two senses: physical space is necessary to break interaction symmetries, and phenotypic bins become more complicated as the production networks that evolve through them advance.

So far, our findings in the ALL chemistry differ quantitatively but not qualitatively from our earlier findings in the SOLO H chemistry. Scientific replications contribute to our confidence in the robustness of the findings, but they do not make for new news. The primary way in which the networks in ALL qualitatively extrapolate beyond the networks in SOLO H, however, is presented in figure 7. ALL chemistry differs from SOLO H chemistry in being capable, through its richness, of generating multiple types of co-existing production networks or technologies, not just one. It in fact does so to the extent measured in figure 7. SOLO H chemistry generated many overlapping cycles, as was apparent in figure 1, but these were always the same cycle, redundantly piled on top of each other. In ALL, by contrast, multiple *types* of network are symbiotically intertwined with each other, dynamically supporting and regulating each other's reproduction. This is not only emergence but also differentiation into multiple subsystems.

'Subsystems' are defined in figure 7 by 'distinct cycles'. A distinct cycle is a cyclic chain of production rules in which no rule is repeated. Distinct cycles can and do share rule or even rule-sequence components; this being the meaning of 'intertwined'. But in rule-chain ensemble distinct cycles are different technologies. Figure 7 reports the average number of subsystems or technologies, so defined, that existed in the surviving autocatalytic networks generated by the experiment indicated. Ours is the only experimental platform of which we are aware that spontaneously generates multiple types of networks.

The specific findings in figure 7 are these: Source reproduction rarely rises above its single type of network of minimal complexity. Stigmergy increases subsystem complexity on average to the co-existence of two types of production networks. Target reproduction in a poor environment also generates on average two subsystems. But target

reproduction, when it operates in a rich resource environment, proliferates networks into subsystem-complexity levels of three or even six distinct cycles! Obviously we are still far removed from the complexity of real organic chemistry, but at least our minimalist model offers some insight into how the rich combinatorics and diversity of living chemistries assemble themselves spontaneously for evolution to work on.

More fine-grained detail on the distinct cycles reported in figure 7 is provided in table 1. There ‘distinct cycles’ are broken down by the length of those production-rule cycles. One can see that not only does target reproduction generate more distinct cycles, it also generates longer and more elaborate distinct cycles.

Why is target reproduction the primary route to complex autocatalytic production networks in ALL? We have already provided one answer – repair. ‘Altruistic’ neighbors in target reproduction, unlike ‘selfish’ neighbors in source reproduction, directly reach in to reproduce low-volume rules threatened with extinction. But this repair answer applies equally to SOLO H and to ALL chemistries. The relative superiority of target reproduction, however, is even more pronounced in ALL than it was in the spatial hypercycle chemistry. Something about the more complex multiple-network context of ALL appears to boost the generic advantage of target’s direct-repair mechanism and to dampen (without eliminating) the generic advantage of stigmergy’s indirect-repair mechanism.

The reason for this extra boost is revealed in figure 8, which plots the volumes of ‘parasite rules’ produced by our various experiments. Recall from our discussion of figure 1 that ‘parasites’ were defined as surviving rules in an autocatalytic network that do not themselves participate in reproductive cycles. These are free riders, in social-science terminology, who live only by feeding on the (re)productive work of others. Figure 8 reveals parasites to be rampant in target reproduction and absent in all versions of source reproduction, including stigmergy. Source reproduction effectively eliminates parasitic free riders. Usually in economics this is considered to be a good thing because efficiency is thereby improved. But in living chemistries, unlike in neo-classical markets or prisoners’ dilemmas, aggressive weeding of free riders stops the emergence of all but the simplest forms of life.

Examination of visualizations of equilibrational ALL production networks, similar to the ones in figure 1, uncovered the reason why. While sometimes parasites in ALL shoot out into unproductive space, as they do in figure 1, other times parasites in ALL become the bridges for linking and symbiotically ‘coordinating’ multiple networks. This is not the only mechanism for multiple-network coordination – shared components or ‘multifunctionality’ is a second device for inter-linking distinct cycles. With parasites as bridges, however, entire cycles reach in to neighboring cycles to support them in the face of weakness. One cycle’s relative growth through competition bleeds into propping up its ‘competitor’, thereby enriching the multiple-cycle ensemble as a whole. The basic repair principle in target reproduction scales up, in other words, from neighboring bins/firms reaching in to help each other, to entire neighboring (through parasites) cycles reaching in

to help each other. There is no collective-good intentionality in this, of course. Self-repairs, both dyadic and inter-cyclic, are unintentional byproducts of target reproduction.

When viewed dynamically, the coordination of multiple networks and the tolerance of parasites thus are causally connected. Today's parasite, from the perspective of production, might evolve into tomorrow's regulatory superstructure, not through intent and certainly not through foresight, but rather because competitive subsystems tolerate each other better through them. In hindsight, one might have preferred to hold on to those parasites who later turned out to be regulators and to eliminate those that never did anything useful. But evolution has no hindsight, there being no end point to be privileged. Like it or not, the road to living complexity is littered with free riders. Certainly too many free riders can kill you, especially in harsh competition. But one polices parasites at the cost of restricting evolvability (cf. March 1982, 1991).

Multi-functionality is when 'parasites' (this time so-to-speak) themselves link up into higher order cycles, of greater length than the subsystems they span. These are the higher-order cycles tabulated in table 1. Overlapping distinct cycles have a 'conflict of interest.' On the one hand, they are competing for the same input resources, just like disjoint cycles. On the other hand, they share components whose growth or decline affects the survivability of both. Target versus source reproduction affects the relative balance of these two cross-pressures, with target tipping the balance toward multi-functional sharing of components, as well as toward parasite bridges. Symbiotic solutions for competitive multiple-network co-existence are not inevitably found, but more time is granted for groping toward solutions. Long cycles essentially are the epiphenomena of this chaining of shorter cycles via shared components. But the fact that these epiphenomena are themselves cycles gives a little boost to the shared components, and hence to the sustainability of multiple networks in the autocatalytic system as a whole.

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. What this chapter 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, Riolo and Axelrod 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:²⁷ dynamic barriers of 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 reap the reproductive rewards, complex technologies are more readily nurtured, because they repair themselves.²⁸ Free riding happens, but that does not threaten system stability. More important than policing free riders is enriching learning feedbacks among core actors, on whom parasites can safely feed if life is strong enough.

In the ALL chemistry, where many more production networks are possible, parasites not only do not impede, but also help the emergence and regulation of multiple networks. Altruistic learning permits the profusion of both parasites and the multi-functional sharing of components. In harsh competition between disarticulated networks, both of these can become liabilities. In harsh competition within articulated networks, however, parasites and multi-functionality can become bridging mechanisms through which emergent economies grope their way toward high-order symbiosis among differentiated technologies.

This conclusion is consistent with anthropological emphases on gift-giving in primitive economies (Mauss 1967 [1899], Sahlins 1972). It is also consistent with sociological observations about the ‘strange’ persistence of generous behavior in modern economies (Macauley 1963, Granovetter 1985, Uzzi 1996, Herrigel 1996, Padgett and McLean 2007). Ours may not be the only explanation of generosity. But repair – both between dyads and between cycles – is one evolutionary reason for the natural selection

²⁷ 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 organizations, the collectivity typically is not renegotiated anew. Rather, within constraints, component parts are transformed and molded into the ongoing flow of action.” (1997: 199-200)

²⁸ 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.

of altruism in competitive economies of all sorts. Altruistic learning stabilizes the reproduction of distributed technological skills, on which all depend.²⁹

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 chains of distributed skills alive. When rule chemistries are restricted, like in SOLO H, stigmergy may be enough to permit the evolution of complexity. As rule chemistries become richer, like in ALL, stigmergy still helps but may be insufficient by itself.

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, and the shapes of the rivers themselves. 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 orchestrate the social forces that produce it.

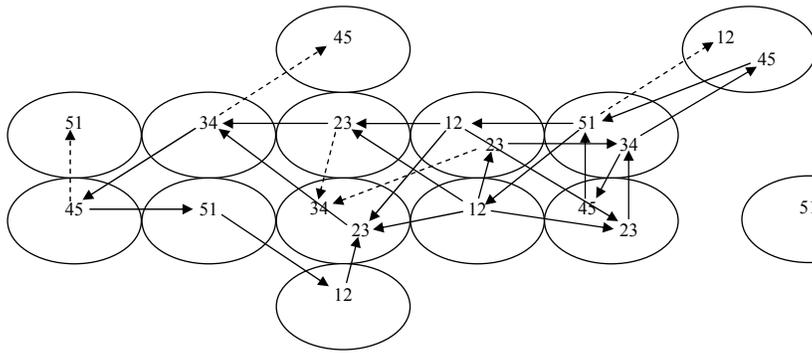
In this article, we have developed a few simple tools, imported from chemistry, that enable us to investigate systematically the co-evolution of distributed technology and social organization. Extreme assumptions about the absence of consciousness are implied by our specification. The payoff of such extreme simplification is the discovery of three social-organizational principles enabling technological evolution. How robust such principles are to alternate specifications remains 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.³⁰

²⁹ This may come as news to some rational choice theorists, but it will not come as a surprise to parents.

³⁰ March and Simon went beyond our point also to argue that social structures enable human cognition, by mapping the world down to levels we can comprehend.

Figure 1. Representative 5-Skill Hypercycles at Equilibrium: Target Reproduction

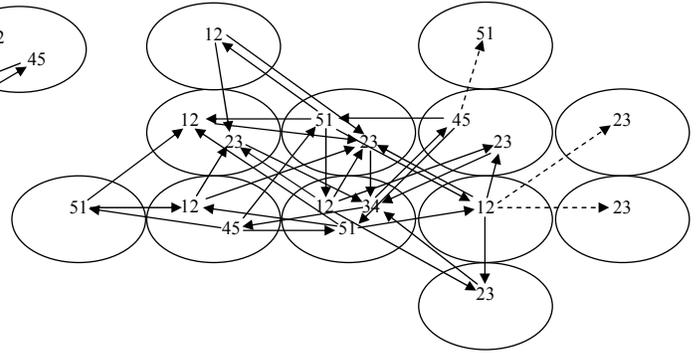
Fixed-Rich Environment; Selective Search:



		13			16
11	23	13	12	25	
18	13	27	3	20	
		6			

Hypercycles= 7

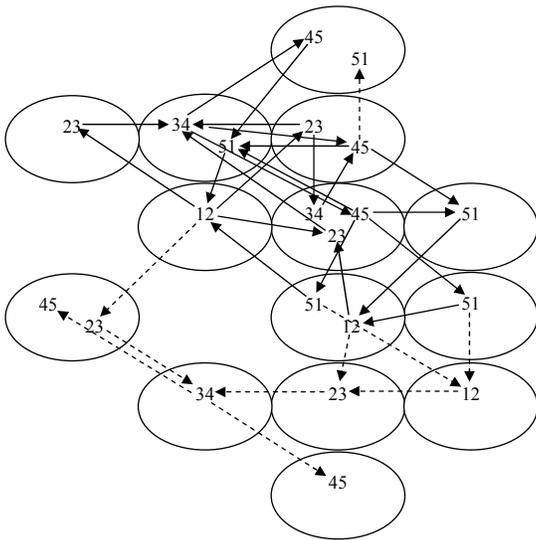
Fixed-Rich Environment; Random Search:



		7		4	
	20	42	23	7	
7	26	38	14	6	
				6	

Hypercycles= 39

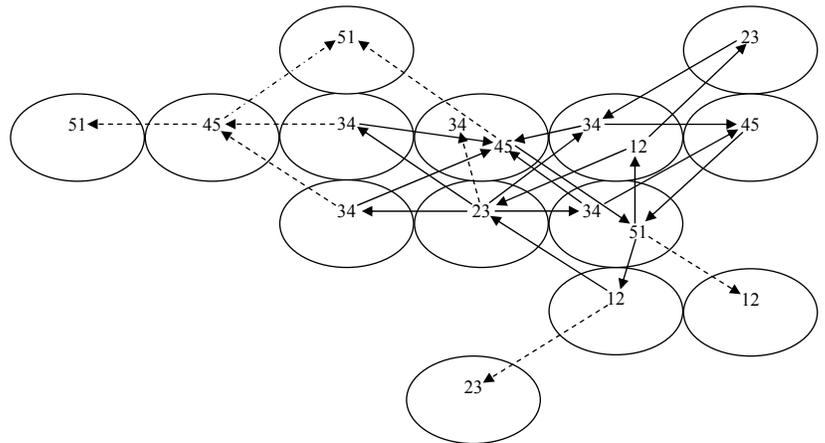
Endogenous Environment; Selective Search:



		15		
10	29	27		
	13	36	12	
14		11	7	
	6	9	4	
		7		

Hypercycles= 17

Endogenous Environment; Random Search:



		14		9
3	7	14	32	29
		14	12	34
				11
				6
		5		

Hypercycles= 14

N.B. 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.

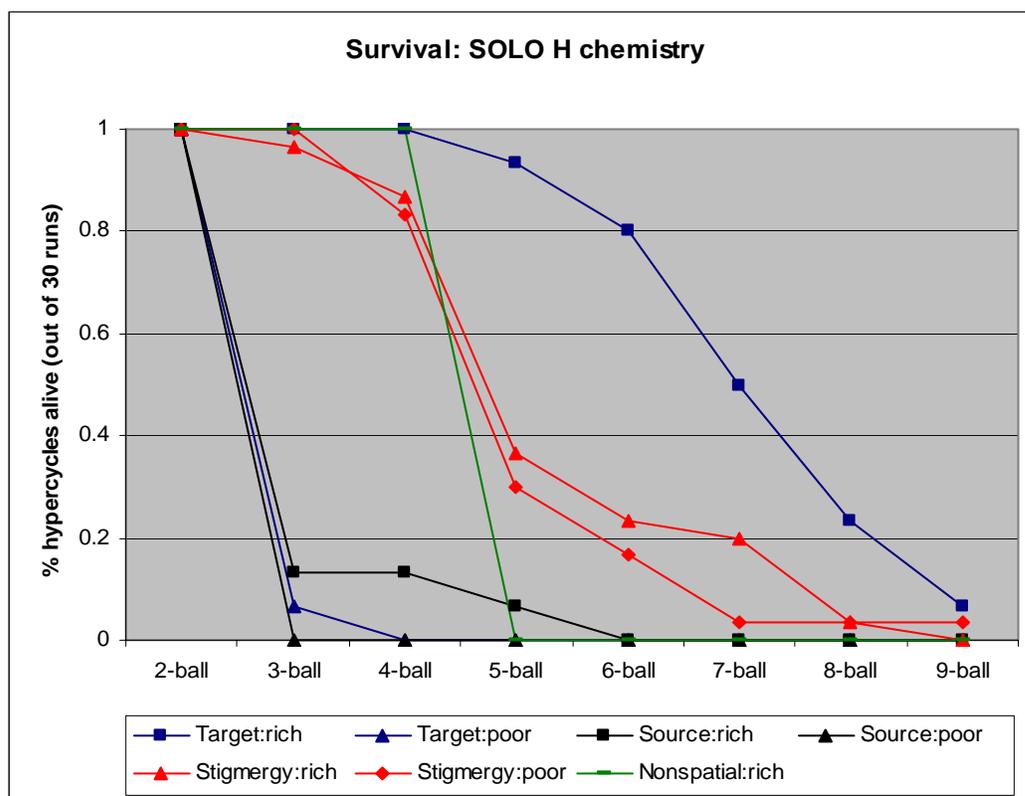
Figure 2. Survival of auto-catalytic networks: Hypercycle (i.e., SOLO H) chemistry

Figure 3. Survival of auto-catalytic networks: ALL chemistry

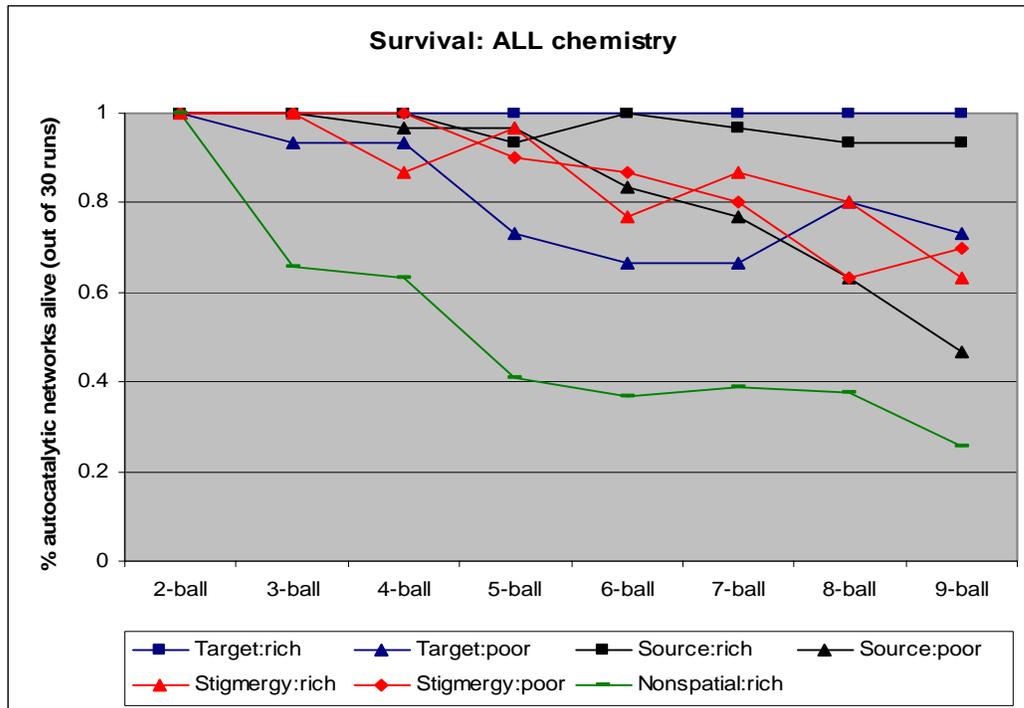


Figure 4. Survival of auto-catalytic networks: ALL chemistry, 3+ cycles

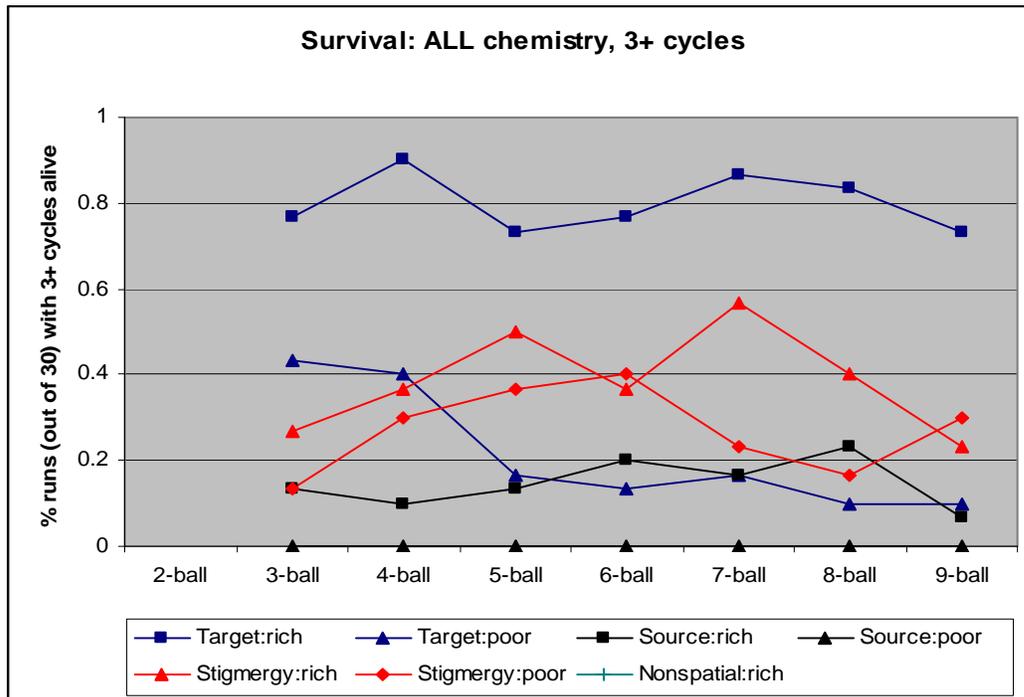


Figure 5. Population: ALL chemistry

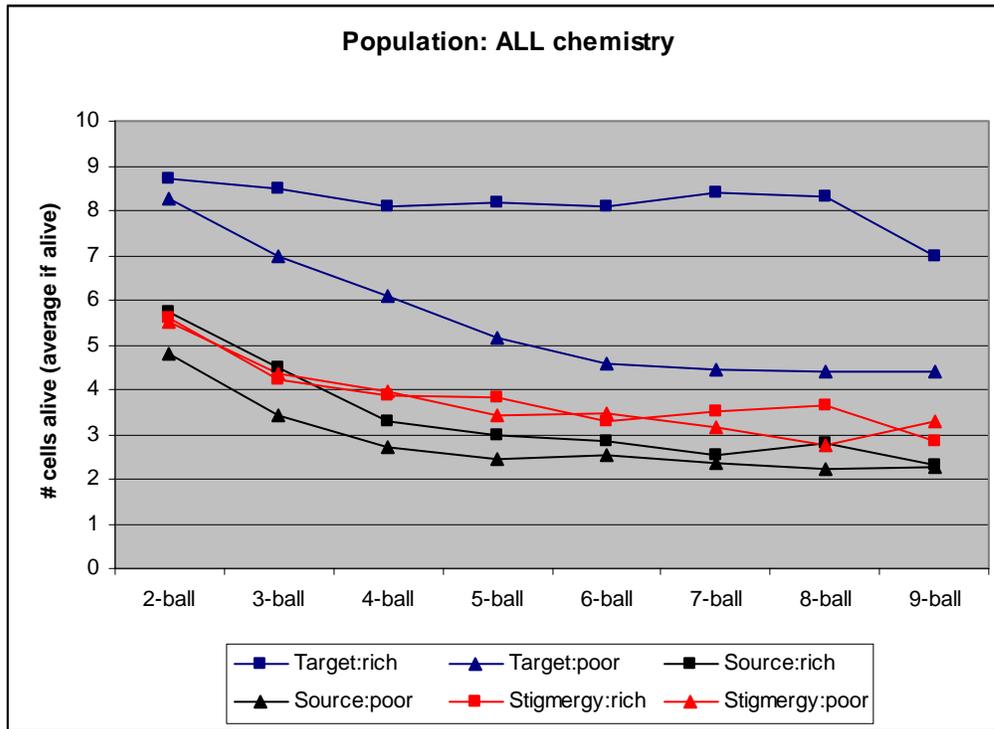


Figure 6. Rule complexity: ALL chemistry

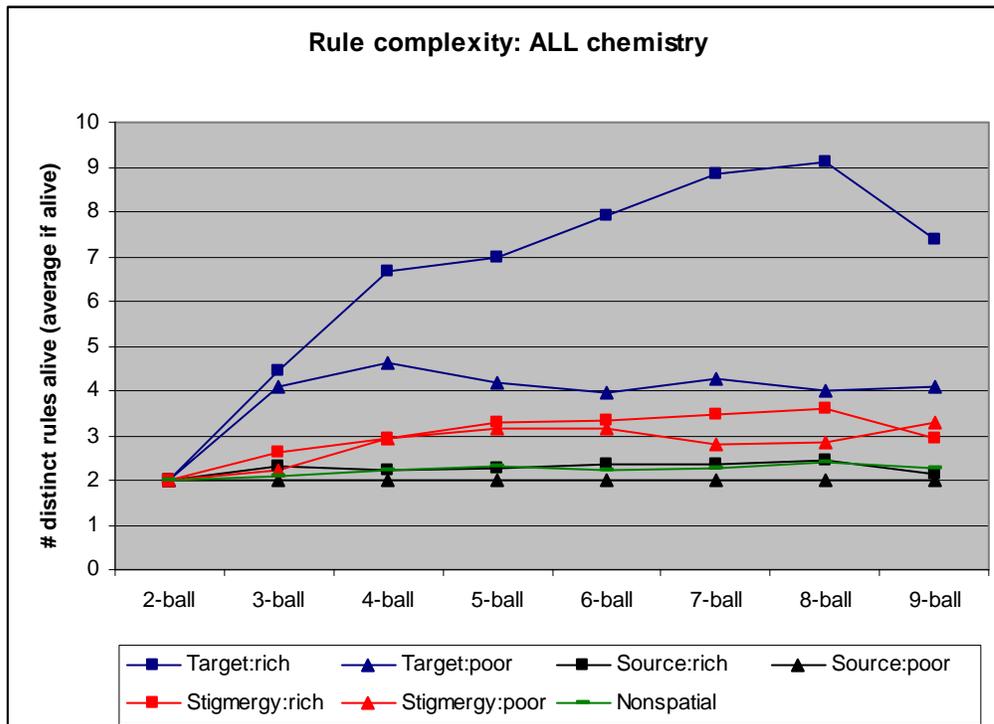


Figure 7. Subsystem complexity: ALL chemistry

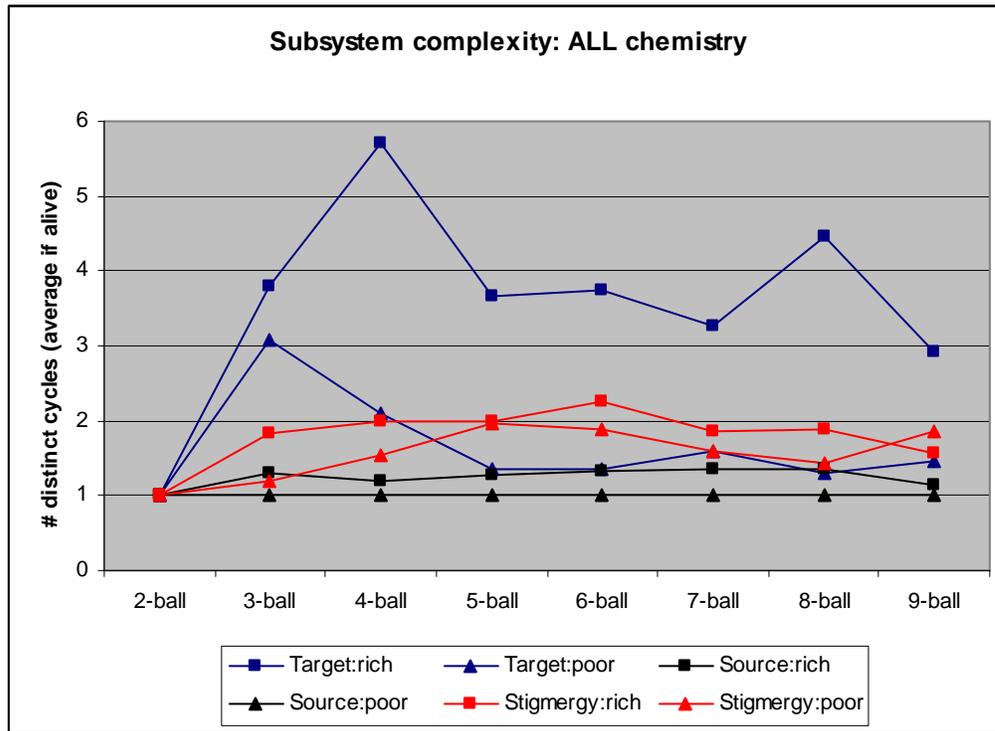
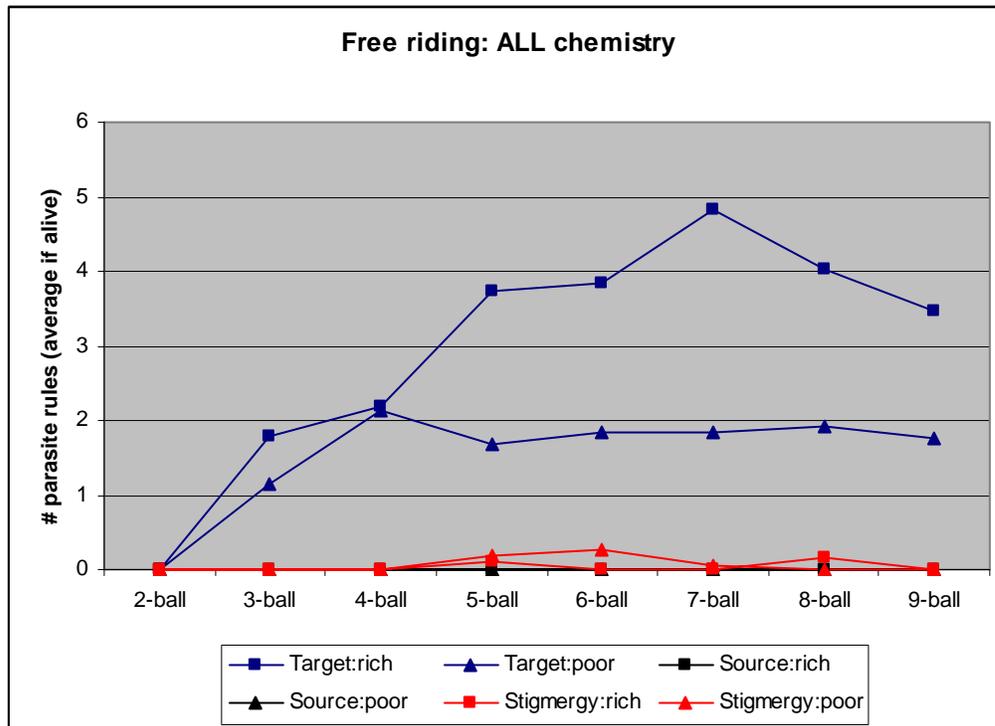


Figure 8. Free riding: ALL chemistry



I. Rich Environment & Random Search:

A. ALL Chemistry:

Network Complexity: cycle success rates (out of 30 runs):

Target:	any-cycle	2-cycle	3-cycle	4-cycle	5-cycle	6-cycle	7-cycle	8-cycle	9-cycle
2-ball	30	30							
3-ball	30	30	18	15		8			
4-ball	30	30	20	20	11	9	4		
5-ball	30	30	13	20	6	3	1		
6-ball	30	27	17	18	8	5	2		
7-ball	30	26	20	13	8	3		1	
8-ball	30	28	14	20	4	9	3	2	1
9-ball	30	28	10	15	2	1	1		

Network Complexity: cycle success rates (out of 30 runs):

Source:	any-cycle	2-cycle	3-cycle	4-cycle	5-cycle	6-cycle	7-cycle	8-cycle	9-cycle
2-ball	30	30							
3-ball	30	30	1	4					
4-ball	30	30		3					
5-ball	28	28		4					
6-ball	30	29	1	5					
7-ball	29	29		5					
8-ball	28	26	2	5					
9-ball	28	28		2					

Network Complexity: cycle success rates (out of 30 runs):

Stigmergy	any-cycle	2-cycle	3-cycle	4-cycle	5-cycle	6-cycle	7-cycle	8-cycle	9-cycle
2-ball	30	30							
3-ball	30	30	2	8		1			
4-ball	26	26	1	10		1			
5-ball	29	29	7	10	1	1			
6-ball	23	23	6	7	1	2			
7-ball	26	22	11	5	1	2			
8-ball	24	23	7	5	1	3			
9-ball	19	18	4	3		1		1	

References

- Adams, Robert McCormick (1966), The evolution of urban society: early Mesopotamia and prehispanic Mexico, Aldine: Chicago, Ill.
- Adams, Robert McCormick (1996), Paths of fire: an anthropologist's inquiry into Western technology, Princeton University Press: Princeton, N.J.
- Boerlijst, M.C. and P. Hogeweg (1991), 'Spiral wave structures in pre-biotic evolution: hypercycles stable against parasites,' Physica B, 48, 17-28.
- Bonabeau, Eric, Marco Dorigo and Guy Theraulaz (1999), Swarm Intelligence: From Natural to Artificial Systems, Oxford University Press: Oxford, U.K.
- Camazine, Scott, Jean-Louis Deneubourg, Nigel R. Franks, James Sneyd, Guy Theraulaz and Eric Bonabeau (2001), Self-Organization in Biological Systems, Princeton University Press: Princeton, N.J.
- Cohen, Michael D., Rick L. Riolo and Robert Axelrod (2001), 'The role of social structure in the maintenance of cooperative regimes,' Rationality and Society, 13, 5-32.
- Coriat, Benjamin and Giovanni 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, U.K.
- Davidow, William (2000), 'Social tornadoes: interconnection, instability, and social policy,' unpublished manuscript.
- Dosi, Giovanni, Renato Giannetti and Pier Angelo Toninelli (1992), Technology and Enterprise in a Historical Perspective, Clarendon Press: Oxford, U.K.
- Dosi, Giovanni, Richard Nelson and Sidney Winter (2000), The Nature and Dynamics of Organizational Capabilities, Oxford University Press: Oxford, U.K.
- Eigen, Manfred (1971), 'Self-organization of matter and the evolution of biological macromolecules,' Naturwissenschaften, 58, 465-523.
- Eigen, Manfred and Peter Schuster (1979), The Hypercycle: A Principle of Natural Self-Organization, Springer-Verlag: New York.
- Elster, Jon (1983). Sour Grapes: Studies in the Subversion of Rationality. Cambridge University Press: Cambridge.

- Fontana, Walter F. and Leo W. Buss (1994), 'The arrival of the fittest: toward a theory of biological organization,' Bulletin of Mathematical Biology, 56(1), 1-64.
- Fontana, Walter F (200x). 'The Topology of the Possible,' this volume.
- Granovetter, Mark (1985), 'Economic action and social structure: the problem of embeddedness,' American Journal of Sociology, 91, 481-510.
- Hayek, Friedrich (1948), Individualism and Economic Order, University of Chicago Press: Chicago, Ill.
- Herrigel, Gary (1996), Industrial Constructions: The Sources of German Industrial Power. Cambridge University Press: Cambridge.
- Hofbauer, Josef and Karl Sigmund (1988), Dynamical Systems and the Theory of Evolution, Cambridge University Press: Cambridge, U.K.
- Hughes, Thomas P. (1983), Networks of Power: Electrification in Western Society, 1880-1930, The Johns Hopkins University Press: Baltimore, Md.
- Hughes, Thomas P. (1987), 'The evolution of large technological systems,' in Wiebe Bijker, Thomas Hughes and Trevor Pinch, eds., The Social Construction of Technological Systems, MIT Press: Cambridge, Mass.
- Hutchins, Edwin (1995), Cognition in the Wild, MIT Press: Cambridge, Mass.
- Jain, Sanjay and Sandeep Krishna (1998), 'Autocatalytic sets and the growth of complexity in an evolutionary model,' Physical Review Letters, 81, 5684-87.
- Kauffman, Stuart A. (1986), 'Autocatalytic sets of proteins,' Journal of Theoretical Biology, 119, 1-24.
- Kauffman, Stuart A. (1993), The Origins of Order: Self-organization and Selection in Evolution, Oxford University Press: New York, N.Y.
- Macauley, Stuart (1963), 'Non-contractual relations in business: a preliminary study,' American Sociological Review, 28, 55-67.
- March, James G. and Herbert A. Simon (1958), Organizations, Wiley: New York, N.Y.
- March, James G. (1982). "The technology of foolishness," in J.G. March & J.P.Olsen (eds.). Ambiguity and Choice in Organizations. Universitetsforlaget, Bergen.
- March, James G. (1991). "Exploration and Exploitation in Organizational Learning," Organization Science, Vol. 2, No. 1, 71-87.

- Malerba, Franco and Luigi Orsinigo (1993), 'Technological regimes and firm behavior,' Industrial and Corporate Change, 2, 45-71.
- Maturana, Humberto and Francisco Varela (1980), 'Autopoiesis: the organization of the living,' in Autopoiesis and Cognition, D. Reidel: Dordrecht, Holland.
- Mauss, Marcel (1967 [1899]), The Gift: Forms and Functions of Exchange in Archaic Societies, Norton: New York, N.Y.
- May, Robert M. (1974), Stability and Complexity in Model Ecosystems, Princeton University Press: Princeton, N.J.
- McKelvey, Maureen (1997), 'Co-evolution in commercial genetic engineering,' Industrial and Corporate Change, 6, 503-532.
- Murman, J. Peter (2003), Knowledge and Competitive Advantage: The Co-evolution of Firms, Technology and National Institutions. Cambridge University Press.
- Nelson, Richard R. (1994), 'The co-evolution of technology, industrial structure, and supporting institutions,' Industrial and Corporate Change, 3, 47-64.
- Nelson, Richard R. (1995), 'Recent evolutionary theorizing about economic change,' Journal of Economic Literature, 33, 48-90.
- Nelson, Richard R. and Sidney G. Winter (1982), An Evolutionary Theory of Economic Change, Belnap Press of Harvard University Press: Cambridge, Mass.
- Padgett, John F. (1986). "Rationally Inaccessible Rationality," Contemporary Sociology 15: 26-28.
- Padgett, John F. (1992), 'Learning from (and about) March,' Contemporary Sociology, 21, 744-749.
- Padgett, John F. (1997), 'The emergence of simple ecologies of skill: a hypercycle approach to economic organization,' in W. Brian Arthur, Steven N. Durlauf and David A. Lane (eds.), The Economy as an Evolving Complex System II, Addison-Wesley: Reading, Mass.
- Padgett, John F. and Paul McLean (2006). 'Organizational Invention and Elite Transformation: The Birth of the Partnership system in Renaissance Florence,' American Journal of Sociology, 111: 1463-1568.
- Padgett, John F. and Paul McLean (2007). 'Economic Credit in Renaissance Florence,' under review at Journal of Modern History.
- Powell, Walter W. (1996), 'Inter-organizational collaboration in the biotechnology

- industry,' Journal of Institutional and Theoretical Economics, 152, 197-215.
- Prigogine, Ilya and Paul Glansdorff (1971), Thermodynamic Theory of Structure: Stability and Fluctuations, Wiley: New York, N.Y.
- Rosenkopf, Lori and Michael 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, Charles F. (1994), 'Learning by monitoring: the institutions of economic development,' in Neil J. Smelser and Richard Swedberg (eds.), The Handbook of Economic Sociology, Princeton University Press: Princeton, N.J.
- Sabel, Charles F. and Jonathan Zeitlin, eds. (1997), World of Possibilities: Flexibility and Mass Production in Western Industrialization, Cambridge University Press: Cambridge, U.K.
- Sahlins, Marshall (1972), Stone Age Economics, Aldine: New York, N.Y.
- Stadler, Barbel M.R. and Peter F. Stadler (2002), 'Molecular Replicator Dynamics,' Santa Fe Institute working paper 02-09-049.
- Uzzi, Brian (1997), 'The sources and consequences of embeddedness for the economic performance of organizations,' American Sociological Review, 42, 35-67.
- Warlien, Massimo (1995), 'Hierarchical selection and organizational adaptation,' Industrial and Corporate Change 4: 161-186.
- Wagner, Andreas (2005). Robustness and evolvability in living systems. Princeton, N.J.: Princeton University Press.
- Wagner, Gunter and Lee Altenberg (1996). "Complex Adaptations and the Evolution of Evolvability," Evolution 50: 967-76.