

Simulating Shocks with the Hypercycles Model of Economic Production

Christopher Watts¹, Claudia R. Binder¹

¹*Human-Environment Relations, Department of Geography,
Ludwig-Maximilians University, Munich, Germany
christopher.watts@geographie.uni-muenchen.de*

Abstract: Padgett's hypercycles model of economic production (Padgett et al. 2003) is an agent-based simulation that demonstrates the emergence of novel forms of social organisation among limited-intelligence agents. In this paper we argue that it can be extended to explore issues of resilience and the responses to shocks in complex systems of agent-environment interactions. Based on learning by doing, a simple search heuristic enables a population of firms to adapt and find system states containing a collectively self-sustaining subset of firms with production rules and mix of materials in the environment. The self-sustaining systems Padgett finds persist in spite of variability from stochastic processes in the model. However Padgett has not yet explored the model's responses to exogenous changes. We here outline experiments to simulate exogenous shocks. So far, the systems have proved highly resilient to the quantitative shocks, that is, to quantitative changes of various sizes in stock levels of rules and of environmental materials. However, systems were less resilient to qualitative changes to the structure of the emergent systems, namely in this paper changes to network links between firms. Suggested future developments include investigating the analogy between resilience and adaptive learning or search, and interpreting the model in the contexts of material flow analysis and industrial symbiosis.

Keywords: Resilience; shocks; adaptation; simulation models

1 INTRODUCTION

1.1 Simulating shocks and resilient systems

This paper summarises an approach to simulating on a computer exogenous shocks applied to a previously self-sustaining system, and exploring the *resilience* of that system, or its ability to return to something like its previous state following the extrinsic perturbation [Holling 1973].

Simulation models can give us insights into the potential outcomes of events that cannot be explored in real-world systems, such as rare events and events with negative consequences that we wish to avoid if possible. Social-ecological systems (SES) [Berkes et al. 2003] in particular are, for practical and ethical reasons, not generally something we wish to experiment with, and their complexity can lead to even well-intentioned interventions having detrimental effects on the sustainability of the system. Simulation experiments on an artificial SES would seem to carry no such risks. But designing a realistic model of a SES may be difficult. As with other systems of complex adaptive systems, a model of an SES will include many components with multiple interdependencies between them. Although empirical studies of an SES can yield much information on these components and their interdependencies, there is a risk that elements of the system missing from the data may play a crucial role in its functioning. Our attempt

at representing a currently sustained SES may result in a simulated SES that is not itself as self-sustaining as the real one.

For this reason, we propose here using an abstract, “toy” model of an economic production system, whose tendency to enter relatively self-sustaining states is already understood. To these production systems we may apply shocks in the form of sudden changes in components, structures or stock levels, and explore the ability of the systems to return to their previous level of relative stability. For this purpose we have selected a model of economic production by Padgett et al. [2003].

1.2 Padgett’s hypercycles model of economic production

Padgett’s hypercycles model draws upon ideas from theoretical biology, especially Eigen’s [Eigen & Schuster 1979] concept of hypercycles and Kauffman’s work on auto-catalysis [Kauffman 1996; 2000], both introduced as part of explanations of the origins of order, especially of living organisms. Padgett provides an interpretation of the model as an economic production system. It simulates a world of firms with variable amounts of production capabilities, contained in a common environment containing variable amounts of production materials. This is much simpler than any SES, but does capture the idea of a system of interdependent parts with some collective ability to adapt to their circumstances and show persistent order. Much interest in these systems concerns their ability to reach, without planned interventions from an observer or designer, self-sustaining sets of states – i.e. emergent system-level behaviour varying due to random processes but over time varying within limits. Earlier, analytical work on hypercycles models had suggested this ability was highly limited and not present among even moderately complex systems [Eigen & Schuster 1979; Hofbauer & Sigmund 1988]. In an agent-based model Padgett was able to demonstrate that by distributing production capabilities among multiple actors linked by a network structure, the complexity limit to emergent order could be raised to some extent [Padgett 1997]. In addition, the limit could be raised further by two variations on the adaptive behaviour of the model, namely how firms benefitted from learning by doing – the main adaptive or search mechanism in the model - and the endogeneity of its environment [Padgett et al. 2003].

In connection with work by one of the current authors on the simulation of innovation [Watts & Gilbert forthcoming], a *NetLogo* version of Padgett’s model has been developed. This reproduces results obtained by Padgett and collaborators with their own versions written in *Java* and *R*. In particular, it can generate self-organising and self-maintaining production systems of agent firms with little or no intelligence behind their actions. In the rest of this paper, we describe the model’s main mechanisms and our representation of shocks to the system, together with some illustrative experiment results obtained so far that point to the systems being relatively resilient, due to their “adaptive capacity” [Berkes et al. 2003] and their redundancy of parts, but they suffer noticeably when network structures between firms are altered. We conclude with some suggestions as to what this implies and how the work can be developed further.

2 METHOD AND MODEL

2.1 How the basic model works

The main entities in the model are firms and a common environment. Each firm has stocks of instances of “skills” or “production rules” of multiple types. Firms also have positions in a network, for which a regular two-dimensional grid lattice is used, with eight neighbours per firm (“Moorean architecture”). The environment, to which all firms have equal access, contains stocks of products, or raw materials, of multiple types. A production rule is a simple transformation rule of the form, “given

a product of type x , transform it into a product of type y ". The total range of types of production rule is the model's "chemistry". Following Padgett we focus our exploration on a simple chemistry in which the types of rule are "0 -> 1", "1 -> 2"... , "n-2 -> n-1" and "n-1 -> 0", where "n" is the number of types of product or material. If all these rules are employed in turn in a production cycle (called a "hypercycle"), the cycle will have begun and ended with the same product type. This may perhaps be thought of as external customers supplying money for a final product to the suppliers of the raw materials that initiated the cycle. The purpose of Padgett's original models was to explore the circumstances under which a population of firms with production rules could self-organise into a subset of firms and rules with at least one such hypercycle realised on a reliable basis. An example system is shown in figure 1, containing 7 firms representing 40 hypercycles.




Figure 1 An example self-sustaining production system, containing 7 firms with rules realising 40 hypercycles. Firms are shown with their stocks of rules. The format of a rule stock is [<output product type> <number of instances>]. 5 firms here have rules of more than one type. The firm in the bottom-right-hand corner is a parasite, not contributing to any hypercycle but receiving inputs from its left-hand neighbour who is contributing to one.

Product transformation and transfer, and rule learning and decay form the key mechanisms in generating self-organised hypercycle systems. The simulation models attempts at production runs. During production runs products are transformed according to production rules, and transferred between firms along their network links. As an example, in figure 1 the firm in the bottom left-hand corner can take a product of type 0 from the environment, transform it into a 1 and transfer it north to a firm who transforms it into a 2. Firms being supplied with the input of one of their rules are able to improve their capability in that rule – a process of learning by doing. Rules not used may be lost, however, due to a decay process.

In more detail, each time step or iteration, one attempt is made at a production run. The simulation samples one of the instances of a production rule to be found among the firms' stocks, and samples one item of raw material from the environment. If the raw material is not of the type required by the production rule, it is returned to the environment and the production run attempt ends. If the material does suit the production rule, a transformation occurs according to the rule and undertaken by the firm whose rule stocks it came from. The result of the transformation, a new product, is then transferred by the firm to one of the firm's neighbours, chosen at random and without preference or intelligence. On receiving a product, a firm checks its own rule stocks for one with this product as input. If

none such is found, the product is deposited in the environment and the production run ends. Otherwise, the product is transformed according to whatever rule has been found, and a new product generated which may then itself be transferred to a randomly chosen neighbour as the production run continues. A production run ends whenever the initiating firm receives a product of the same type as the raw material it started with. The received product is then returned to the environment, perhaps, as we suggested, symbolising financial resourcing. The long-term availability of raw materials is assumed to be indefinitely large for these simulations.

Learning by doing occurs whenever one firm transfers its output to a firm able to make use of it, i.e. to a firm with at least one instance of a rule taking this product as input. Under so-called “target reproduction”, or “altruistic learning”, the receiving firm is able to improve its capability. This is simulated by increasing by one instance the firm’s stock of the rule to be used. Whenever this happens, to keep constant the total number of rule instances in the system, one randomly chosen instance of a rule “decays”, or disappears from the system. This could represent expertise being forgotten, experienced staff leaving, or resources being reallocated to higher priority uses. If the firm to which the decayed instance belonged is now left with no instances of any rule, that firm “dies”, or leaves the system. Its network links to other firms are also removed, potentially leaving some firms cut-off from others. The disappearance of under-used rule-types and firms means that the remaining firms and rule-types are now more likely to receive use in future. Thus, in a process akin to stigmergic learning among ant colonies [Corne et al. 1999], previously successful cycles of firms and rules become more successful, and gradually become more visible as well. The system in figure 1 has emerged, after 50000 iterations, from an initial population of 100 firms, each with 2 rule instances of randomly set type.

In the *NetLogo* version, charts are updated periodically with the current stock levels of materials in the environment and rules among the firm population. In addition, a test is made for the presence of complete hypercycles among the network of remaining firms and rules. If there are no longer any hypercycles present, the simulation run halts. Otherwise, a fixed number of simulation iterations is run, known previously to almost always be sufficient for the emergence of self-sustaining hypercycle systems.

This concludes our summary of the original model. Padgett’s papers [Padgett 1997; Padgett et al. 2003] examine several variations on the mechanism described, and vary the number of types of products (“ n ”) in the production rules. In the space permitted by a conference paper, we cannot explore all of these parameter settings. For our purpose of illustrating the principle of using the hypercycles model to study resilience to shocks, we chose one point in the parameter space: “solo” chemistry, a regular, eight-neighbour network, a “rich”, “endogenous” environment, random selection of initial material, and “target reproduction” as learning. Those interested in what this means and what alternatives existed must consult the source papers. The number of types of product, n , and hence number of types of rule (“rule complexity”), was set to 5. This is known to almost always result in a self-sustaining hypercycle system surviving at the end of 50000 iterations, given these other parameter settings, but the survivability of systems declines rapidly as n increases beyond 5. As in Padgett’s experiments, 100 firms are each given two rule instances which may or may not be of the same type of rule. Hence there are 200 rule instances in total across the system. The environment is initialised with 200 instances for each of the five types of product material.

2.2 Simulated shocks and their impact

In our extension to the basic hypercycles model, we simulate the emergence of a self-sustaining hypercycle system, apply some change to the system representing

a shock, then continue the simulation for some more iterations and record how the final system state differs from the state immediately prior to the shock.

Although relatively abstract and simple, Padgett's model contains several types of component which might be altered in a simulation of a shock. In our experiments, we tested three types of shocks which can be related to real world situations of firm networks. Just as real materials may suddenly become scarce or plentiful, so in the model stock levels of materials in the environment might be altered. *Environmental shocks* are the removal or addition of given quantities of a randomly chosen type of product material from the environment's stock levels. This can be related to the impact of resource scarcity or sudden surplus of a resource. Two versions of environmental shock were tested. In the first the removed or added product instances are redistributed to or from randomly chosen other product stocks, keeping the total number of instances in the environment constant. In the other, no redistribution occurs, and the total number in the environment changes. A *rule shock* is defined as the removal or addition of given quantities of instances of a randomly chosen rule held by a single firm. Just as firms may suddenly lose capabilities through machine breakdowns and staff departures, or gain them through mergers and acquisitions, so in the model stock levels of rules held by firms might be altered. Again two versions of this shock were tested: one in which the total number of rule instances in the system was kept constant at 200 by using random rule cloning and decay. In the other version, the total number of rules was allowed to be changed by the shock. Finally, *Link shocks* are the random rewiring of given numbers of randomly chosen network links. Links represent perhaps the spatial proximity relations between firms' production sites, their transport links, or the inter-personal relations employed when suppliers approach potential customers.

Rule shocks and environment shocks will both alter the chances of particular types of product and rule being sampled at the start of the simulation iteration. Link shocks alter to whom products are transferred, and hence which firms are likely to learn by doing. Rule shocks also mean that a firm can potentially lose all its instances of a particular type of rule, and hence loses the ability to process a type of product. If it has no other types of rule, it will disappear from the network.

Real firms might exit the market for reasons external to the system, or new entrants might arrive with some basic capability – i.e. non-zero stock level – in at least one type of rule, and occupy a new position in the network. However we leave these possibilities for future work.

For the purposes of this paper, we focused on 12 hypercycle systems to administer shocks to, of which one is depicted in figure 1. These were each the result of 50000 simulation iterations or production run attempts. By repeating random number seeds the simulation model is able to reproduce these systems exactly. From each of these 12 starting points, 100 simulation runs, each of 50000 further iterations, were then made for each type of shock, having reinitialised the random number generator with a new random seed for each continuation run. Statistics were recorded for comparing final system state with pre-shock state, including the numbers of firms, instances of rules and firm-rule-type combinations. Also included was the presence or absence of hypercycles, from which the chance of *hypercycle survival* is estimated as the percentage of the 100 continuation runs containing self-sustaining hypercycles. A fall in the number of firm-rule-type combinations would represent a qualitative change in the system, since it implies at least one firm has lost all capability in at least one rule or skill. Firms' stock levels of rules may vary over time, but they cannot recover if they hit zero.

3 RESULTS

3.1 Resilience in response to stock level changes

For both rule shocks and environment shocks, in general the 12 example systems proved to be highly resilient, with most surviving all 100 attempts to hit them with a shock. Removal of 8 or more instances of a rule-type, both with and without compensating random additions, had some impact on the % of firm-rule-type combinations being lost between pre-shock and final system states, but the difference between losing 95% without shocks and losing 93% with them does not seem particularly dramatic. As figure 2 illustrates, even when, as in this example (applied to the system from figure 1), a rule shock has been strong enough to remove a firm, the impact on overall rule-type stock levels is hard to discern.



Figure 2 Numbers of instances of rules, labelled by each rule’s input product type. In this particular example run, at time = 50000 one firm lost all its instances of its only type of rule. However, any impact on the overall balance between rule stocks seems lost in the general variation over time.

Environment shocks involving the removal or addition of up to 400 instances of material were tested, but no correlation was found between shock size and loss of firms and types of rule. There was almost no loss of hypercycles during the experiments with environmental shocks. Figure 3 shows how the impact of an environment shock on the environment stock levels is short-lived as the system readjusts the stock levels.

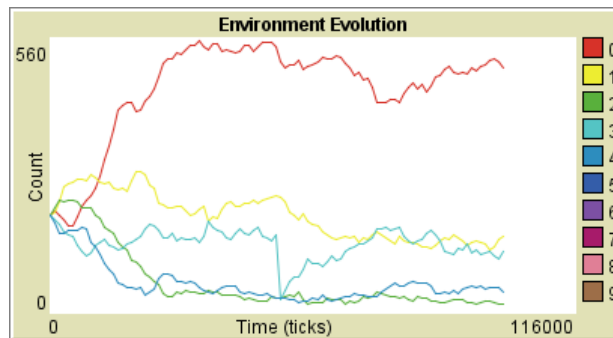


Figure 3 Impact of an environment shock on the stock level of materials in the environment. Although a large proportion of product type 3 have been removed at time = 50000, the system rebalances the stock levels quickly.




Figure 4 Link shock impact as shock involves increasing numbers of links: mean % hypercycle survival for the 12 test systems (left), and % of firm-rule-type combinations lost between pre-shock and final system state (right).

3.2 Sensitivity to firm network changes

Random rewiring of links did have a noticeable effect on hypercycle survival and on the numbers of firm-rule-type combinations. Figure 4 shows how hypercycle survival falls as more links are rewired, and also shows the impact on the numbers of firm-rule-type combinations. These results are averaged over the 12 test systems, which showed much variability between them.

4 DISCUSSION

4.1 Adaptation and redundancy

Why were the systems so resilient to rule and environmental shocks, despite the lack of intelligence on the part of the firms? Recall the analogy between these firms engaged in learning by doing and the stigmergic learning of ants altering their environments with pheromones. Just as ants can collectively find efficient routes to sources of food through a combination of *exploration* of new variations on paths and *exploitation* of previously travelled paths, so a system of firms can collectively seek and then track self-sustainability. The emergent hypercycle systems represent attractors in the state space of this simulation model. Sudden changes to rule and environmental stocks were not able to move the simulation from the basin of attraction of these regimes. So, after the shock the firms' and environment's stock levels re-converged on self-sustaining behaviour. The production systems retained their capacity for collective search or adaptive learning.

The production systems are also benefiting from a degree of redundancy. Most rule-types are instantiated at more than one firm. Most firms have instances of more than one rule-type. Although during rule shocks it is possible for a firm to lose all instances of a particular rule-type, provided that there are some instances of that rule-type still present elsewhere in the network of firms, the hypercycle capability can survive these shocks. Likewise, firms can survive the loss of instances of one rule-type, providing that they have instances of other rule-types. They suffer a change in role, but not a loss of all roles. So, although wiping out all of a firm's instances of one type of rule can alter the system's state space, it usually leaves behind a new attractor, still with at least one hypercycle in it, and the firms' and environment's stock levels soon converge on this.

Rewiring links, on the other hand, had a much more dramatic effect on the state space. Which firms are connected to which is not something that can be modified without loss. Although redundancy can help maintain the existence of hypercycles, firms are likely to experience alteration in their roles, and because of this, may lose one or both of their types of rule. The rewiring of a production system may be viewed as change in a search task. Kauffman has written about search task difficulty in the contexts of biological and technological evolution. Genetic algorithms, another heuristic search mechanism, can not only find optimal designs, or peaks on a moderately *rugged* fitness landscape, but can also track them as the fitness landscapes change, as can happen when co-evolving species are interlinked in a single eco-system [Kauffman 1996; 2000]. Depending on the amount of redundancy, or duplication of roles among the firms, rewiring the firm network may alter the number of attractors of self-sustainability and affect how easy it is to track these for the collective search capabilities of the remaining network of firms.

4.2 Future development

This paper has covered just a few types of shock applied to a relatively small set of example systems, generated using just one of a number of combinations of parameter settings. Clearly much more exploration must follow before we can start making more general points about what makes a production system resilient and by how much. For example, if firms were to change their behaviour during a simulation run, such as by modifying the learning-by-doing process, this might affect their collective ability to keep adapting and searching for self-sustaining system states. Future experiments should try more test systems, perhaps compensating for the extra computer work by reducing the number of simulation runs from 100 per system and parameter setting. Further investigation should also follow up the idea of an analogy between, on the one hand, SES resilience, and, on the other, search task evolution and difficulty.

In addition, much remains to be done towards bridging the gap between Padgett's abstract model based on artificial chemistry, and real-world social-ecological systems. Padgett & Powell [forthcoming] will relate the model to empirical cases of organisational emergence and evolution, and propose extensions of the basic model to allow dynamic networks between the firms or actors, and flux in the form of new firms entering the system. Once the chemistry has been extended to include multi-input, multi-output rules, we suggest it may also be possible to calibrate a hypercycles model using empirical data from material flow analysis (MFA). Given how Padgett's model has shown random variability to interact with network structure in determining hypercycle survival [Padgett 1997; Padgett et al. 2003], agent-based models of production systems may provide important new insights into the sustainability of systems of material flows, that system dynamics models neglect. Furthermore, this model might be used to explore the emergence, resilience and decay of industrial symbiosis or eco-industrial parks [Chertow 2000], in which firms are connected to each other through their by-products or wastes. If several firms produce several types of products or are flexible enough to use different inputs, this industrial symbiosis system is likely to be resilient to changes in production schemes of some firms and able to react to resource scarcity as long as other pathways for producing the necessary input or substitutive inputs can be found.

REFERENCES

- Chertow, M. R., Industrial symbiosis: literature and taxonomy, *Annual Review of Energy and Environment*, 25, 313-37, 2000.
- Corne, D., M. Dorigo and F. Glover eds., *New Ideas in Optimisation*, McGraw-Hill: London, 1999.
- Eigen, M. and P. Schuster, *The Hypercycle: A Principle of Natural Self-Organization*, Springer-Verlag: New York, 1979.
- Berkes, F., J. Colding and C. Folke (ed.), *Navigating social-ecological systems: building resilience for complexity and change*, Cambridge University Press: New York, 2003.
- Hofbauer, J. and K. Sigmund, *Dynamical Systems and the Theory of Evolution*, Cambridge University Press: Cambridge, U.K, 1988.
- Holling, C. S., Resilience and stability of ecological systems, *Annual Review of Ecology and Systematics*, 4, 1-23, 1973.
- Kauffman, S., *At home in the universe*, Penguin: London, 1996.
- Kauffman, S., *Investigations*, Oxford University Press, Oxford, 2000.
- Padgett, J. F., The emergence of simple ecologies of skill: a hypercycle approach to economic organization, in W. B. Arthur, S. N. Durlauf, D. A. Lane (eds.), *The economy as an evolving complex system II*, Addison-Wesley: Reading MA, 1997.
- Padgett, J. F., D. Lee and N. Collier, Economic Production as Chemistry, *Industrial and Corporate Change*, 12, 843-877, 2003.
- Padgett, J. F. and W. W. Powell, *The Emergence of Organizations and Markets*, forthcoming.

Watts, C. and N. Gilbert, *Simulating Innovation*, Edward Elgar Publishing:
Cheltenham, UK, forthcoming.