J Evol Econ (1993) 3: 89-107

Journal of
Evolutionary
Economics

Artificial worlds and economics, part I

David A. Lane

School of Statistics, University of Minnesota, 270 Vincent Hall, 206 Church Street S.E.,
Minneapolis, MN 55455-0488, USA

Abstract. In this and a companion paper (Lane 1993), I describe a class of models, called artificial worlds (AWs), that are designed to give insight into a process called emergent hierarchical organization (EHO). This paper introduces the ideas of EHO and AWs and discusses some of the interferential problems involved in trying to learn about EHO by constructing and studying the properties of AWs. It concludes by introducing two abstract AWs that address important general problems in EHO: the relation between structure and function, and the dynamics of evolutionary processes. The companion paper will discuss several AWs expressly designed to model particular economic phenomena.

Key words: Emergence – Organization – Evolutionary dynamics – Lambda calculus
JEL-classification: D 70

1. Emergent hierarchical organization

Many systems, in chemistry and biology as well as in human society, appear to have the capability of achieving, over time, a more and more complex organization. The process through which this organization is achieved, emergent hierarchical organization, typically displays two characteristic features.

First, the organization is hierarchical. That is, the systems are composed of a number of different levels, each level consisting of entities that interact with one another. Lower-level entities may actually be components of higher-level ones. The higher in the hierarchy is the level, the longer is the time-scale and the more extended the space-scale in which it is natural to describe the interactions between the relevant entities. For example,

- biological systems include entities and processes at levels ranging from molecular to cellular to organismic to ecologic;
- economic activities involve interactions between individual "decision-makers", firms and households, industries, and national economies.

Second, the systems appear to produce their own order. The actions of lower-level entities are channelled – in effect, coordinated – by higher-level structures that themselves arise from the lower-level entities' interactions. For example,

- informal trading networks transform into formally organized impersonal markets;
- neurons firing in response to sensory stimuli or the firing of other neurons with which they are connected produce predictable organism-level behavioral responses to particular patterns of environmental activity – or may even give rise to action-guiding "concepts".

The order induced by this kind of hierarchical coordination is never static, since the interactions between higher-level entities change the environment in which lower-level interactions take place, and hence in the higher-level structures that develop out of them. Thus, the system as a whole is characterized by perpetual novelty at all its levels.

2. What are artificial worlds – and what might we learn from them?

Artificial worlds (AWs) are computer-implementable stochastic models, which consist of a set of "microlevel entities" that interact with each other and an "environment" in prescribed ways. AWs are designed so that they themselves may, under some conditions, *manifest* EHO. As a result, AWs represent an engineering approach to the study of EHO.

The entities built into an AW and their modes of interaction may be quite abstract, or they may be closely linked to objects and relations occurring in some real-world system of interest. In the former case, the AW may be used to investigate general principles underlying EHO, while in the latter the AWs may help us to understand how particular aggregate properties of the modelled real-world system depend on the characteristics of the lower-level processes that underlie them.

Formally, an AW consists of a set of *microlevel entities* (MEs), an *environment* and a *dynamic*. Each ME has attributes and modes of interactions with other MEs. The environment has a state.

When two or more MEs interact, their attributes may change. The changes are determined by the MEs' interaction modes. In addition, they may depend on the MEs' current attributes and the current state of the environment. Interactions between MEs can also change the state of the environment.

The dynamic, which may be in part stochastic, specifies the order in which interactions occur. The dynamic also imposes rules that determine when MEs die and when new ones come into the World (and with what attributes).

The initial conditions of an AW determine a state of the World: the state of the environment, a population of MEs, and the attributes of each of the MEs. These initial conditions, together with the dynamic of the AW, generate a history – that is, a time-ordered sequence of states of the world. (With a stochastic dynamic, of course, the same initial conditions generate a probability distribution over a space of possible histories.)

The aim of AW modelling is to discover whether (and under what conditions) histories exhibit interesting *emergent properties*. An emergent property is a feature of a history that (i) can be described in terms of aggregate-level constructs, without reference to the attributes of specific MEs; (ii) persists for time periods much greater than the time scale appropriate for describing the underlying micro-interactions; and (iii) defies explanation by reduction to the superposition of "built in" micro-properties of the AW.[1]

For example, imagine an artificial economy in which MEs represent traders exchanging a set of commodities according to some prescribed rules that do not single out any particular commodity as a medium of exchange: the replacement of a barter system with the exclusive use of one of the commodities as a "money" would be an emergent property (Marimon et al. 1990). Similarly, in an artificial economy in which some MEs produce machines for sale to other MEs who in turn produce consumer goods for sale to other MEs (who work for one or the other producer MEs), the evolution of a stable growth rate for "GDP", or of sector-specific Pareto-distributions for firm size, might be an emergent property (Chiaromonte et al. 1992; Eliasson 1985; Lane 1993).

As these examples indicate, some emergent properties can be described in terms of variables that aggregate over the attributes of many MEs (like GDP), while others refer to "real" higher-level structures (like money). Both give evidence of self-organization in the AW – coordination among the MEs induced by their interactions, leading to system meta-stability. More is possible: higher-level "entities" may arise. These entities are composed of sets of MEs that display coordinated patterns of behavior. They may even reproduce themselves (section 3.2) and develop modes of interaction between one another (section 3.1), leading to even higher-level emergent properties. In such cases, the AW exemplifies EHO.

What can we hope to learn from AWs? We have to begin by considering "about what" we can learn. First, the AW itself might be the primary target of inference, and we might want to discover just which emergent properties it manifests, and how they depend on the system rules and initial conditions. Second, the AW might be regarded as a model of some real-world phenomenon in which we might be interested. In this case, we might want to determine whether (and if so, how) certain "lower"-level interactions in the real-world "cause" higher-level structures and processes to arise – and how these higher-level structures and processes then change the nature of the lower-level interactions. Third, we might want to learn about EHO as an abstract phenomenon, investigating such questions as the following:

- What properties must a system have for EHO to occur?[2]
- Is there a taxonomy of possible forms of emergent organization? In particular, are all emergent organizational forms hierarchical?

---

[1] Obviously, what "defies explanation" to one person may be explicable by another. What is required here is a negative assertion by the modeller, to the effect that the aggregate-level property in question is not deducible from the model's micro-properties by any argument substantially shorter than producing that property by running the model. I will discuss later some manoeuvers that might lend "public" credibility to such an assertion. Notice that the modeller's assertion is not equivalent to the statement that he assigns low a priori probability to the property manifesting itself when he runs the model: after all, he may have other reasons than deductive argument for believing that systems with the micro-properties he built in to his model tend to exhibit aggregate-level regularities analogous to the property in question!

[2] See Kauffman (1990) and Rasmussen et al. (1990) for some interesting speculation on this question.

- How do the properties of emergent higher-level entities and their interactions depend on the properties of the lower-level entities from which they arise?
- What kinds of interactions are possible between the levels of a hierarchically organized system? In particular, how autonomous are the processes of different levels? Under what circumstances can the evolution of a system process be predicted on the basis of observations only of the attributes of entities at the same level as the process (that is, without detailed information about processes at lower or higher levels)?
- What are the dynamical properties of emergent processes: For example, are "punctuated equilibria" (section 3.2) generic?

While computer scientists might be interested in an AW for its own sake, economists presumably would study AWs in order to get insights into what might be going on in economies. Whatever the goal, to learn anything useful about any of the three inferential targets described above, we need strategies for designing appropriate AWs and for generating and processing useful data from them. There are some formidable difficulties standing in the way of this endeavor. I conclude this section by mentioning four of them.

The need for computer-implementation

AWs are well-defined mathematical models, but it is unlikely that interesting theorems about their emergent properties will be proved with tools currently available. I offer three reasons for this assessment:

First, AWs are designed to be innovatory or open-ended systems. Their emergent properties are only meta-stable, not equilibria or asymptotic states. By changing the environment of the lower-level entities that give rise to them, emergent structures induce processes leading to their own transformation (or demise). As a result, it will be difficult to apply the rich repertoire of mathematical methods that compute equilibria or asymptotic states, and there is no corresponding methodology for studying the properties of transient phenomena.

Second, emergent properties are necessarily complicated functions of the history of the attributes of the ME's from whose interactions they are formed (if this were not so, it would be easy to explain them by superposing the AM's microproperties, and they would not qualify as emergent properties!). Since the dynamics of AWs are specified in terms of these micro-interactions, it is hard to imagine that the mathematical description of emergent properties will be analytically tractable.

Third, it seems to be a plausible (albeit ill-defined) hypothesis that the capability of a system to produce EHO is a function of its complexity, either in the attributes or arrangements of its component entities or in their patterns of interaction. As a result, the mathematician's ploy of constructing a higly simplified, tractable model that can be proved to display an interesting behavior observed in some more complicated system will not work in the context of EHO phenomena.

Thus, it seems likely that we will learn about EHO from AWs only by implementing them computationally and observing what happens. As a result, we can learn about their emergent properties only inductively, and our success in that enterprise will depend on our ability to develop appropriate statistical tools, for the design as well as for the analysis of "evolutionary" experiments.

Identifying emergent properties

The very nature of emergent properties makes it problematic for us, as observers of the AW, even to formulate them, let alone discover whether or not they in fact obtain. Emergent properties represent innovations in the organization of the AW, and, to describe them, a new vocabulary is required, beyond the modelling language used to express the attributes and interactions of the AW's micro-entities. After all, emergent properties cannot be compactly expressed in the modelling language itself – and, by definition, they "defy explanation" in terms of the constructs of that language. So how do we develop the right aggregate-level language to define – and guide our search for – potentially emergent properties?

AWs that model a real-world system have a natural vocabulary to express potentially emergent properties: the language that describes higher-level patterns and structures observed in the modelled system. Some[3] of these higher-level constructs may suggest AW analogs that can be expressed as functions of AW histories, and the words that describe the real-world constructs may be appropriated to define these functions. Thus, the modeller can build a glossary that semantically links higher-level real-world constructs with particular functions of AW histories. Any real-world phenomenon that can be described by these constructs translates, via the glossary, to a candidate for an emergent property of the AW – provided, that is, that it satisfies the metastability and "explanation-defying" definitional requirements. Candidates generated in this way might be described as "expected emergent properties" of the AW.

"Unexpected emergence" – an aggregate-level coordination phenomenon in the AW unmotivated by any real-world analogy – is harder to find. This is particuarly troublesome for abstract AWs, which lack a natural real-world reference vocabulary. In fact, most of the work that goes into studying such AW models as Coreworld (Rasmussen et al. 1990), Tierra (Ray 1992) and Function-Object Gas (Fontana 1992 – see section 3.1 below) consists in poring over output, attempting to identify features that display the "right" kind of coherence and temporal stability – and then formulating a vocabulary, with both mathematical and "natural language" variants, in which to express them. Whether this search can be in some way "automated" is an important conceptual and practical problem.[4]

---

[3] But certainly not all. After all, the modeller abstracts only a small subset of entities, attributes and interactions to incorporate into the artificial world, and only those higher-level constructs for which it is meaningful to aggregate only over this subset can be translated as a function on artificial world histories. The determination of which higher-level constructs are meaningful in the artificial world – and how – can be an important exercise for understanding the meaning and role of these constructs in the real world system itself.

[4] Bedau and Packard (1992) propose a statistic whose purpose is to diagnose the arrival of an "innovation" into an artificial world. Their statistic seems to depend on a genotype-phenotype distinction: the microentities in the world are replicators, whose behaviors are coded by a genome; selection operates on the coded behaviors; innovations in behavior depend on the introduction of a new genotype; and successful innovations are marked by the initiating genotype's ability to persist in the population over time. The Bedau and Packard statistic tracks such persistence at the genomic level. But the generality of this approach seems questionable: not all higher-lever innovations depend upon the persistence of single micro-innovations, even in biological evolution. To paraphrase the evolutionary perspective persuasively set forth in Buss (1987): on an evolutionary time scale, genotypes are transient, while phenotypic organization is here to stay.

Finding conditions of emergence

When potentially emergent properties have been identified and translated into the behaviors of appropriate functions on histories, the next question to ask is: under what initial conditions (and, for stochastic dynamics, with what probability) will they obtain? Developing strategies to answer this question is difficult, since the space of initial conditions typically has a very high dimension, and interesting emergent properties may well depend on complicated interdependencies among the system parameters that define these dimensions.

Moreover, the relevant search space is even larger, because it has a time dimension. Well-defining the function on histories that determines whether a particular property emerges requires a specification of how long that property must persist – and this specification must always be somewhat arbitrary. In addition, whether a particular property emerges or not depends not only on initial conditions, but on the length of time the history is observed – so negative results may just mean that longer observation times are required, not that the initial conditions are insufficient to support the emergent property in question.

Causality and emergence

Suppose a potentially emergent property of an AW has been identified and defined in terms of some function of histories – and, with some set of initial conditions, a history has been generated and the property obtained. What kind of claim can be made about what "caused" this property – in particular, is it meaningful to think of emergence itself *as* a cause?

To interpret emergence as a cause, we mean to say that the property formed because of the interactions amongst a dense network of entities – and this formation depended on the density of this network, and perhaps the richness of the structure of the entities and their interactions. Thus, it is not enough merely to produce the property in the AW from some particular set of initial conditions: that set would have to be embedded in a hierarchy of sets, ordered by a "complexity" measure that increased with the network's density and the structural richness of the MEs and their interactions. Emergence as a cause would then require demonstration that the property fails to appear for low values of this measure – but does, beyond some threshold value.[5]

Such a complexity measure imposes a structure on the high-dimensional AW parameter space. Without this structure, it is hard to see how one could begin to infer about what causes emergent properties – and it is equally hard to see how any causal inference could be made that is independent of the particular measure used to induce the structure.

Now suppose we know how to infer about emergence-as-cause inside the AW. Suppose further that we believe that a particular aggregate-level feature in the AW is indeed an emergent property, and we have determined how "complex" the AW

---

[5] One might suspect that typically, as the complexity measure increases above this value, a second threshold might be obtained, beyond which the system again fails to manifest the property in question – just as turning up the heat applied to the bottom of a beaker of fluid results first in the formation of convection cells and, at even higher temperature, their degradation into a regime of turbulence. See Kauffman and Johnson (1992) and Langton (1992) for stimulating discussions on this theme.

needs to be in order to support the feature's emergence. Suppose in addition that this emergent property is semantically linked to some real-world higher-level pattern or structure: what can we infer about the "cause" of this feature in the real world?

At the least, we can certainly argue against the necessity of any alternative explanation that assigns a causal role either to other real-world aggregate-level features that do not have analogs in the AW or to attributes of lower-level "agents" that are not possessed by the MEs of the AW. For example, an artificial economy in which, say, a stable growth path for GDP emerged from sufficiently rich patterns of micro-interactions would thus argue against the necessity of invoking the existence of Walrasian equilibrium to explain macro-coordination – or against the proposition that such macro-coordination depended upon the assumption of optimizing agents capable of forming rational expectations.

But we would like to infer more than this. Can we argue that the real-world aggregate regularity is indeed "caused by" the entities and interactions we abstracted out of it and built into the AW, in which the analog of that regularity was identified as an emergent property? That is, can we infer emergence as a "causal mechanism" in the real world, once we have so identified it in the AW?

Certainly, the AW demonstration ought to raise our probability that such a mechanism operates in the real world, just as it diminishes the probability of alternative causal stories that credit features and attributes not detected or built into the AW. But the real world necessarily contains many more entities and interactions than the AW, operating at levels below, at and above that of the focal regularity. Surely, it is possible that the causal mechanism hinted at in the AW is swamped by the additional "turbulence" in the real world, and some entirely different sets of interactions or direct effects drive the formation of the features of interest. It is not clear how to determine how plausible is this possibility – but of course, the more specific one can be about just which additional interactions or effects might provide the alternative causal story, the more plausible it would appear to be.

3. Abstract AWs and the lawfulness of EHO

In this section, I describe two abstract artificial worlds, Walter Fontana's Function-Object Gas (Fontana, 1992) and Kristian Lindgren's Evolutionary Prisoner's Dilemma (Lindgren 1992). Function-Object Gas is directed primarily to an exploration of the relation between structure and function, Evolutionary Prisoner's Dilemma to the dynamics of evolutionary processes.

While much work remains to be done before AWs yield deep insight into these two themes, the themes themselves are fundamental to an understanding of many real-world processes. The section concludes with a discussion of an economic example of such a process, the coming into being of a new industry.

3.1 Function-Object Gas: function and organization

Function-Object Gas (FOG) is designed to explore how higher-level structure emerges from micro-level function. The notion of function on which FOG is based is abstracted from chemistry. A chemical entity functions by acting on other chemical entities to produce new chemical entities. Similarly, in FOG, all interactions between MEs are of a single type: a ME A acts on a ME B to produce a new ME A(B).[6]

FOG also abstracts from chemistry the relation between structure and function at the micro-level. Which new entities are produced when chemical entities interact are completely determined by the structure of the interacting entities: the components from which they are built up and the way in which these components are arranged. Thus, a chemical entity is both a syntactic and a semantic object. Syntactically, it is built up from component objects, according to well-defined rules. Semantically, its "meaning" (that is, its function), coded by its structure, is revealed in the chemical reactions in which it partakes. The dual character – syntactic and semantic – of chemical entities is most striking in catalysis: the syntactic form of the catalyst is unchanged, even as it accomplishes its function of transforming the structure of other chemical entities.

In FOG, each ME has a syntactic representation in terms of more elementary components. This representation never changes during the lifetime of the ME. An ME's representation codes for its semantics, in that the representations of the interacting MEs determine the outcome of the interaction. That is, the representations of the MEs A(B) and B(A) can be "computed" from the representations of A and B, for every pair of allowable syntactic representations A and B.[7, 8] In FOG, all interactions are doubly catalytic: neither A nor B is "destroyed" by their interaction. So, $A+B \rightarrow A+B+A(B)$.

Thus, in chemistry and FOG alike, micro-level function is determined by micro-level structure. However, this is by no means the end of the function-structure story: micro-level function can in turn give rise to *higher*-level structures. Consider an autocatalytic network: a set of chemical entities that (perhaps in the presence of some "food set") catalyze reactions among its members (and the food set), such that each member of the network is a product of at least one of these reactions. Thus, an autocatalytic network reproduces itself – collectively, not necessarily individually. Take away some of its members, and an autocatalytic network may "disappear" as one after another of its members fail to be produced by reactions involving remaining members; while the removal of others of its members may not matter, as they are soon replaced from transformations among the "survivors". Thus, even though the functionality of a particular chemical entity may be latent in its structure, the organizations of chemical entities to

---

[6] The interacting entities are ordered: A(B) need not be the same as B(A). In addition, A(B) is not defined for all MEs A and B.

[7] Technically, this is achieved by using Alonzo Church's λ-calculus to represent MEs as λ-objects – mathematical functions in intensional form, that act on other functions to yield new functions according to nine axioms of construction and syntactic transformation. Computationally, then, a λ-object is both function and data. The components of a λ-object are variable names, the abstraction symbol λ, and three structural symbols (period and left and right parentheses). The set of λ-objects are defined recursively by the three construction axioms: variables are λ-objects; if x is a variable and M an λ-object, then $\lambda x.M$ is a λ-object; and if M and N are λ-objects, so is M(N). The semantics governing function evaluation are incorporated in the other five axioms. The λ-calculus is computationally complete; every recursive function can be represented as a λ-object. See Barendregt (1984) for details.

[8] For λ-objects A and B, B is not in the domain of A if the computation implied by the transformation axioms applied to A(B) does not halt. In FOG, there is a limit placed on transformation steps, and any interaction whose associated computation exceeds this limit produces no product.

which this functionality may give rise are really aggregate-level or population concepts.

To see how FOG can be used to address the problem of the emergence of higher-level structure from micro-level function, I first describe how to generate a FOG history. Start with a population of MEs[9] (λ-objects: see footnote 7) – these are typically generated at random. Next, select a pair of these MEs at random, say A and B, and let them interact as described above. If the computation for A(B) terminates, add this ME to the population and select another ME at random and remove it from the population. This dynamic keeps the population size constant. Now iterate the interaction-deletion steps many times.[10]

The population of MEs in the FOG after many interactions may display structure at the syntactic or the semantic level. Syntactic structure refers to common features of the representations of the members of a set of MEs. For example, the set of λ-objects of the form $A_{ij} = \lambda x_1.\lambda x_2 . \ldots . \lambda x_i .\, x_j, j<i$, exhibits syntactic structure.

Semantic structure depends on the production pathways involving reaction products from interactions between members of the set. For example, suppose A, B, C, and D are MEs, with A(B) = C, B(C) = D, C(D) = A and D(A) = B. Then, regardless of the other interactions of these MEs, the set (A, B, C, D) is *self-maintaining*, in that each can be formed from interactions between members of the set. (This property is analogous to the concept of an autocatalytic network). Note also that (A, B), (B, C), (C, D) and (A, D) are all *seeding sets*, in that the entire set can be reconstructed by interactions involving the elements in each of these subsets and their "descendant" products. A set that contains all of the products from interactions between set members is *closed*. Closed self-maintaining sets are *self-reproducing*.

Self-maintaining sets are not guaranteed to survive under FOG dynamics, since MEs are removed randomly from the population. Clearly, MEs that belong to a self-reproducing subset with several small seeding sets have a better chance of persisting in a population that contains that seeding set than does an ME that belongs to no such subset. One way in which a FOG population can display semantic structure is if it can be decomposed into a number of such self-reproducing subsets. These subsets in turn can have a variety of semantic structures, which may be represented by means of interaction graphs, as in Fontana (1992).

So far, there have been no constraints imposed on interactions in FOG, except for the upper bound on allowable computation time (see footnote 8). It turns out, however, that what higher-level structures form depends crucially on which interactions are allowed to take place. For example, some MEs may reproduce themselves (that is, A(A) = A) or other MEs (A(B) = B). Clearly, if the set of MEs reproduced by an ME A contains A, it is self-maintaining, in a trivial way. Fontana (1992) reports that, without constraints in interactions, FOG tends to organize around production pathways that end in an ME that reproduces every ME in the pathway. Starting with 1000 random MEs, after tens of thousands of collisions, the FOG population is typically closed and consists of one or more self-reproducing subsets, each with its own identity function.

---

[9] There is no (external) environment in FOG.

[10] Note that with this dynamic, FOG interactions are "on average" singly, not doubly, catalytic, since A is removed from the system with the same probability as it is selected to form a product A(C), for all C in the population.

Thus, to explore a greater range of interesting emergent structures in FOG, Fontana has begun to investigate what happens when he constrains the permissible set of interactions. He does this in two ways, which correspond to syntactic and semantic constraints. For example, barring copy reactions is a semantic constraint, since whether a reaction copies one of the reactants is a function of the interaction, not just the product of the reaction. In general, though, it is difficult to formulate semantic constraints. Syntactic constraints bar interactions that produce reaction products with specified structure. Thus, they amount to restricting the FOG population to particular subsets of λ-objects.

To determine which products to prohibit, Fontana has taken advantage of a peculiar finding: FOG tends to produce organization on *both* the syntactic and semantic level. That is, when the FOG achieves a metastable, closed population, this population exhibits patterns both in the structure of their MEs and in their production pathways. Thus, it is possible to prevent a particular semantic organization from occurring by prohibiting reaction products that have its corresponding syntactic features.

For example, when copy reactions are prohibited, families consisting of MEs of the form $A_{ij} = \lambda x_1.\lambda x_2. \ldots \lambda x_i.x_j, j < i$, as described above, proliferate. Their syntactic structure is clear. Semantically, according to the transformation rules of λ-calculus, these so-called projection functions satisfy

$$\begin{aligned} A_{ij}(A_{km}) &= A_{i-1,j-1}\,, && \text{if} \quad j > 1 \\ &= A_{k+i-1,m+i-1}\,, && \text{if} \quad j = 1 \end{aligned}$$

Thus, start with, say, A_{i1}: this ME acts on itself to produce $A_{2i-1,i}$, which then acts on itself (or on any other member of the family) to produce (in turn) $A_{2i-s,i-s}$, for $s = 2, \ldots, i\text{-}1$. These i MEs form a simple semantic structure, organized around the cycle $A_{i1} \rightarrow A_{2i-1,i} \rightarrow A_{2i-2,i-1} \ldots \rightarrow A_{i1}$. Note that any member of this cycle is a seeding set for the cycle. According to Fontana (personal communication 1992), FOG without copy reactions organizes into one or more of these families, with transient random selection between families (and victory tends to go to the largest).

So the next organizational question to investigate is: what structures emerge when all MEs of the form A_{ij} are prohibited? Once these are discovered and their syntactic regularities are found, a further constraint can be imposed, and additional organization forms obtained. By continuing in this way, Fontana is uncovering a hierarchy of increasingly complex organizational forms that can emerge in FOG, under increasingly complex constraints on allowable interactions. He is attempting to associate with each of these forms an underlying algebraic structure that describes its interaction graph. The hope is that these structures will provide the basis for a mathematical theory of organizational form.

Another direction of current research with FOG is to search for the emergence of structures at a higher level than the sets of MEs so far described. For example, can self-reproducing sets interact with one another to produce other sets with some metastable structure? An interaction between sets of MEs can be defined trivially to produce the union of all the pairwise interactions between elements of the two sets. It is not clear that this is a useful definition; nor is it yet clear what a reasonable alternative might be. It may also be necessary to introduce noise into the system, for example by occasionally perturbing the structure of individual MEs or the products of their interactions. This may "destabilize" emergent organizations, especially those that involve many MEs with complicated production pathways, with the results that the system will support more, smaller structures that may support or inhibit one another through their mutual interactions. At any rate, EHO is so far a one-level phenomenon in FOG.

To conclude this discussion of FOG, consider an alternative way of building a computational system in which entities interact with entities to produce new entities. An obvious strategy is to decide how many entities you want to have in the system, say n, and then randomly construct an n-by-n lookup table that gives the products of all possible pairwise interactions. Representing MEs as λ-objects has two principal advantages over this "random lookup" strategy:

- The λ-based system is computationally open-ended.[11] You are not limited to any pre-fixed number of MEs, and you can represent any imaginable relation between MEs, since any computable function can be expressed as a λ-object.
- In the λ-based system, the representation of MEs codes for their function. Thus, it is possible to explore relations between structure and function that have no counterparts in the "random lookup" scheme. In particular, any syntactically correct expression or family of expressions can be inserted (or deleted) from the system and the effects on organization monitored. Put another way, the λ-representation provides a true genotype-phenotype[12] distinction – and a way of experimentally determining which "genes" are responsible for which "body plan" characteristics.

On the other hand, experiments with FOG alone cannot tell us whether the structure-function relations that they reveal depend upon the λ-representation of its MEs. That is, we need other arguments to determine whether the algebraic structures of organization that Fontana is discovering are general principles of emergent organization or merely artifacts of his model (and perhaps reducible to theorems in λ-calculus itself). These arguments must be inductive in character. Can these structures (and not others!) be observed in other systems, from real or artificial worlds, in which functional interaction can be interpreted as the creation of new entities?[13]

3.2 Evolutionary Prisoner's Dilemma: the dynamics of evolutionary processes

Evolutionary Prisoner's Dilemma (EPD) is a simple example of an evolutionary process. The leading natural example of an evolutionary process is, of course, biological, and it is far from simple. It is hard to think about biological evolution now without taking account of its rich organizational structure, in particular the hierarchy of descent (replicating genes, interacting organisms, evolving species –

---

[11] At least in principle; in practice, one must introduce constraints on the number of steps in a computation, the length of the representation of objects and so forth.

[12] Here a self-maintaining set of λ-objects represents the "organism", with the syntactic structure of each λ-object representing a gene. The phenotype is the (semantic) structure of the interaction graph of the set and its reaction products.

[13] Fontana and biologist Leo Buss are currently translating some organizational experiments with FOG into the language of evolutionary biology, with promising results (Fontana and Buss 1993). In particular, they provide new interpretations of the significance of "life cycles".

and beyond) and the economic or ecological network, with its complex of relations between organisms, revolving around energy production and exchange[14].

The concept of evolutionary process on which EPD is based abstracts away from all this structure. It starts with the notion of an entity as a set of attributes. Entities are capable of self-replication: that is, they can produce other entities that have the same set of attributes as themselves. Entities with the same set of attributes form an entity type. The entities in an evolutionary process form a population, and the population consists of more than one entity type. Different entities replicate at different rates, so that the distribution of entity types in the population changes over time.[15] The probability that an entity replicates at any given time depends not only on its own attributes but also on those of the other members of the population at that time. Finally, evolutionary processes include mechanisms whereby entities with new kinds of attributes enter the population. Frequently, these mechanisms depend upon innovation-generating errors that take place in the process of replication.

Thus, evolutionary processes are characterized by *replication* (the reproduction of existing entities), *selection* (the differential replication rates of different entity types), and *variation* (the generation of new entity types). To determine a particular evolutionary process, it is necessary to specify the following elements:

- a *set of entity attributes*;[16]
- a *fitness function* (which may be stochastic) that gives the replication rate for each entity type, given the current distribution of entity types in the population;[17]
- *variation mechanisms* whereby new entity types enter the population; and
- an *initial population* of entities.

For example, in population genetics models used in theoretical evolutionary biology, entity attributes are typically defined at the genotypic level. The variation mechanisms include such genetic operators as mutation and recombination. The most problematic element in these models is the fitness function, since relative replication rates depend on the interactions at the phenotypic level. Thus, a genotype's relative replication rate is a function not only of how phenotype is determined by genotype,[18] but also of the kinds of ecologic relations that different phenotypes have with one another (competition, predation, symbiosis and so

---

[14] For introductions to the literature on hierarchical views of evolution, see Hull (1988, 1989), Salthe (1985), and Eldredge (1985).

[15] Entities may also leave the population, for example by dying.

[16] Note that if the process is truly open-ended, S is an infinite set.

[17] Note that the domain of the fitness function is not the set of individual entity types, but the set of possible *populations* of entity types. The process described here is coevolutionary: the fitness of each entity type depends on what other entity types share its world. In this sense, the population is an "individual", with entities as its "parts", which itself undergoes evolution. Thus, no "landscape theory" that fixes a "fitness function" over the set of entity types can describe the dynamics of the kind of evolutionary process I am defining here, since such a "landscape" is continuously deforming as the distributions of the entity types in the populations change.

[18] Which may of course depend in part on what other genotypes are in the relevant population, since this determination is "environmentally" mediated – and the other entities in the population form part of a given entity's environment.

forth). These underlying processes are not at all well understood, and so it is impossible to derive the form of the fitness function from first principles. In contrast, if entities were taken to be organisms (or even species), the relevant attributes might be structural or functional properties that could be directly related to relative replication rates – but then the variation mechanisms could be modelled only phenomenologically.[19]

The designer of an AW evolutionary process faces two difficult challenges: how to determine the fitness function for an arbitrary population of MEs, and how to create variation mechanisms that can supply new types of MEs indefinitely. Lindgren solved these problems, and also provided a natural language in which to describe his AW, by building EPD aroud a version of Iterated Prisoner's Dilemma.

Each EPD ME represents a strategy for playing a two-person game with two possible actions (say, 0 and 1).[20] This strategy is the only attribute of the ME. Each generation, MEs interact with one another in a round robin tournament; every ME in the population uses its strategy to play a particular version of Iterated Prisoner's Dilemma against every other ME.[21] The MEs then receive their average reward from these encounters, and they replicate in such a way that the expected number of replicates of each ME is proportional to its average reward.[22] Thus, the fitness function is determined by the representation of MEs, via the pairwise interaction rule of the round robin tournament and the interpretation of the representation as a Prisoner's Dilemma Strategy.

---

[19] An alternative approach to modelling evolutionary processes begins by positing two different types of entities: replicators and interactors. Replicators have a fixed structure that can be exactly replicated; variation mechanisms then introduce new types of replicators. On the other hand, replicators do not interact directly with one another; interactors do. So selection operates on interactors. The key modelling problem in this approach is to relate the replicators to the interactors: in particular, how do the functional properties of interactors depend upon the structure of replicators, and how do the interactions between interactors determine the differential rates at which the replicators replicate? The answers to these questions determine the analog of the fitness function described in the text. Hull (1988, 1989) argues exhaustively and convincingly for this approach to modelling biological evolution. In EPD, the MEs (or strategies, see text) are both replicators and interactors.

[20] Each EPD ME is a string of 0's and 1's of length 2^m, where m is an integer. The strategy encoding for the ME works as follows: write the last m moves (in reverse order: the opponent's last move, your last move, the opponent's next-to-last move, . . .); read what you have just written as a binary number; go to that coordinate of the your strategy vector – and play the number you find there.

[21] The version has the following features: (a) the play is noisy: that is, if a player's strategy dictates that he play a "0", say, he plays a "1" with probability p (p is small, and does not depend on the player or the history of the game); (b) the payoff per play is as follows: if both players choose 0 ("defect"), they each win 1; if they both choose 1 ("cooperate"), they win 3; otherwise, the one who chooses 0 wins 5 and the one who chooses 1 wins nothing; (c) the iteration is infinite, and the reward to each player in the iterated game is average payoff per play given above.

[22] In Lindgren's version of EDP, population size is kept constant and the proportion of each entity type in the next generation is proportional to its average reward. If the proportion of any entity type falls below 1/N, where N is the nominal population size, the entity type is dropped from the population. In effect, rather than setting the probability of replication for each ME to be proportional to its average reward, Lindgren substitutes the expected number of replicates per type. While Lindgren's version gains computational efficiency at the cost of failing to be a true evolutionary process, it shares the qualitive dynamical features described below with the truly evolutionary probabilistic replication scheme.

Variation in EPD arises from three kinds of replication error, each of which occurs with a fixed probability, independently for each transcription event. First, any given bit may be transcribed incorrectly (here the probability is per bit transcription, so the greater is the length of the string representing the ME, the higher the probability of replication error). Second, the string may get adjoined to a copy of itself, doubling its length (for example, "01" is incorrectly copied as "0101"). This error is particularly important, since it makes the set of possible MEs infinite, so that EPD is potentially open-ended. Because of the way in which strategies are encoded (see footnote 20), the offspring ME resulting from this error has exactly the same strategic behavior as its parent.[23] However, its doubled length means that it takes account of one more previous move than its parent does – and a subsequent transcription error in any of its bits will give rise to a different kind of strategic behavior than could arise from any transcription error in the parent type. Finally, the string may be cut in half, with either half chosen at random as the viable offspring (for example, "1101" might be incorrectly copied as either "11" or "01").

EPD dynamics exhibit interesting emergent properties. First, a succession of stable ecologies – that is, distributions of entity types that persist for many generations – form, dominate the EPD population, and then degrade. Both the individual ecologies and their succession may be regarded as emergent higher-level structures. Each ecology may possess one of a number of possible organizational forms: some are dominated by a single entity type; some have several symbiotic or competitive dominant types; in others, the dominant role is distributed among a number of "quasi-species" that share some key features and differ in others.

Second, the periods of stasis or "quasi-equilibrium" in which a stable ecology persists are interrupted by shorter periods of destabilization, which also display certain characteristic features. During a destabilization period, the number of entity types in the population fluctuates rapidly. Frequently, these periods begin with a large "extinction", in which the number of entity types drops rapidly. It is also typical that the average reward that MEs receive drops during the destabilization periods. In EPD, there is no exogeneous environment, so that all destabilizations are endogeneously generated: that is, such phenomena as mass extinction and structural disintegration do not necessarily require exogeneous causes (like asteroid collisions or volcanic eruptions!). Destabilization periods end with the formation of a new stable ecology, in which the leading entity types were not present (or present only at low frequencies) in the previous "quasi-equilibrium".

Contingency plays an important role in EPD ecological succession. While it is easy to compute which strategies have relative advantages over which, it is not easy to predict which sets of strategies will dominate the emerging stable ecologies. Start with the same initial populations, and quite different successions can occur. For example, starting with particular values for the system parameters (growth and error rates) and an initial population consisting entirely of memory 1 strategies, with probability[24] about 0.9 EPD will end up (by 30,000 generations) in an ecology dominated by many different memory 4 entity types that share common features in their representation (1××10×××0××××001): Lindgren

---

[23] For example, 0101 is the same strategy as 01, since its play depends only on the opponent's last move, regardless of its own previous move.

[24] These probabilities, as reported in Lindgren (1992), are of course obtained as frequencies over many runs of EPD.

argues that this particular ecology cannot be destabilized by the low-frequency introduction of any possible entity type. On the other hand, with probability 0.1, this ecology will not form, and the system will follow some other succession, leading to ecologies whose dominant types have memory lengths of 5 or greater.

These features of EPD dynamics – a contingent succession of "quasi-equilibria" interrupted by "catastrophic" destabilization periods – resemble the "punctuated equilibrium" version of the history of biological evolution, as put forward by Eldredge and Gould (1972).[25] Their appearance in such a simple evolutionary process as EPD suggests that they may be generic, at least in some very general subclass of evolutionary processes. An important goal for future work with abstract AWs is to try to discover the defining properties of this subclass and to gain a better understanding of punctuated equilibrium dynamics. What characterizes the set of possible stable ecologies? How large is the set? To which perturbations is a stable ecology robust – and which destabilize it? Why are the destabilization periods relatively short-lived, compared to the "quasi-equilibria"? Why are destabilization periods frequently initiated by rapid mass extinctions – and what endogenous mechanisms drive these events? What determines the order of succession of stable ecologies – and which successions are contingent and which (at least conditionally on some predecessors) necessary?

I conclude this discussion by pointing out two important phenomena in biological evolution that do not arise in EPD but could be the targets of future AW research. To explore these two phenomena would require evolutionary AWs with more structural possibilities for higher-level organization than are present in EPD:[26]

- A key ingredient of the "punctuated equilibrium" story is that fundamental structural innovation seems to arise only in brief destabilization periods, not in the intervening "quasi-equilibria", in which various "implications" of the fundamental innovations are worked out. Most dramatically, all existing animal phyla (and many more, since lost) appeared in the Cambrian explosion, a period lasting less than two million years, over 500 million years ago (Gould, 1989). That is, biological evolution seems to produce big differences first, in quick bursts, and slowly fills in the details.

- In biological evolution, selection operates at more than one level at the same time. Thus, within organisms, cellular selection continues to occur (for example, cancers represent successful selection at the cellular level that can be fatal at the organism level); and, at the same time, higher level entities – like colonies, species, or even ecologies – complete for resources, reproduce themselves and generate new attributes that lead to new colonies, species, or ecologies. The coexistence of all these processes constrains the structure and direction of each of them.[27]

---

[25] Somit and Peterson (1992) contains a very interesting series of essays on the meaning and scope of punctuated equilibrium.

[26] An interesting evolutionary AW that addresses at least the first of these issues is Thomas Ray's Tierra (see Ray 1992).

[27] According to Leo Buss (1987), the two phenomena are related: the bursts of structural innovation coincide with the emergence of a new level of entity, which has successfully developed mechanisms that control the selection processes operating on its component entities so that they do not favor variants that are harmful to the larger entity of which they are a part.

3.3 Economics, EHO and abstract AWs

In general, abstract AWs are designed to study processes whereby higher-level structure emerges from lower-level functional interactions. The two abstract AWs described in this paper, FOG and EPD, focus on two different aspects of these processes: the characterization of types of structure that can arise as a function of constraints on allowable interactions; and the dynamics of emergent structure. Clearly, far more exploration of both of these themes, by these and other abstract AWs, must be carried out before we can expect to gain useful insights into the lawfulness of EHO processes. Once obtained, such insights will serve as a background against which it might be possible to understand what is generic and what particular to real-world processes in which these themes appear to play a role.

Here I offer an economic example of such a real-world process: the coming into being of a new industry.[28] This process is central to economic growth and development. The point is not that we can apply Fontana's and Lindgren's investigations to learn anything interesting about this process. Rather, I want to call attention to those of its features that appear to exhibit EHO and to argue that these features are fundamental to understanding what the industry comes to "be" and to "do". Furthermore, the most interesting questions that arise about the process in my description involve precisely the themes that FOG and EPD were designed to investigate.

The emergence of industrial structure. I begin by sketching what I mean by the structure of an industry. An industry can be described in two complementary ways. First, the industry can be identified with the set of products that it produces. These products are related to each other functionally, by the uses to which they can be put, and technologically, through the processes by which they are made. These two kinds of relations induce a structure to the industry's product set.

An industry's product set changes over time, as new products and ways to make them are developed. Since new products may come from the modification of existing ones (or their production processes), products also are related to one another by descent. Descent relations induce a hierarchical structure on the product set, with higher-level "taxa" defined in terms of successively more remote "common ancestors". As is the case in biology, the members of higher-level families of products also may share attributes, for example, functional complementarities (such as computers that share software) or similar production processes (so that expertise accrued in making one of the family carries over to making others).

The second way of describing an industry is as a collection of economic entities or "agents". These entities have a variety of structural relations with one another, all oriented towards developing, making and exchanging products in the set described above. At least six classes of entities enter into these relations: producers, demanders, suppliers, financiers, scientists, and governments. While the industry has an organization induced by the relations between its component entities, these entities themselves (firms, universities, research centers, regulatory agencies) have internal structure as well. Thus, an industry exhibits hierarchical structure. For example, a firm may have, subordinate divisions – marketing,

---

[28] The formulation of this process, sketched here, is described in detail in a forthcoming paper by the author, Franco Malerba and Luigi Orsenigo.

production, R&D – and may also belong to a superordinate entity like a research consortium or a trade association.

The entities that make up an industry and the kinds of relations between them also change over time, as a result of the interactions between the entities. Thus, the industry's organization is an emergent phenomenon. Consider, for example, the case of biotechnology.[29] By 1975, research funded by NIH and NSF and carried out by scientists working in the biomedical centers of several American universities had resulted in the development of recombinant DNA and hybridoma technologies. With financing obtained initially from venture capitalists (a relatively new kind of financial entity, swollen with profits from prior investments in microelectronics), some of these scientists set up new firms designed to exploit the economic possibilities of the new technologies. There were some formidable obstacles to be overcome, especially in product selection and development and "scaling-up" production volume.

Lured both by the promise of the technologies and their potential competitive threats to existing products and production methods, some older, established firms explored a variety of techniques to acquire proficiency in the new technologies – ranging from research contracts with individual scientists and their universities or with the new biotech firms, to buying into the new firms, to setting up in-house biotech R&D units. The most active of these established firms were pharmaceutical companies, which had long-standing ties to the research centers where the new ideas originated and thus were well positioned to appreciate their implications; and companies with experience in fermentation techniques, which were crucial to "scaling up". The background and competences of these firms played a key role in reinforcing the orientation of the new technologies towards medically-related products and, later, extending them to agricultural products. By the mid-1980's, the interactions between the new research-oriented firms, the pharmaceutical companies, the chemical companies with expertise in fermentation, the venture capitalists, the universities, and the government regulators had produced a distinctive organization of "biotechnology" entities, with a burgeoning (if still largely prospective) product set.

Connections between entities take many forms. Of course, some of the interactions between entities take place in impersonal markets. But many more involve direct and longer-lasting relationships. Pharmaceutical and chemical companies fund university research, place representatives on the boards of smaller, research-oriented companies, send their in-house researchers to scientific meetings. Producing firms carry out extensive market research into the needs and preferences of current and potential customers and use special price and service incentives to consolidate long-term relationships with suppliers and buyers. Competing firms cooperate in various research initiatives, form consortia to jointly produce particular products, work together through their trade associations to lobby legislatures and develop international markets for their products.

Industry structure is then the totality of the connections between the economic entities that make up the industry. To understand how an industry develops, this structure matters, for at least two reasons:

---

[29] For an excellent analytic account of the emergence of the biotechnology "industry" through 1985, see Orsenigo (1989).

- Not everyone knows how to do everything. The competence to perform economic tasks is embodied: particular entities have acquired skills, particular ways of doing things, through experience and over time. It is not generally possible to transfer these skills without immersion in the experience that gave rise to them. To solve new economic tasks, like those that arise in the early days of a new industry, it is necessary to patch together solutions to old problems, as embodied in the entities with the requisite skills. That is, new economic tasks require new entities, which consist of old entities connected in new ways. For example, the research-oriented biotechnology firms combined the technological skills of the university researchers with business plans put together under the auspices of the venture capitalists – and when these firms developed products, they formed partnerships with older firms that embodied competences in production, marketing and regulatory management.

- To decide what to do next – what new products to make or how to improve production processes – a producer has to ferret out opportunities, which requires knowledge outside the producer's current competence. That knowledge is embodied somewhere else – in the tastes or experiences of users of the industry's products, in the theories or experiments of scientific researchers, in the factories or design studios of competitors. And the knowledge can be obtained only through the connections that already exist between the producers and the entities that embody it. Without the mutual experiences that arise from these connections, it is not even possible to conceive of what one needs to know about. So who is connected to whom (and how) determines in part what directions will be explored and how those explorations proceed.

Thus, the process whereby new industries come into being links two interdependent processes, both of which can be viewed as evolutionary in the sense described in section 3.2. The first takes place in the product set; in it, technological and functional relations between existing products give rise, through the interactions of different kinds of agents, to new products. The other occurs in the set of agents, amongst whom new connections create new structures that embody the solutions to the economic problems posed by developing, making and using the industry's new products. The kinds of structure to which these linked processes can give rise and the dynamics by which they do so ought then to be fundamental objects of economic inquiry. Abstract AWs can provide an important modelling tool in this enterprise, particularly by shedding light on what is peculiarly economic about these evolutionary processes.

Acknowledgement. Work on this paper is supported by grants from the National Science Foundation and Citicorp. For many helpful conversations around the ideas presented in this paper, I would like to thank my colleagues Jim Dickey of the University of Minnesota and Brian Arthur, John Holland, Walter Fontana and Jim Pelkey of the Santa Fe Institute; my collaborators Franco Malerba and Luigi Orsenigo of Bocconi University and Giovanni Dosi and Marco Lippi of the University of Rome; and my students in the Artificial Worlds seminar at the University of Minnesota. I would also like to thank Rick Riolo, Michael Wolfson, and Leo Buss for comments on an earlier draft of the paper.

References

Barendregt H (1984) The lambda calculus: its syntax and semantics North-Holland, New York

Bedau M, Packard N (1922) Measurement of evolutionary activity, teleology, and life. In: Langton C et al. (eds.) Artificial life II. Proceedings of the Workshop on Artificial Life, hold February 1990 in Santa Fe, New Mexico. Addison-Wesley, Redwood City, CA, pp 431–462

Buss, L. (1987) The evolution of individuality. Princeton University Press, Princeton

Chiaromonte F, Dosi G, Orsenigo L (1993) Innovative learning and institutions in the process of development: on the microfoundation of growth regimes. In: Thomson R (ed.) Learning and technological change. MacMillan, London

Eldredge N (1985) (in press) Unfinished synthesis. Oxford University Press, New York

Eldredge N, Gould S (1972) Punctuated equilibria: an alternative to phyletic gradualism. In: Schopf T (ed.) Models in paleobiology. Freeman, Cooper, San Francisco, pp 305–332

Eliasson G (1985) The firm and financial market in the swedish micro to macro model – theory, model and verification. IUI Stockholm

Fontana W (1992) Algorithmic chemistry. In: Langton C et al. (eds.) Artificial life, II: Proceedings of the Workshop on Artificial Life, held February 1990 in Santa Fe, New Mexico. Addison Wesley, Redwood City, CA, pp 159–210

Fontana W, Buss L (1993), What would be conserved if the tape were plaxed twice? Yale University Department of Biology preprint

Forrest S (1990) (ed.) Emergent computation: self-organizing, collective and cooperative computing networks. MIT Press Cambridge

Gould S (1989) Wonderful life: the Burgess shale and the nature of history. W. W. Norton, New York

Hull D (1988) Science as a process: an evolutionary account of the social and conceptual development of science. University of Chicago Press, Chicago

Hull D (1989) The metaphysics of evolution. State University of New York Press, Albany

Kauffman S (1990) Requirements for evolvability in complex systems: orderly dynamics and frozen components. Physica D 42: 1–11

Kauffman S, Johnson S (1992) Co-evolution to the edge of chaos: coupled fitness landscapes, poised states, and co-evolutionary avalanches In: Langton C et al. (eds.) Artificial life, II: Proceedings of the Workshop on Artificial Life, held February 1990 in Santa Fe, New Mexico. Addison Wesley, Redwood City, CA, pp 325–370

Lane D (1993) Artificial worlds and economics part II. J Evol Econ 3 (in press)

Langton C (1992) Life at the edge of chaos, In: Langton C, et al. (eds.) Artificial life II. Proceedings of the Workshop on Artificial Life, hold February 1990 in Santa Fe, New Mexico. Addison-Wesley, Redwood City, CA, pp 412–92

Langton C, Taylor C, Farmer JD, Rasmussen S (eds.) (1992) Artificial life II: Proceedings of the Workshop on Artificial Life, held February 1990 in Santa Fe, New Mexico. Addison Wesley, Redwood City, CA

Lindgren K (1992) Evolutionary phenomena in simple: Proceedings of the Workshop on Artificial Life, held February 1990 in Santa Fe, New Mexico. Addison Wesley, Redwood City, CA, pp 295–312

Marimon R, McGrattan E, Sargent T (1990) Money as a medium of exchange in an economy with artificially intelligent agents. J Econ Dynamics Control 14: 329–373

Orsenigo L (1989) The emergence of biotechnology: institutions and markets in industrial innovation. Pinter, London

Rasmussen S, Knudsen C, Feldberg R, Hindsholm M (1990) The coreworld: emergence and evolution of cooperative structures in a computational chemistry. Physica D 42: 111–134

Ray T (1992) An approach to the synthesis of life, in Langton et al. (eds.) Artificial life, II: Proceedings of the Workshop on Artificial Life, held February 1990 in Santa Fe, New Mexico. Addison Wesley, Redwood City, CA, pp 371–408

Salthe S (1985) Evolving hierarchical systems: their structure and representation. Columbia University Press, New York

Somit A, Peterson S (1992) The dynamics of evolution: the punctuated equilibrium debate in the natural and social sciences. Cornell University Press, Ithaca, NY