THOMAS YOUNG

TOWARDS EMERGENT REPLICATORS IN A MOLECULAR ARTIFICIAL CHEMISTRY

A THESIS SUBMITTED FOR THE DEGREE OF DOCTOR OF PHILOSOPHY IN COMPUTER SCIENCE IN THE UNIVERSITY OF CANTERBURY

NOVEMBER 1, 2017

Abstract

All evolutionary systems, natural or artificial, are built from essentially the same three elements: variation, inheritance and selection. What then distinguishes the process of biological evolution by natural selection, which produces such impressive outcomes, from the relatively underwhelming results of artificial digital evolution? We focus on one aspect of this: the emergence of simple replicators from a lower-level foundation in an artificial chemistry. Previous work has either supplied a handbuilt basic (or shortcut) replicator for evolution to work upon, or has provided direct support for replication in the chemistry itself.

Our first research question concerns the relationship between heredity and selective pressure in a theoretical replicator. We construct a simple model of generalized evolution that shows complex and non-obvious emergent behaviour. We show by simulation that inheritance in this model is a target of evolution, and that it evolves under a range of conditions. The degree of inheritance is related to the predictability of environmental change, and the degree of inheritance is tuned by evolution to balance fitness and robustness. Fitness is maximized in unchanging environments where there is little penalty to reduced diversity, while a more diverse population is maintained in changing environments to provide robustness to environmental change. This balance emerges unprogrammed from the underlying model.

Our second research question regards the practicality of realising replication in an artificial system. The investigation is founded on ToyWorld, a highly-modular artificial chemistry that allows us to explore the effect of different combinations of modules, such as for reactant or product selection, upon replicator formation. Our underlying hypothesis is that replicators can form from sequences of linked reaction cycles, where the stochiometry of the sequence is necessarily greater than one for replication.

We first test the influence of two strategies to select reactants for a reaction, and two other strategies for selecting the resulting products post-reaction from the alternative product-sets. Our first reactant selection strategy is to choose reactants with equal-probability from the set of possible reactants without consideration of position; our second strategy is spatial, where reactants are chosen if they are spatially co-located. For product selection, the first strategy is again based on an equal-probability choice from the alternative product sets, while the second biases the choice towards the product set with the greatest energy return, or least energy input (a "least-energy" strategy). Of the four possible combinations of reactant and product selection strategies, the combination of spatial reactant selection and least-energy product selection strategies maximizes the number and cycle-length of the resulting reaction cycles.

Next we search the ToyWorld reaction network for evidence of exact multipliers: re-

peated increasing sequences of the same type of reaction cycle, connected by one or more shared reactants and products. We develop an algorithm for the detection of multipliers in a reaction network, and using that, we find that multipliers do form in ToyWorld, but at low rates, and without great longevity. They occur as a result of Product and Reactant selection strategies, and not by chance alone.

Finally, we examine the reaction network for variable replicators: multipliers that can take any of a set of structural states that appear equivalent under selection. We extend our earlier model of external environmental change, and search for variable replicators in the resulting reaction network. Our results are inconclusive. Candidate variable replicators emerge, but each is endemic to a single run.

This thesis has been a long time in the making, including periods of part-time study and even some earthquakes. I would like to thank my supervisors, Kourosh Neshatian, Tim Bell and, in the early stages, Brent Martin, for their unstinting guidance and support and in particular for their calm assurance that my thesis would one day be done. My great appreciation also to my external examiners, Alan Dorin and Peter Whigham, whose suggestions immeasurably improved the work. And finally, my heartfelt thanks and apologies, to my wife, Susan, and family who I've neglected shamefully at times; your continued encouragement has made all the difference.

Contents

List of Tables

List of Figures

List of Algorithms

1 Introduction

Evolutionary systems are commonly described in terms of three foundational elements, usually some form of variation, selection and heredity[1]. Selection is the mechanism by which possibilities are pruned, and focus is maintained. Selection guides evolution by directing the search into promising areas, acting upon the phenome, or the form developed from a genome under the influence of the environment. However, it is not a generator of new ideas. Novelty instead is the outcome of the second element, variation. Variation extends the search into different areas, and creates new alternatives for selection to evaluate. The final element, inheritance, elevates evolution beyond random search. Instead of starting afresh at each generation, inheritance causes the search to be cumulative, with gains preserved. Variation and inheritance together are the source of new forms; each variation is a modification of the previous generation, forming a spreading tree of related entities when viewed over time. Unlike selection, variation and inheritance are processes that act upon the genome directly.

An open-ended evolutionary system, capable of ongoing evolution without practical limit, needs these components to be themselves subject to evolution, change and adaptation[2]. A system where variation, selection and heredity emerge endogenously from lower-level elements appears to be a promising approach to open-ended evolution in artificial systems. After all, this is the approach taken by life itself.

However, creating the conditions under which evolution emerges endogenously has proved to be extremely difficult in artificial systems, and remains an open problem[3]. Instead of directly addressing this broader problem, therefore, our focus is on only one of the necessary elements for open-ended evolution, the mechanism for heredity within artificial evolutionary systems.

In this work we explore two related aspects of an endogenous mechanism for heredity:

[1] Other related formalisms exist, as discussed in section 2.2; however, none alter these elements fundamentally.

[2] This relates this work to evolvability; see for example Pigliucci (2008), Wagner and Altenberg (1996), and Wagner (2008).

[3] See recent review in Banzhaf et al. 2016

1. In what way does selective pressure drive changes in heredity in a population of evolving replicators?

2. Can variable replicators emerge from a molecular artificial chemistry?

1.1 Context - the emergence of biological replicators

Biological evolution has produced ecosystems of astonishing variety and range, occupying essentially every viable niche on, above, and within the Earth from a common origin in the prebiotic world many millions of years ago. This has been far more than a working out of a single theme; instead a hugely impressive radiation of form and function has left that original ancestor far behind. The only connection that remains is the unbroken lineage of genes that links that earliest ancestor to every descendent organism alive today.

The origin of life was almost certainly contingent, and there is an absence of evidence from early stages (Pross and Pascal 2013). There are many possible pathways, and unless some record remains somewhere (either geological or phylogenetic), the actual path is essentially lost to history.

However, a consensus is forming that early life began with chemoautotrophs fuelled by energy from inorganic redox couples and biomass from CO_2, and that innovations in carbon-fixation created the main branches in the tree-of-life (Braakman and Smith 2012).

The initiation of selection is marked by the advent of the Initial Darwinian Ancestor (IDA), probably from an RNA world, followed substantially later by the Last Universal Common Ancestor (LUCA) (Yarus 2011). It is important to note for clarity that LUCA was almost certainly not a single entity or even species, but is rather a construct of evolutionary genetics because of the likely predominance of Horizontal Gene Transfer in archaic biology (Doolittle and Bapteste 2007). Under Horizontal Gene Transfer, also known as Lateral Gene Transfer, genetic material is incorporated into an organism's genome by methods other than by reproduction along the vertical line of descent.

Horizontal Gene Transfer is thought to have been so common in early life that there was no single common ancestor; instead genes from multiple lineages intermixed during this early stage into all lineages today (Ragan, McInerney, and Lake 2009). The starting point for the development of modern replicators might therefore have begun with horizontal inheritance, although the changes can then be heritable in a vertical sense: in organisms with genes, by definition under Horizontal Gene Transfer the change becomes part of the target's genome and hence heritable. In entities without genes, there are forms where the horizontal change can be subsequently inherited vertically[4].

[4]Mechanisms for horizontal transfer don't require replication (by definition). Therefore it's possible that they can act as a precursor for the development of replication. That is, horizontal inheritance plus selection might be sufficient for replication, rather than replication being required for inheritance.

Although the advantages of a distinction between genome and phenome are discussed by many, including (Taylor 1999a, section 7.2.3) and indirectly Von Neumann (1966); there is no inherent dependency on this in Evolution by Natural Selection (ENS). Early evolution may have involved the inheritance of complete portions, or components, of the phenome before the advent of a distinct genome, while research into Horizontal Gene Transfer (*e.g.* Ochman, Lawrence, and Groisman (2000), Pace et al. (2008), and Ragan, McInerney, and Lake (2009)) has shown that not only was component transfer between species a major driver of early evolution, but a horizontal component-based mechanism continues to exist even in many of today's organisms that have a genome built from DNA.

Two alternative models exist for the step from the prebiotic world to IDA (fig. 1.1): replication- or genes- or RNA-first, and metabolism- or protein-first. Both metabolism and replication were almost certainly required for IDA, however. Self-replicating RNA enzymes are described in Lincoln and Joyce (2009), forming the basis of a selective system (also see Cheng and Unrau (2010) and Powner, Gerland, and Sutherland (2009) for formation of RNA in prebiotic conditions).

Some elements of IDA are thought to still be with us in lineages of informational (for protein synthesis and RNA transcription) and operational genes (for some standard cellular processes) (Ragan, McInerney, and Lake 2009), for example the ribosome and ribonuclease P (RNase P) (Wilson and Lilley 2009). The next major transition was to the protein world, although predominance of RNA transcripts leads to suggestions that it should more accurately be called the RNA-Protein world (Altman 2013).

A self-sustaining autocatalytic network[5] is generally considered essential (Pross and Pascal 2013), but not sufficient (Hordijk, Kauffman, and Steel 2011). Both competing models—replication-first and metabolism-first—build on this. In the case of replication-first, through autocatalysis as expressed by self-replication of oligomeric compounds; in metabolism-first, by cycles and networks. From another perspective, metabolism-first privileges function, while replication-first privileges descent.

The main difficulty with the replication-first model concerns the sizeable step required from abiotic compounds to template-based, or information-based, replication (although ribonucleotides conceivably could form in pre-life conditions (Powner, Gerland, and Sutherland 2009)). Templates encode information in biology, so template-based replication requires an encode/decode mechanism to store information in the template and to later retrieve it, as well as an information code. This is significantly more complex than simple multiplication. By contrast, the main issue with the metabolism-first model concerns the

[5]Perhaps in the form of a Reflexively autocatalytic and F-generated (RAF) set, a "set of molecules and reactions which is collectively autocatalytic in the sense that all molecules help in producing each other (through mutual catalysis, and supported by a food set)" (Hordijk, Kauffman, and Steel 2011).

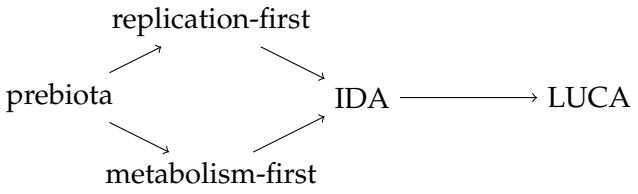

Figure 1.1: Some of the stages proposed in the evolution of early life

necessary shift from composomal or holistic[6] inheritance to template-based, and the ability of holistic systems to represent a sufficient range of stable states for unrestricted evolution (Vasas, Szathmáry, and Santos 2010).

In summary, the origin of life can be seen as the transition from chemistry to biology, and this transition seems tantalizingly similar to the artificial evolutionary goal of moving from simple, uninteresting, systems to creative systems which evolve.

1.1.1 Replicators in evolutionary theory

Although artificial systems are the main focus of this thesis, there is a natural overlap with research in biology into the origins of life. In particular the transition from chemistry to biology has several central concerns in common with our field of emergent replication in artificial systems. For example, how does the interaction of lower-level elements form long-lived composite structures? What are the necessary and sufficient conditions for evolution? How does an evolutionary process begin in a non-evolutionary system? This overlap means that much of the most relevant previous work is to be found in origins-of-life research.

In this particular section we adopt a *biological* classification of replicators from Zachar and Szathmáry (2010) to structure our review of *artificial* replication, and in other sections of this thesis we shall extend our scope as needed to include pertinent results from biology. However, it's important to stress that we are informed by biology, but do not hope or intend to contribute to biology in return. This remains a work concerned primarily with *artificial* systems.

In the transfer from the prebiotic world to the biotic one, the proportion of information held by a predecessor that could be passed on to its successor increased, over many generations, from "none" to "nearly all". Vasas and Fernando (2012) relate heritability to the correlation between the parent and child entities. Heredity is therefore a matter of degree,

[6] See discussion in section 1.2.3.

rather than being a binary relationship (related or not-related), and that opens up the possibility that a series of gradual changes might over time transform a very poor replicator into a very good one. Of course, one possible process for this transformation might well be ENS.

But what exactly does it mean for something to be a replicator? Is a rock that erodes to form grains of sand a replicator[7]? Is a set of autocatalytic reactions that splits into two replicating? Dawkins (1976) was the first to define replicators, including a range from biological genes to non-biological ideas (memes) in the scope of the definition. Many other definitions and formulations followed as various properties or features were examined.

Relatively recently, Zachar and Szathmáry (2010) saw a need to reexamine the definition primarily to resolve issues of discrimination between entities which are clearly replicators or not replicators and those which are borderline, and between biological replicators and non-biological or cultural ones.

They conclude that a *replicator* is "any autocatalytic entity for which there is a selection process defined", using autocatalysis in the usual sense of a reaction, generally cyclical, that increases a reactant (say A) n-fold, but in this case the products are not necessarily *identical* to the reactant As, but instead *equivalent* when subject to selection. Selection is the mechanism used to determine if entities are equivalent. Instead of producing more As, replication produces something indistinguishable by selection but which is not identical—we can call these products Bs to make the distinction clear.

Replication therefore has the form $\Sigma x_i + A \rightarrow \Sigma y_j + \Sigma B_k$ (where some n of the $B_0, B_1 ... B_k$ products are equivalent to A under selection, $\rightarrow$ means "yields", and x and y are some reactants and products, respectively, as needed), and where selection has the fairly standard definition of "a process, acting on a particular population of entities in a particular environment, which sorts entities according to their phenotypes." (Zachar and Szathmáry 2010, p.21)

Note that replicators in this definition are subsets of autocatalytic entities: replication implies multiplication, but replicators must also have some property (a phenotype) that is visible to an external process (selection.) Zachar and Szathmáry further discriminate between types of replicator based on two further factors: variability and heredity.

If the replicator can change and yet still be equivalent under selection, then it is a variable replicator, which can exist in several equivalent states. An exact replicator on the other hand is one in which any change causes the entity to fail, or to respond differently under selection. Variability is the first stage towards metabolism and a genotype/phenotype distinction. Heredity describes which elements of variability are passed on to a descendent (the genotype.) If variability can be passed on, then the replicator is informational—the

[7] In Bourrat (2015) rocks are given as examples of *persistors*, unable to reproduce and subject to only a "weak" form of selection for hardness.

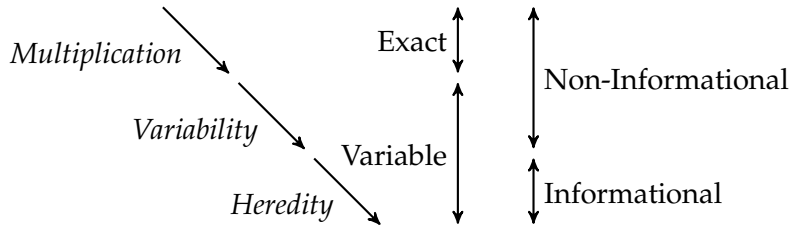

Figure 1.2: Replicator systems according to the hierarchy of Zachar and Szathmáry. Diagram redrawn from Zachar and Szathmáry 2010, Fig.13.

change in information can persist. If all changes are passed on, then the replicator is an *ideal* replicator, where the genotype encompasses the entire entity, while if they are not, the replicator is *real*.

Building on this, Zachar and Szathmáry, extending earlier work by Szathmáry (1999) and Szathmáry (2006), go on to develop the following general classification of replicators as a logical sequence of forms (Zachar and Szathmáry 2010, p.21, line breaks and numbering added for emphasis):

1. "The simplest of replicators is the *exact replicator*, which is *non-informational*, and any change made to it causes a change in the phenotype.

2. If a variation can arise in the structure in such a way that it does not change equivalence of the entity, then it is a *variable* replicator, with more than one stable state.

3. If such changes can be passed on to the offspring then the replicator is *informational*.

4. If the non-heritable part is constructed by a developmental process, then the replicator is a *reproducer*."

An alternative hierarchy, from the same source (Zachar and Szathmáry 2010), is shown in simplified form in fig. 1.2. Here the hierarchy is by a cumulative arrangement of features from simple multiplication through to heredity; the main discriminators from the earlier definition are arranged to the right of the figure.

1.2 Previous work

We now move from biological systems to artificial ones. Many modern artificial replication systems trace their intellectual ancestry back to one of two main starting points—Bagley and Farmer (1991) and Kauffman (1986)—both motivated by the goals of origins-of-life

research. Although these works have been since overtaken by later research, for completeness they are mentioned here.

Bagley and Farmer investigates Auto-Catalytic Metabolisms with catalysts included as the mechanism to provide the permanently elevated concentrations required. In contrast to many later methods, the model is based on differential equations to determine concentrations, rather than simulation. This approach severely restricts the model's ability to describe open-ended scenarios as the equations must be defined in advance.

Kauffman describes a binary polymer model to investigate the formation of autocatalytic sets, and in particular how this is affected by the probability of catalysis. In Kauffman polymers are modelled as strings of binary digits of maximum length, n, and molecules catalyse randomly chosen reactions for ligation or cleavage, with a constant probability (P) that a molecule will catalyse any given reaction. Again, the model is primarily analytical and logical, although in discussion simulation is suggested as the means by which the proposed extensions to the model might be implemented.

Before examining other systems that show at least some degree of replication, according to the classification of Zachar and Szathmáry, there is one other interesting system in which replication was a goal, although in this case it was not achieved. *RBN-World* (Faulconbridge 2011) is an artificial chemistry where the entities take the form of Random Boolean Network (RBN) (Kauffman 1969), with the addition of a bonding mechanisms to allow for their composition and decomposition. The resulting form of RBN is called a bonding RBN or bRBN. Larger structures are formed by "bonding" two independent bRBNs at each bRBNs bonding node. "All reactions are between two reactants; it is assumed that more complicated reactions can be expressed as a series of two-reactant reactions with intermediate structures." The choice of reactants is described as "Gillespie-like" (*i.e.* Gillespie (1976)), and essentially random, uncorrelated in any way with reaction energies or rates (Faulconbridge 2011, chap.8).

Each bRBN is a synchronous RBN, made up of a number of nodes, each with an initial state (*true* or *false*) assigned randomly and with an input/output matrix assigned randomly. Finally k(=2) inputs are established per node. The bonding method uses "cycle length as the bonding property and equality as the bonding criterion....bonds only exist between bRBNs that have the same cycle length." After initial bond formation the algorithm recalculates cycle lengths, and checks again for equality. This might result then in decomposition (records are kept of composition operations so that the reverse decomposition can be easily done.)

A number of parameters affect the behaviour of the chemistry, and so a series of experiments sampled from the parameter-space, and then used a GA, to search for interesting variants as measured by non-catalysed "loops" (as the preferred measures of auto-catalytic sets and hypercycles are too rare for use as a measure) (Faulconbridge 2011, chap.8).

1.2.1 Replication without self-replication

The systems in this section are capable of replication, but not self-replication. The difference is that they rely for replication on mechanisms that are completely external to the entity. As such, they are of limited applicability to our work, although they do illustrate various methods by which simple chemical entities might replicate before the establishment of even minimal self-replication.

The model for *Chemical Evolution by Natural Selection* in Fernando and Rowe (2007) and Fernando and Rowe (2008) is driven by origin-of-life objectives ("the evolution of chemical networks that lead to autonomous systems"), and takes the form of a simulation of laboratory experiments of lipid aggregates ("phase separated 'individuals', e.g. liposomes") in a reactor. The molecules and food molecules (that make up an autocatalytic cycle) share a common representation and underlying chemistry, while replication, unlike the (rare) autocatalytic cycles in Faulconbridge (2011), is by " division by externally imposed agitation, i.e. replication rather than self-replication.".

Mutations in cycles are generally not heritable as mutant copies are rarely functional in autocatalysis (Vasas and Fernando 2012). Mutation is thus problematic as a source of variation in these autocatalytic systems; one alternative to mutations, introduced in Fernando and Rowe (2007), is "avalanches", to form new cycles from temporarily-preferred side-products of the base networks. Two existing species are chosen at random, tested to see if a reaction between them is possible by thermodynamics, and then the resulting products seeded in the reactor at low concentrations. These new reactions may be autocatalytic, or the products may complete a food-set for another autocatalytic reaction, or they may enable a side-reaction from an existing autocatalytic one. Introducing these new species into the reactor therefore can trigger the rapid formation of a series of novel products—a chemical avalanche.

1.2.2 Exact replicators

Remember that autocatalysis by definition is replication; if the entities are molecules, the autocatalytic reaction $\Sigma x_i + A \rightarrow \Sigma y_j + \Sigma A_k$ replicates the molecule A (*e.g.* Lifson (1997)).

An exact replicator is where the As on the left-hand side of the equation and the right-hand side are identical; any change results in a functionally different entity (that is, different under selection), and so each entity exists in only one stable state. This fundamentally limits the evolutionary potential of an exact replicator as there is no heritable variation.

1.2.3 Variable replicators with limited heredity

If states are to be a store of information for a replicator, information must be preserved from generation to generation. The number of alternative stable states forms an upper

bound on the total information held by a variable replicator. A limited heredity replicator (in the original terminology of Szathmáry) is one in which the possible number of states that the replicator may take is relatively small, and so the amount of information that can be maintained is likewise restricted. This limits the possible forms that the phenotype, derived from the genotype, may take, and hence the long-term evolutionary potential of the replicator.

Gánti (2003) and Eigen (1971) showed that distinct, organisationally different alternative autocatalytic networks in the same environment might compete, and the fittest would prevail. A number of models have been proposed since where autocatalytic networks form stable components that can be inherited in a modular fashion in a process called compositional inheritance. All these models are analytical in nature, in many cases based on the Ordinary Differential Equation (ODE) model of Farmer, Kauffman, and Packard 1986.

In the most well-known of these models, the Graded Autocatalysis Replication Domain (*GARD*) model (Segré et al. 1998), highly catalytic molecules determine the properties of the compotype (compositional genotype), and these are not necessarily inherited equally. Instead a child may, or may not, inherit one of these molecules and so its properties may be similar to or very different from its parent (Vasas et al. 2015; Vasas et al. 2012; Vasas and Fernando 2012). Information fidelity varies widely; the Eigen threshold (Eigen 1971) applies, and mutation rates overwhelm selection (Vasas et al. 2015; Vasas et al. 2012; Vasas and Fernando 2012).

The model of Vasas et al. (2012), derived from Farmer, Kauffman, and Packard (1986), and extended in Vasas and Fernando (2012) and Vasas and Fernando (2012), tests the GARD hypothesis that compositional inheritance is possible where there is a parent-offspring correlation in molecular composition. Autocatalytic cores made up of one or more linked autocatalytic loops to provide the compotype. The core forms an attractor, where one core equals one attractor, but multiple cores are required of course for selection.

Multiple cores provide multiple attractors, but the attractors must be stable for selection to be stable and meaningful. Unlike GARD, which generates only single-core networks, the model in Vasas is capable of multiple cores, but as cores are the equivalent of a single bit of heritable information, it's hard to see core-based inheritance being capable of unrestricted heredity as there are practical limits to the number of stable cores that can co-exist in a system.

1.2.4 Variable replicators with unlimited heredity, but with shortcut replication mechanisms

We now move on to systems where the entities are phenotypically-rich, but where the replication mechanism itself, although part of the entity, is provided by a *shortcut*; the mechanism is directly implemented by the experimenter rather than being an emergent property

(Banzhaf et al. 2016). As the mechanism is independent of the entity it is not under selection; shortcut replicators are not, by our definition, open-ended as they have a limited ability to self-improve. They differ from systems in the earlier section on Replication without self-replication in that systems with a shortcut are capable of self-replication, but the replication system remains independent of the evolutionary pressures acting on the entity.

One of the most well-known systems with a shortcut replication system is *Squirm3* (Hutton 2007; Hutton 2002), an artificial system capable of life-like open-ended evolution (creativity), initially developed with the goal of testing the hypothesis of Taylor (2001) (Hutton 2002, p.341).

All elements in Squirm3 are constructed from atoms defined by *types* (e.g., a, b, c...) and *states* (e.g., 0, 1, 2...). Atoms and hence molecules are located on a 2D grid, and molecules cannot overlap or pass through each other. All reaction rules, such as *R1: e8 + e0* $\rightarrow$ *e4e3*, are pre-specified (Hutton (2007, p.4) and Faulconbridge (2011, p.49)), and a shortcut consisting of a set of eight rules in Hutton (2002) is sufficient to replicate single molecules. As all reaction rules, including those responsible for replication and hence inheritance, are exogenous to the model, the inheritance mechanism in Squirm3 is not evolvable.

Individual entities in Hutton (2002) are simply single molecules. Hutton (2007) introduces cells made up of a collection of molecules and bounded by a membrane of a particular atom type of limited reactivity; the membrane is intended to allow individuals to benefit from innovations by protecting their internal reactions from others. With a greatly increased set of predefined reaction rules (from the original eight rules to thirty-four in Hutton (2007) for the cell replication shortcut, extended again in Lucht (2012)), each cell has the capacity for division and mutation. Selection is purely by indirect competition for the raw materials (atoms in the environment) required by the reactions in a cell; interactions between individuals are purely through this indirect competition (niche construction without direct interaction). Cells in Hutton (2007) are incapable of making use of resources from other cells (as they are effectively protected by non-reactant membranes) and so an intermittent exogenous mechanism ("floods") is used to return the atoms in a number of randomly chosen cells to the environment.

1.2.5 Variable replicators with unlimited heredity

Our final class of artificial replicator systems are where there are no fundamental limitations known to prevent unlimited self-improvement, but where, for one reason or another, it remains unproven or undemonstrated.

Replication in these systems emerges from base rules [8] without being directly specified by the experimenter or designer. However, even here most systems still jumpstart this by

[8] *Level-0* rules in the terminology of Banzhaf et al. (2016)

providing a seed or universal ancestor (*e.g.* Ofria and Wilke (2004)) that contains a working module for replication that can be then modified (for good or ill) in each subsequent generation.

The small number of proposed unlimited heredity systems mostly fall into two closely-related types: string-manipulation systems, where entities are algorithms encoded in strings of primitive operators, and pure automata, where the algorithms are similar to assembly-language programs, running directly in some virtual machine. In both cases, the phenotype, or evolutionary-value, of the entity is the result obtained by running its algorithm.

Not all unlimited heredity systems fall into these classes however; we conclude this section with a summation of two other systems, StringMol and GraphMol, that are neither pure string-manipulations systems nor pure automata but which contain elements of each.

String-manipulation systems

The most influential string-manipulation system is Fontana's Algorithmic Chemistry, or *AlChemy* (Fontana 1992), originally intended as an exploration of innovation in biological-like complex systems. Entities in AlChemy are represented by strings of symbols that form an algorithm according to a LISP-like language. Entities interact through applying the algorithm (or function in AlChemy terminology) of one entity to that of another through a recursive evaluation procedure.

A new entity or object, h, is formed by the composition of two parent functions, f and g: $h(x) = f(g(x))$. More formally, replication is defined as the interaction expression, h, of two randomly chosen objects f and g, if, and only if, the interaction expression contains at least one variable and one primitive operator, and is shorter than some maximum length (Fontana 1992, p.173–p.180). New objects in AlChemy are therefore the children of two parents.

How then is the interaction expression $((' f)(' g))$ between f and g evaluated to produce h? AlChemy is a form of pure LISP (with some "minor idiosyncrasies") based on toyLISP, with six primitive operators defined in Fontana (1992, p.205). The interaction expression is defined in Fontana (1992, definition A.9, p.204) as $V[((' f)(' g)), ()] = (V[f, (a \leftarrow g)])$, using the notation $V[e, L] = v$ to mean the expression e with the "association list" L (a list of value assignments between atoms and expressions) evaluates to v. The result h is described as $f(g)$ and the process as $f + g \rightarrow ((' f)(' g)) \rightarrow h + f + g$.

Clearly reproduction in AlChemy is self-referential and endogenous: the process to construct a child object is defined in the code of the parent objects. However, inheritance doesn't follow straightforwardly as the reproduction process is unusual in some important ways. Producing new objects as a function of two parents means that the relationship between parents and child is not a straightforward mutation or other syntactic difference, but rather a complex functional relationship. What this means for the relationship between

the parent's fitness and the child's fitness is not obvious. It seems that the fitness differences in AlChemy might be more extreme than in other systems where parent and child have a more straightforward relationship.

Automata

The previous systems we've described have all, in one way or another, been simulations of evolution. With Automata, the technique moves from simulations towards direct evolution, even if still within a simulated, digital, world. Entities are algorithms competing against other algorithms for computer resources. The major insight behind these systems is that by relaxing the restrictions that prevent programs from modifying themselves or others, all of the elements thought necessary for ENS are present in the modern digital computer.

Coreworld (Rasmussen et al. 1990), inspired by the early computer game "Core War"[9], set segments of simplified assembly code into competition in core memory. The assembly command to copy values from one memory location to another can spontaneously introduce errors into the copying, and hence can introduce evolutionary variation. However, as reviewed by Ofria and Wilke (2004), the system "collapsed into a non-living state" perhaps because organisms could copy over each other in the shared system memory (Ofria and Wilke 2004).

In *Tierra* (Ray 1991) mutations are introduced during replication by randomly flipping bits during the copy operation (at a given rate of generally between 1 bit flip per 1,000 and 2,500 instructions copied). This rate is set by the experimenter, and is not evolvable. Mutations can also be introduced by the copy algorithm itself; as it is an algorithm defined in the organism in standard Tierra instructions (and hence fully embedded), mutations in the algorithm during a copy will be inherited by the child. The initial copy algorithm is part of the 80-instruction ancestral creature documented in Ray (1991, app.C).

In Ray's words, "...this approach involves engineering over the early history of life to design complex evolvable organisms, and then attempting to create conditions that will set off a spontaneous evolutionary process of increasing diversity and complexity of organisms" (Ray 1991, p.3). As Taylor criticises though, the problem with "engineering over" is that we don't understand the natural examples well enough to engineer them at all (Taylor 2001)

Tierra has been the testbed for a number of other works. For example, Sugiura et al. (2003) converted Tierra into a string manipulation systems, introducing a set of 140 regular-expression based rewriting rules where each rule encoded one or more of the original 32 Tierran instructions. The initial rewriting ruleset was manually generated by the experi-

[9] See `corewar.co.uk`

menters, although details are unclear. Unlike in Tierra where the instruction set is fixed during a run, the rewriting ruleset for each organism itself could evolve through a separate genetic algorithm. This algorithm removed the least applied rules and inserted the same number of new rules generated by mutating (through duplication, removal and addition of operations) a selection of the most applied rules. Although results support the claim that the ruleset as well as the genome evolves, the use of a separate genetic algorithm for ruleset evolution artificially separates the ruleset from the genome; the feedback loop from genome back to ruleset is broken.

Tierra was also the starting point for Taylor (1999b) and Taylor (2001) to explore the creation of Artificial Life (Alife), by adding cell regulation, parallel processes and energy modelling.

Avida (Ofria and Wilke 2004), introduced in the summer of 1993, is based on Tierra with improvements including better metering and measuring, and a 2D lattice or well-stirred reaction vessel topography (unlike the shared linear memory of Tierra, for example).

Like Tierra, the automata engine in Avida follows a Turing-tape metaphor, with instruction, read, write, and flow control heads that can be moved forward and backwards through memory using relative rather than absolute addressing. Instructions are grouped into instruction sets, with the default set containing 26 instructions, and by definition every program is valid. Each organism runs on its own virtual automata; the only interaction between organisms is via resources in the shared environment and through competition for virtual machine CPU cycles based on "merit" or fitness. Direct Tierra-style interactions by insertion of code into another organism is not enabled by default, but is possible through configuration.

Phenotypes take the form of computations (entities take in resources, perform computations that result in merit, and perhaps produce output or by-product resources): "...by inputting numbers from the environment, performing computations on those numbers, and outputting the results. The organisms receive a benefit for performing specific computations associated with resources" (Ofria and Wilke 2004). Crucially, the resources in the environment are not the same as the elements of the organisms (instructions from an instruction set.) Avidan organisms are not fully embedded in their environment.

New organisms are created asexually by the parent first allocating memory for a child. The parent's read-head is placed at the beginning of its code, the write-head placed at the start of the newly allocated memory and successive `h-copy` instructions copy the instruction from the read-head to the write-head and advance both. After all instructions have been copied, `h-divide` splits the child from the parent (all instructions between read-head and write-head go to the child) and start execution in both parent and child from a clean state. Variation comes through mutations which can be introduced through either `h-copy` (the write-head writes a random instruction rather than the instruction at the

read-head) or in `h-divide` (a single random instruction may be deleted or added from the child code). Both forms of mutation happen with a fixed probability set by the inventor: COPY_MUT_PROB, INS_MUT_PROB, and DEL_MUT_PROB for `h-copy`, and DIVIDE_INS_PROB, DIVIDE_DEL_PROB for `h-divide`. However, there is a second, evolvable source of mutations during replication: the replication process itself is embedded in the organism, as a set of instructions, and so changes to this algorithm during the copy will persist in the child. The self-replication algorithm is initially defined in the ancestral organism used to seed a run, and as documented in Ofria and Wilke (2004, A1.3) consists of 15 instructions.

Avida is extremely configurable, but provides little guidance or theoretical justification for any particular configuration. Indeed this variability allows it to function most usefully as a general testbed for experiments, *e.g.* "in one experiment we wanted to study a population that could not adapt, but that would nevertheless accumulate deleterious or neutral mutations through drift" (Ofria and Wilke 2004).

Amoeba-II (Pargellis 2001) shares similar features to Tierra and Avida in that it is an instruction-set based automata, but unlike in those systems, replication spontaneously emerges in Amoeba-II. The replication process requires four steps: "register initiation, memory allocation, copying of the parent's instructions to the child (embryo), and division where the child is initiated as a cell on its own" (Pargellis 2001, p.69). Two specific instructions in the parent's genome are required at a minimum. The MALL command allocates memory, and a virtual CPU to the child, and DIVD performs the division. Register initiation and copying of the parent's instructions can be done in a variety of different ways, and a good deal of the interest of the Amoeba-II system lies in the evolvability of these mechanisms. Selection is by efficiency of replication, where faster and more efficient replicators replace less efficient ones in the population (as Pargellis (2001) says, "Amoeba has only one task: replication") and so there is selective pressure from the least efficient, but functioning, mechanisms in the direction of improved replication. To complement selection and inheritance, variability is provided by the DIVD instruction, which introduces a mix of instruction substitutions, deletions and insertions into the child's program at a fixed but low rate.

Inheritance in Amoeba-II is almost too effective. Fit organisms rapidly evolve into extremely rapid reproducers and out-replicate all other entities, leaving a mono-culture (Pargellis 2001).

StringMol and GraphMol

The related automata *StringMol* (Hickinbotham et al. 2011) and *GraphMol* (Nellis 2012; Nellis and Stepney 2014) explore computational novelty through embodiment: "Our aim is to improve novelty-generation algorithms by making their biological models richer." No

measure is proposed for novelty. The authors state that they're not even sure it is possible, but an informal definition is used that sees novelty as the outcome of increasing embodiment (Nellis 2012, p.87).

A persistent theme in Alife attributes the differences in richness between biology and Alife to the grounding of biology in the physical world, or in other words, embodiment. Nellis (2012) examines a pre-existing definition of embodiment, concerning the relationship that exists between a "system" and an external "environment", before suggesting a modification based on the specifics of biological genome-copying where the complexity and subtleties of the copying derive directly from the embodiment of an abstract concept (copying) in physical-world mechanisms. The modification suggested by Nellis (2012) is to define embodiment in terms of a world, encompassing both environment and system, in which mechanisms of interest exist (such as a template-copying replicase) that produce abstract phenomenon (such as copying.)

The base elements in StringMol are single-character symbols, each representing a microcode instruction. These combine to form strings representing molecular microprograms (the strings of the name). The general arrangement feels very similar to that in Avida or Tierra, as seen for example in the various pointers, or indexes, into the code (instruction, read, write, and flow for iterations).

GraphMol replaces the strings in StringMol by graphs: "[t]he world defined by GraphMol contains chemicals (represented as graphs) that bind to each other via multiple binding sites, and then run simple computer programs (encoded in the graphs) that modify the binding of these chemicals.". No explicit rationale for graphs is presented, other than as a natural extension from StringMol given the stated importance of a rich binding mechanism.

An underlying principle in both systems is that the mechanism of evolution must be itself evolvable; functions such as template copying must be embodied mechanisms in the world so that they can be affected by evolution, and so evolved. Crucially both StringMol and GraphMol have embodied template copying, where a "replicase" molecule can copy another by an algorithm such as in alg. 1.

Although the algorithm itself is unremarkable, the underlying functions allow StringMol and GraphMol to explore various forms of embodiment. Each of the four functions referenced in alg. 1 (*start*, *at-end*, *char-copy* and *next*) can be either "crisp" (*i.e.* perfect or precise) or embodied (variable or fuzzy, and subject to evolution). As an example of an embodied function, StringMol's *start* function uses pattern-matching to determine the binding region between the replicase and the other string where the replicase should begin copying. By changing the sub-sequences in either string the strength of bind can be varied, with a corresponding shift to the beginning of the copy region even though the pattern-matching algorithm itself does not change (it is a "level-0" component in the terminology

of Banzhaf et al. (2016)). The same pattern-matching is used in StringMol's *at-end* function to determine the end of the region to be copied.

```
1 i ← start(string B)
2 while i not at-end(string B) do
3     string A(i) ← char-copy(string B(i)))
4     i ← next i
5 end
```

Algorithm 1: The canonical algorithm for template copying used by StringMol and GraphMol, from Nellis and Stepney (2014).

StringMol includes an embodied *start* and *at-end*, crisp *next*, and stochastic *char-copy*; GraphMol has an embodied *start*, *at-end* and *next*, with a crisp *char-copy*. The primary difference between the two is opposite approaches to *next* ("in order to investigate a method of embodying the copying process that would be completely different from Stringmol's." (Nellis 2012, p.145))

These mechanism differences (the choices of which functions are crisp and which are embodied) result in different outcomes for the overall system: "Stringmol exhibits macromutation and two chemical copying; GraphMol exhibits two types of quasispecies, cooperative and parasitic. These two systems use the same domain (emergent evolution) and metamodel (machines copying strings), but different computational models." Other combinations would presumably show different behaviour again.

Despite the use of the term "embodied", those functions that aren't "crisp" are not in fact fully self-referential as the pattern-matching algorithm itself remains unaffected by evolution. There is also a disconnect between the properties of the targets of the matches and the functioning of the algorithm. The algorithm matches on the symbol, and is completely unaffected by the meaning or properties of those symbols, the underlying microcode. In this it differs significantly from a completely endogenous system such as a biological replicase, where the match or bind is actioned by the same chemical rules that construct and maintain the replicase from component atoms.

1.2.6 Conclusions

Of the systems discussed in section 1.2 (summarised in table 1.1), the most successful and influential approach to date has been through automata. In these systems, entity multiplication is by an algorithmic process where code within the parent, when executed, repeatedly copies the instructions from parent's code to a child using a specific copy operator (*e.g.* `h-copy` in Avida.) Deliberate errors introduced during copying by the operator (erroneous copying, see Zachar and Szathmáry (2010, p.16) for the biological equivalent)

Table 1.1: Artificial chemistry replicators in selected previous works, categorized according to Zachar and Szathmáry (2010). ODE-based models such as GARD and Vasas et al. (2012) are not included as they are not constructive and hence incapable of unlimited evolution; RBN-World (Faulconbridge et al. 2010) is not included as self-replication has not been reported.

Artificial replicator	Model form	Multiplication	Variability	Heredity
Squirm3 (Hutton 2002)	Molecular	Exogenous rules	Exogenous (hand-crafted)	Unknown—depends on rules
AlChemy (Fontana 1992)	λ-calculus	Endogenous (function composition)	Variable (functional equivalence)	Holistic
StringMol (Hickinbotham et al. 2012)	Automata, encoded as strings	Endogenous (algorithmic)	Variable (functional equivalence)	Endogenous
GraphMol (Nellis and Stepney 2014)	Automata, encoded as graphs	Endogenous (algorithmic)	Variable (functional equivalence)	Endogenous
Coreworld (Rasmussen et al. 1990)	Automata	Endogenous (algorithmic)	Variable (functional equivalence)	Unknown
Tierra (Ray 1991)	Automata	Endogenous (algorithmic)	Variable (functional equivalence)	Endogenous
Avida (Ofria and Wilke 2004)	Automata	Endogenous (algorithmic)	Variable (functional equivalence)	Endogenous
Amoeba-II (Pargellis 2001)	Automata	Endogenous (algorithmic)	Variable (functional equivalence)	Endogenous

provide one source of variation; another comes from the potential for different instruction sequences to be functionally equivalent, so providing variability as defined by Zachar and Szathmáry.

The initial multiplication process in each of the reviewed automata copies all instructions from the parent; therefore the entity's genotype is the same as its instruction set. Heredity is thus holistic and in the notation of Zachar and Szathmáry, $H = V$ (discussed further in chapter 2). For example, in Tierra, the 80-instruction ancestral creature includes code to determine its beginning and end in memory and then to copy everything between the two to a child entity. However, this multiplication code is part of the entity and subject to variation; it is itself evolvable and in principle at least capable of evolving towards a model where only a proportion of the entity is copied.

1.3 Research questions

Of the forms of replicators defined by Zachar and Szathmáry (2010) and described in section 1.2, only unlimited heredity informational replicators provide both the representational range and the evolutionary flexibility for open-ended evolution:

- Exact replicators lack heredity as they cannot pass on variability to their descendents.

- Limited heredity replicators are capable of only a limited number of distinct states and so suffer from a restricted representational range.
- Informational replicators based on replication "shortcuts" are in practice still limited as their inheritance/heredity mechanism is not evolvable.

This thesis is therefore motivated by the following research questions:

RQ1: *In what way does selective pressure drive changes in heredity in a population of evolving replicators?* This is driven by two considerations. First, existing work in evolutionary theory (see chapter 2) has suggested that heritability is the result of selection and variation under constant selective pressure. There is also a broad body of work in biology exploring aspects of the connection between fitness (the counterpart to selection) and environmental variability (for example, canalization). We are interested in a model of heredity that incorporates environmental variability as variable environments are the superset of fixed ones, and as a response to variability appears to be a factor in shaping the inheritance mechanism in biology. We believe that the factors underlying this relationship between heredity and variability also apply in artificial systems. Second, the response of heredity to selection is clearly evolutionary, and our hypothesis is that heredity can emerge through evolution in a molecular artificial chemistry. A coherent model would illuminate the mechanism by which this may be achieved in practice in an artificial system. Therefore, we investigate a model for the response of heredity to selective pressure in chapter 2.

RQ2: *Can variable replicators emerge from a molecular artificial chemistry?* Although replicators have been studied extensively, previous models have provided, or assumed, the provision of a replication mechanism rather than having the mechanism evolve. We instead will test the position that replicators can in fact emerge from lower-level reactions between molecules in an artificial chemistry. We begin our examination of this hypothesis with a review chapter to introduce artificial chemistries in chapter 3, then in chapter 4 we introduce a molecular artificial chemistry, modularized for experimentation, and show that certain configurations are capable of reaction cycles. We then investigate the emergence of multiplication, the simplest form of replication, from cycles in chapter 5, and conclude by examining variability in the sense of Zachar and Szathmáry (2010) in chapter 6.

1.4 Methods

In this section we explain the methods used in the remainder of this work. For each research question we follow a common process:

1. Relevant previous work is used to identify areas where the research question is well understood, and those areas where further work would be beneficial.

2. We construct a simulation model that attempts to capture our understanding of the problem for the areas where further work is needed.

3. Based on the context formed by the research question and the previous work, we form a hypothesis of how the simulation should function if our understanding is correct.

4. We then proceed to test this hypothesis by experiment using the simulation.

Each simulation model is parameterised. Parameters are elements in the simulation model that can take different possible values, where the different values may (or may not) lead to quite distinct behaviour of the simulation. The purpose is two-fold: first, to allow for the investigation of a range of models simply by changing parameter values (rather than changing entire models) so as to broaden the *scope* and hence the applicability of the results, and second, to permit us to test the *sensitivity* of the simulation to different parameters overall. The combination of these two allows us to robustly justify the scope and strength of any claims that arise from the experiments.

As each parameter may take any of a range of values, and given that there will be a number of parameters, we need a way to limit the size of the parameter space for experimentation.

We do this by first testing the response of the model to each parameter and identifying those that do not make a statistically significant difference; the set of these parameters, those to which the model is insensitive, allows us to establish the scope over which all claims will hold. Those to which the model is on the other hand sensitive, or responsive, are tested separately in all later experiments; any claim must be made conditional on the particular level of each of these parameters.

Second, we use statistical "design of experiments" (*e.g.* Montgomery 2009) methods to simplify the number of separate experiments needed. There are many design of experiments strategies, but they mostly fall into two standard groups.

The most common approach is via some form of factorial design, where each parameter of interest is represented by a factor taking some small number of discrete values, or levels (two levels being most common) and the analysis model constructed from runs that systematically work through a series of combinations of factors at different levels. The emphasis here is on the response given particular factor (and hence parameter) values.

A variant of this is a fractional-factorial design, which takes a well-defined sample of factor levels to further reduce the number of level combinations while still maintaining an acceptable level of experimental strength. Latin square (and Graeco-Latin and Hyper-

Graeco-Latin) designs are well-known forms [10] of fractional factorial design for the case of a single factor of primary interest combined with one or more "nuisance" factors which are of no interest to the experiment but have an effect on the result. However, in this work the relatively small number of representative levels required, combined with the reasonable speed of the simulation, is such that we can remain with a simpler full-factorial design.

Statistical analysis must be carefully considered in any multi-factorial experiment design as pair-wise tests of the mean such as Student's t-test cannot be extended to multiple groups (factors) without significant loss of accuracy. Instead we will use Analysis of Variance (ANOVA) when analysing the relative contribution of each factor, or combination of factors, to the observed values of the response variables.

Note that there are some differences between the design of experiments in the physical world and in simulations, with the most significant being the sources and understanding of experimental errors. In simulation, experimental runs are exactly reproducible, absent of any dependency on factors external to the simulation. Variation is explicitly introduced usually through a random number generator, which can be seeded to produce the same sequence of numbers again and again. This means that the practice in real-world experiments of "blocking" to control external variation is not required in simulation experiments. However, replicates where the same combination of factor values is run several times each with a different random seed value, remain valuable, but in this case less to control for experimental error and more to record the variation across a series of runs and the sensitivity of the model to parameter settings.

1.5 Contributions

This thesis makes the following novel contributions:

1. A model for the effect of selection pressure upon heritability, incorporating an exogenous property for fitness suggested by Bourrat (2015).

2. A new and modular artificial chemistry, ToyWorld (chapter 4), developed to facilitate experimental comparisons between artificial chemistry components, to a degree unusual in other experimental artificial chemistries (*e.g.* table 3.1).

3. An assessment of the influence on cycle formation, as a step towards replication, of various reaction and product selection strategies in a molecular artificial chemistry (chapter 4.)

4. A demonstration of the emergence of simple replicator forms from a molecular artificial chemistry (chapters 5 and 6.)

[10] For example, see discussion at `http://www.itl.nist.gov/div898/handbook/pri/section3/pri3321.htm`

5. An algorithm for reaction cycle detection based on sampling to generate a set of seed molecules alg. 8.

6. An algorithm for the identification of exact multipliers from reaction cycles, based on a specific definition for exact multipliers as two or more copies of the same reaction cycle species, where the reaction cycle species has stoichiometry greater than one, and where each cycle in the multiplier is connected to at least one other multiplier cycle by a molecule that is a product of one cycle and a reactant in the other ()alg. 7.)

7. An algorithm for the identification of variable replicator candidates (without consideration of selection) from reaction cycles, where we define variable replicators as multipliers that can occupy any of a limited set of states without losing their underlying identity (alg. 11.)

A version of section 4.7 was published as Young and Neshatian (2015), and material from section 4.6 appears in Young and Neshatian (2013).

The source code for all software developed for use in this work is available on GitHub under the GNU GPL v3 open source licence:

- For ToyWorld, the artificial chemistry from chapter 4, see Young 2013a.
- For the updated version, ToyWorld2, used in the experiments in chapters 5 and 6, see Young 2013b.
- Finally, the model source from chapter 2 is available at `http://github.com/th0mmeke/thesis-models`.

2 Heritability, Heredity and Inheritance

This chapter attempts to establish an overall high-level framework for the evolution of replicators; in the following chapters we shall test this framework through stepwise enhancements of a basic artificial chemistry along the way towards the long-term goal of a full replicator.

In an earlier section, we described the model of Zachar and Szathmáry (2010) for heritability and variation (in the sense of changes to which selection is blind.) Replicators are classified by Zachar and Szathmáry by reference to the elements of this model—variable replicators are those that can acquire variation during their lifetime, informational replicators are those that can pass on these changes to their offspring—but the model does not address causality.

Now, if variability is represented by V, and heritability by H, then an informational replicator is *ideal* if $H = V = 1$ (all variation is heritable, and all changes are inherited) and *real* if $H \leq V \leq 1$ (Zachar and Szathmáry 2010). Our working hypothesis is that a theoretical replicator (rather than a specific actual one, which may be constrained by its structure or evolutionary history), is not fixed *ideal* or fixed *real*, but instead lies at a point between the two determined by selective pressure.

Therefore we have the following goal for the work in this chapter: can we describe and model the relationship between heritability and selective pressure for a theoretical replicator?

2.1 The relationship between heritability, heredity and inheritance

There is a subtle difference between heredity, heritability and inheritance; inheritance is the process that results in heredity (for example, Lamarckism or ENS), while heredity is the relationship between entities linked by inheritance and is measured at the population-level by heritability. As noted by Griesemer (2005), "[o]ne must clearly distinguish between heredity (a relation), heritability (a capacity), and inheritance (a process)."

Heredity describes the similarity of an offspring to its parent, and depending on the

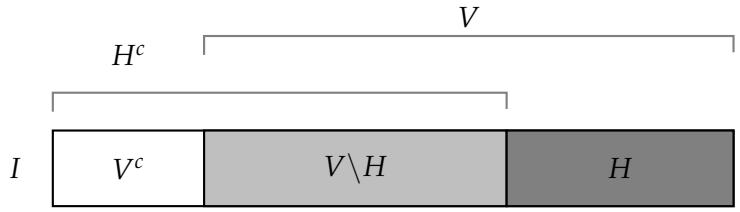

Figure 2.1: Reproduction of Figure 10 from Zachar and Szathmáry (2010). This diagram of the total information content of a replicator (I) illustrates the relationship between heritable information (H) and variable information (V). Areas V^c and H^c indicate regions which are fixed or conserved; V^c because change would cause the entity to become either non-functional or to be treated differently under selection, and H^c the portion of variability that is not heritable.)

context, can refer to either a single property or to a group of properties shared between offspring and parent. Heredity is related to variation: low variation implies high heredity and vice versa. For the single property (or in biology, trait) case, when the value for a parent's property and an offspring's property are identical we say we have complete or full inheritance and if this is the case for all entities in the population, then $H = 1$, where H is the heritability for the trait or property. Conversely, if there is no relationship there is no heritability ($H = 0$) and the entities are effectively unrelated. We can extend the measure to a group of properties simply by averaging the degree of similarity for each property across all properties.

2.2 Evolution by Natural Selection

As has been said before, "[w]e take it as given that biology instantiates ENS" (Watson 2012), but that doesn't mean that the algorithm of biology is in fact ENS. Adaptation in biology appears to precede Natural Selection, reinforcing that adaptation is possible without ENS (Watson et al. 2010), and other forms of evolution altogether have been proposed in artificial systems (*e.g.* the compositional evolution of Arthur (2009)), and on the boundaries of ENS in the domain of living systems (*e.g.* composomal inheritance and Horizontal Gene Transfer). These variants are examined later, but first we concentrate on canonical ENS.

Godfrey-Smith (2007) examined a number of summaries of ENS taken from the most influential items in the literature. The purposes of the summaries varied, but have interest for us "as attempts to capture some core principles of evolutionary theory in a highly concise way." Incidentally, as an illustration of the difficulty of definitions, although the usual aim for a summary is to "give conditions that are sufficient ceteris paribus for a certain kind of change occurring", Godfrey-Smith (2007) notes that the scope of most summaries

is somewhat ambiguous. They can be read as either being discriminatory—"this process is ENS"—or alternatively in providing conditions that will result in ENS when it is assumed that the meaning of ENS is clear. In other words, these are alternative *constitutive* or *causal* readings.

The main examples selected by Godfrey-Smith (2007) are:

1. "Owing to this struggle for life, any variation, however slight and from whatever cause proceeding, if it be in any degree profitable to an individual of any species, in its infinitely complex relations to other organic beings and to external nature, will tend to the preservation of that individual, and will generally be inherited by its offspring. The offspring, also, will thus have a better chance of surviving, for, of the many individuals of any species which are periodically born, but a small number can survive. I have called this principle, by which each slight variation, if useful, is preserved, by the term of Natural Selection, in order to mark its relation to man's power of selection." (Darwin 1859)

2. "If there is a population of entities with multiplication, variation and heredity, and if some of the variations alter the probability of multiplying, then the population will evolve. Further, it will evolve so that the entities come to have adaptations...." (Maynard-Smith, in Griesemer (2001))

3. "Any entities in nature that have variation, reproduction and heritability may evolve" (Lewontin 1970) and "1. Different individuals in a population have different morphologies, physiologies, and behaviors (phenotypic variation). 2. Different phenotypes have different rates of survival and reproduction in different environments (differential fitness). 3. There is a correlation between parents and offspring in the contribution of each to future generations (fitness is heritable)." (Lewontin 1970)

4. "...evolution will occur whenever and wherever three conditions are met: replication, variation (mutation), and differential fitness (competition)" (Ofria and Wilke 2004, quoting Daniel Dennett)

Godfrey-Smith (2007) concludes that the core requirement for ENS is some "combination of variation, heredity, and fitness differences", although he identifies a number of differences between the summaries. For example, the most commonly cited summary is Lewontin (1970), but unusually that formulation states that "fitness is heritable" whereas typically phenotypic heredity (as appears in Lewontin's later 1980 summary) is stated as sufficient for a trait to evolve.

These differences are also discussed in Griesemer (2001), in particular with reference to the variations between the concept of inheritance in Darwin which includes both heritability (a capacity) and inheritance (a process carrying the capacity); Lewontin, which stresses

the heritability while assuming inheritance, and Maynard Smith's multiplication which is actually about inheritance; his heredity is both (Griesemer 2001).

Quite apart from these differences in interpretation, the major theoretical difficulty in a literal application of ENS to artificial systems is captured succinctly by Griesemer (2005): "Darwin's theory of evolution by natural selection is restricted in scope. One sense in which it is restricted is that it refers to organisms." Organisms are not defined, but the context and scope is clearly biological.

Although ENS is usually only discussed within the context of modern day biology, as we've seen when examining the definition of ENS there is nothing exclusively biological about the standard formulations. By abstracting the concepts of variation, selection, and in particular inheritance or heredity, a generalised form of ENS can be developed that goes beyond the usual biological readings.

A specific example lies in the transition from non-living to living things. One of the main problems (of several) in extending ENS to prebiotic entities is that the meaning of the foundational elements—such as variation, selection, heredity, multiplication—are generally couched in biological terms. For example, heredity is often discussed in terms of alleles, or traits, rather than in more general ways. However, there must have been some point at which the prebiotic processes transitioned to ENS, and it is clear that this transition did not happen abruptly in a population of fully-fledged modern organisms. The pathway to ENS, the evolution of evolution itself, is one of the main open problems in origins-of-life research, and it involves an extension of those foundational elements back into the prebiotic world.

This example of the transition to life is an example of *extending* ENS; two further examples are now presented that are *alternatives* to ENS: the evolution of technology, or evolution in an entirely non-living domain, and DaisyWorld, demonstrating regulation without selection.

Arthur (2009) describes a mechanism for the evolution of technology, where evolution is used in the sense of "all objects of some class are related by ties of common descent from the collection of earlier objects". Evolution in technology occurs by using earlier technologies as building blocks in the composition of new technologies, and these new technologies then become building blocks for use in later technologies, and so on. Arthur calls this "combinatorial evolution." But what is the starting point? How is this regression grounded? Arthur proposes that the capture and harnessing of natural phenomena starts each lineage and provides new raw components for inclusion in later technologies. Bourrat (2015) comments that distributive evolution (where only the distribution of elements changes, as result of selection or drift) cannot result in novelties: Arthur anticipates this objection by stating that novelty comes from this incorporation of new phenomena from the source, the natural world.

The first obstacle to a more general scope is the existence of innovations such as the jet engine, laser, railroad locomotive, or QuickSort computer algorithm (to name Arthur's examples.) Innovations seem to appear without obvious parentage; they do not appear to be the result of gradual changes or adaptations to earlier technologies. Arthur's reply is to look inside the innovation and to recognise that each is made up of recognisable components or modules[1]; the key lies in the nature of heredity in technology. Technologies are formed by combining modules of earlier technologies. These groupings start as loose assemblages to meet some new function, but over time become fixed into a standard unit (for example, the change in DNA amplification mechanisms from assemblages of laboratory equipment to standard off-the-shelf products.)

However, even Arthur in his rejection of "Darwinian evolution" describes a process that is still ENS, relying on selection, variation and inheritance. This is not obviously the case in our final example, DaisyWorld.

Lovelock and Margulis (2011) propose DaisyWorld as a example of an alternative to selection for regulation, based on two feedback loops. Saunders (1994) explains how regulation can emerge in DaisyWorld without selection: the planet's temperature is adjusted to meet the conditions for maximum growth of the daisies without the daisies adapting to the planet. "As a result, regulation, one of the most fundamental and necessary properties of organisms, appears without being selected for. What is more, it appears as a property not of the daisies, on which natural selection may have acted, but of the planet, on which, as Dawkins rightly points out, it could not" (Saunders 1994).

The fundamental insight in DaisyWorld is that individuals modify environments (by niche construction): the daisies adapt the planet (specifically the temperature for maximum growth) to suit themselves, rather than adapting themselves to the planet, and in fact there's little benefit to adaptation by the daisies to the planetary conditions. As Saunders (1994) states, "the ability to withstand a greater variability is not the result of Darwinian adaptation. On the contrary, it exists because of the absence of Darwinian adaptation."

Daisyworld describes an alternative mechanism for adaptation, without invalidating ENS as the usual or predominant model. ENS remains our best model for evolution overall, while we recognise that in some specific environments or scenarios other mechanisms may also be possible.

2.3 Previous work

Over the years a significant body of literature has accumulated on models of biological inheritance. While standard population genetics describes the dynamics over time of genotype frequencies, it remains deeply rooted in the biology of genotypes and alleles.

[1]Similar in many ways to composomes, or compositional replicators (Segré et al. 1998)–see section 6.1.

Different forms of inheritance, such as acquired characteristics (Lamarckian inheritance), have been compared and contrasted to the canonical non-acquired form in ENS by a number of authors (including for example Jablonka et al. (1995), Paenke et al. (2007), and Gaucherel and Jensen (2012)), going back to the competing models for inheritance of Darwin and Lamarck.

Early hypotheses in the metabolism-first or replicator-first debate led to the recognition of the significance of error threshold rates for mutations during copying (Eigen (1971) in biochemistry and Muller's Ratchet (Muller 1964) in population genetics, compared in Wagner and Krall (1993)), and later investigation of the interaction between inheritance (of genotypes, when non-acquired), selection (on the phenotype), and development (linking genotype to phenotype) inspired the theory of neutral landscapes (Kimura 1968).

By comparison, little explicit modelling of inheritance has been done in Alife. Most relevant work has already been reviewed in section 1.2 where it generally forms one element of a wider investigation into open-ended evolution. In relation to Evolutionary Algorithms (EAs), many models for inheritance from variation and recombination have been proposed. These however are less relevant to our needs as they rely on exogenous, or external, predefined mechanisms rather than emerging from the properties of the genotypes.

Paixão et al. (2015) sought common principles between population genetics and evolutionary computation with the goal of unifying the two fields, starting from a broad description of evolutionary processes as a "population undergoing changes over time based on some set of transformations". A transformation can be decomposed into a collection of (stochastic) operators, with an operator representing a probability distribution of potential outcomes; some operators act on phenotypes, others on genotypes. Various operators drawn from evolutionary computing are defined: for selection (uniform, proportional, tournament, truncation, cut, replace); variation from mutation (uniform, single-point), and variation from recombination (one-point crossover, k-point crossover, uniform crossover, unbiased variation).

Evolution is then a trajectory through a space of distributions; it can therefore be seen both as a sequence of population transformations and distribution transformations. When considered as a series of distribution transformations, Paixão et al. (2015) sees a correspondence with the Estimation Distribution Algorithms (EDA): "In an Estimation Distribution Algorithm (EDA), the algorithm tries to determine the distribution of the solution features, e.g. probability of having a 1-bit at a particular position, at the optimum. Some EDAs can be regarded abstractions of evolutionary processes: instead of generating new solutions through variation and then selecting from these, EDAs use a more direct approach to refine the underlying probability distribution. The perspective of updating a probability distribution is similar to the Wright-Fisher model."

Paixão et al. (2015) concludes by demonstrating how existing "classical models in theo-

retical population genetics and in the theory of evolutionary computation" can be mapped into the framework of classified and categorised operators. Most of the various population genetics models can be represented, with the exception of some topic-specific EA models, while genetic programming models are omitted for reasons of balance between "simplicity and inclusiveness." However, Paixão et al. (2015) is of limited relevance for this work as it is in effect a constitutive framework, providing a series of tests to classify existing models based on a general meta-model of evolution, rather than a causative model which is where our interest lies. Variation, in the form of reproduction and mutation, is only one element of the framework, and Paixão et al. (2015) does not specifically address the emergence of heredity from variation.

More specific models of inheritance and heredity can be found in three works comparing the adaptive value of acquired and non-acquired characteristics. In the first two works, Jablonka et al. (1995) and Paenke et al. (2007), the comparison is made with respect to an environment that alternates between two states, E_0 and E_1. Each environment is associated with a corresponding adapted phenotype P_0 and P_1, respectively. In both works variation only affects the ratio of each phenotype in the population (directly in Jablonka et al. (1995), indirectly via a "predisposition" or tendency in Paenke et al. (2007)); no new phenotypes are created, and as inheritance is modelled only at a population level, that is without reproduction, we cannot use these models to investigate the emergence of heredity.

Inheritance however is modelled in Gaucherel and Jensen (2012), the third work, in relationship to individuals. Two models are presented, the first and simplest describing a non-spatial scenario conceptually similar to that in Jablonka et al. (1995) and Paenke et al. (2007), while the second examines a spatial world based on DaisyWorld (Lovelock and Margulis 2011). Focussing on the first and most relevant of the two models, reproduction is the middle of three repeated stages: first "annihilation" where the population size is adjusted to some practical level, then reproduction, and finally development.

In the second stage (reproduction), individuals, represented by a single trait, or phenotype, value, are chosen for reproduction with some probability (based on the trait value), and, to model mutation, the child given a trait value that slightly varies from the parent's value: $trait_{child} = trait_{parent} + \delta$, where δ is described as taken from a uniform distribution of given range around the parent's trait value. [2])

The most relevant previous work is that of Bourrat (2015) who, in the course of examining the difference between evolution, natural selection and ENS, models the emergence of heredity in unchanging environments. Bourrat (2015) demonstrates that imperfect inheritance is not compatible with ENS using an argument by contradiction (Bourrat 2015, p.96): he lists the three conditions for a population to evolve solely by ENS, and then continues on to show that at least one of those conditions is incompatible with imperfect inheritance

[2] although this appears to be an error; δ should instead be centred on 0.

(as it happens, no production of new variants). The context is unequivocally related to biology; his examples involve genes, traits, phenotypes and drift.

2.3.1 Bourrat's model of biased inheritance

The six applicable models in Bourrat (2015, chap.3) are designed to explore the implications of bias on inheritance, where *unbiased* means a trait is uncorrelated between child and parent, in practice it means that trait is taken as a random choice between some lower and upper bounds (Bourrat 2015, p.153). *Biased* inheritance is naturally the opposite: there is some correlation between parent and child values for a trait, and so some prediction of traits is possible—a parent can "pass it on". Note that Bourrat (2015, p.173) expressly notes that his "biased inheritance" is not the "transmission bias" of the second term in the Price equation: "The bias in 'biased inheritance' is in reference to the type of the parent(s) (biased toward the type of the parent), while the bias in 'transmission bias' refers to a departure from an event of perfect transmission." Bourrat unfortunately doesn't formalise his model descriptions; instead there is a reference (Bourrat 2015, p.129) to a NETLOGO 5.02 implementation, and textual narrative descriptions of each model and the results.

Model 1 begins with a population of 5000 "persistors" (that is, entities without the ability to reproduce), each of which has a survival rate (viability or the likelihood of surviving at each time step) between 0 and 0.99. Unsurprisingly, all eventually die. Model 2 introduces a single "procreator", capable of reproduction with both survival and fertility rates (offspring per unit time), into the population of persistors. The traits of the offspring of the procreator are uncorrelated (that is, unbiased inheritance) to those of the procreator, and the model now consists of *selection* $\rightarrow$ *reproduction* $\rightarrow$ *check-for-overcrowding* (Bourrat 2015, p.141) at each timestep. Again, all entities eventually die.

Model 3 begins to get interesting: Bourrat (2015) replaces the procreator by a "minimal reproducer" which differs from a procreator in that the ability to procreate is itself a heritable trait, although as "minimal" it is an unbiased trait. As such, the offspring of the minimal reproducer may be either minimal reproducers or persistors without the ability to reproduce. Now the population size drops then increases to maximum size with about 10% of the population being minimal reproducers. However, the proportion of high fitness (that is, high viability) entities doesn't increase beyond about 0.05, so there is no cumulative adaptation.

Biased (in fact, perfect) inheritance of viability is introduced in Model 4; the offspring inherit the viability of their parent. The proportion of high fitness entities rapidly approaches the upper limit of 1.0, as expected as high viability entities live longer and so produce more offspring while low viability entities die sooner and so produce less. Fertility is random, but viability is heritable.

The most significant model is Model 5 which adds a variable ability to procreate to

Model 4, while viability remains inherited with perfect fidelity from the parent. The variation is provided by the addition of a *mutation* stage so that the model now consists of these stages at each timestep: *mutation* → *selection* → *reproduction* → *check-for-overcrowding* (Bourrat 2015, p.153). At each mutation stage, there can be an increase or decrease in both the ability to transmit the ability to procreate, and in the degree of bias (that is, the relationship to the parent's ability) in the ability to procreate. The first defines the proportion of offspring of the parent who are themselves able to procreate; a low trait value for the parent means a low proportion of siblings can procreate, while bias represents the correlation between the parent and offspring's abilities to transmit procreation: low bias means the offspring's ability is only weakly correlated with the parents ability. The initial population contains entities with viability in the full range $[0, 1)$, an ability to procreate in $[0, 0.2)$ and initial bias of 0. Both the ability to procreate and the bias *increase* over time in the population towards the upper limit of 1.0. The model is asymmetric with respect to the change of inheritance of ability to procreate: reductions lead to extinction of a line, while increases lead to increased population. Bourrat's conclusion is that an initial population of unreliable reproducers (a low proportion of procreating offspring, no bias) will evolve into one of reliable reproducers that is, inheritance can emerge.

This is extended in the final Model 6, where Bourrat demonstrates the combination of inheritance of procreation with that of another trait, viability. The model begins with a population produced by the end of Model 5 - a set of entities that can reliably transmit the ability to procreate to their offspring. Using the same structure for trait inheritance as in Model 5, Model 6 shows that inheritance of viability or fitness can also emerge. In total then, the entities at the conclusion of Model 6 have full inheritance of multiple traits.

2.3.2 Limitations to the model presented by Bourrat

There are some limitations however to these otherwise insightful models. First, heredity and fitness (viability) are treated as independent traits. But the mechanism for heredity is the thing that copies the information that generates an offspring's traits, so in practice they are not independent.

Second, while Model 5 has perfect inheritance on viability and demonstrates emergent inheritance on procreation, Model 6 shows emergent inheritance on viability while beginning with perfect inheritance on procreation (as it begins with a shortcut population of entities assumed to have been produced by Model 5.) Thus Bourrat does not show in either model the simultaneous emergence of inheritance of both procreation and viability.

Finally, the model assumes that the problem learnt by evolution is amenable to perfect understanding, and that there is one and only one optimal solution. This is a corollary of the model design where fitness is absolute and unchanging. As fitness represents an implicit relationship between an entity and its environment, then in Bourrat's model this

relationship is also fixed and unchanging. Evolution is omniscient with full visibility into the world. However, in the real world and in the artificial domains of interest, the relationship between entity and environment is less sure. The environment itself may either change, or be uncertain. Under these conditions it is unlikely that values of 1.0 would be possible, or indeed helpful, as a perfect bias value effectively is removing any source of variation from the population. This is unexplored by Bourrat.

The work of Bourrat (2015) therefore has the following limitations:

- Heredity and fitness are unconnected in the model.
- The simultaneous emergence of inheritance of procreation and viability is not addressed.
- Most importantly, the effect of changing environments on inheritance is unaddressed.

2.4 Hypothesis

Bourrat (2015) argues that heredity may in fact be a product of evolution rather than a precursor, and that the process of inheritance emerges from the action of selection and variation upon a population (section 2.3.) In other words, *variation and selection are sufficient for inheritance*.

The degree of variation between generations is important; if there is no correlation it effectively means evolution is operating as an unguided, random, search while complete correlation means that each generation is a copy of the previous one, and there are no novelties.

Our intuition is that the average heritability of the population will be inversely related to the degree of environmental change. More specifically, if we consider evolution to be a means by which a population learns how to adapt to an environment, the degree of environmental change can be described as the degree to which it's possible to learn the environment. We intuit that this is related to predictability, and also to complexity.

This seems a reasonable supposition. Adami (2002) recasts population fitness in terms of complexity, specifically his physical complexity measure: "It is probably more appropriate to say that evolution increases the amount of information a population harbours about its niche (and therefore, its physical complexity)" (Adami 2002). Prokopenko, Boschetti, and Ryan (2009) discusses the information-theoretic view of the benefits of complexity.

The initial hypothesis of Bourrat (2015) assumes a perfectly predictable environment that can be exactly "learned" by an evolutionary algorithm. If we assume that learnability is related to predictability, then we can restate and expand the hypothesis in terms of predictability:

Hypothesis 2.1. *Variation and selection are sufficient for inheritance, where the degree and variance of inheritance is proportional to the predictability of the environment.*

The conceptual basis for the experiments in this chapter is outlined in fig. 2.2.

2.5 Evolutionary model relating variation and inheritance

We now introduce a general model of the relationship between *Variation*, *Selection* and *Inheritance* in a population of evolving abstract entities (alg. 2) where the key elements, *heritability* and *fitness*, are represented as explicit parameters:

- *Heritability* is the likelihood that a child's value for a property will resemble that of its parent. The range is $[0, 1]$, where a value of 0 means that the value for a child's property has no inherent relationship to its parent's value, while $heritability = 1$ means the child's value will be identical to the parent's. This is a slight modification of the classic definition of heritability, which is concerned with the population-level attribution of the sources of variation, but co-option, rather than inventing new terminology, seems appropriate given both have to do with the source of heritable variation.
- *Fitness* represents the probability that an entity will survive and possibly also reproduce, and has the usual range for a probability of $[0, 1]$.

This strategy of making otherwise derived variables explicit is also followed by Bourrat, who describes it as "explaining variables which have previously been taken for granted in a model (such as reproduction and inheritance), by reference to other, more fundamental variables present in the model" (Bourrat 2015, p.129). Our model owes a direct debt to Bourrat in the representation of heritability as an explicit parameter (related to the two parameters "heredity of the ability to procreate" and "transmission of the ability to procreate" in Bourrat (2015)).

In other respects, however, it resembles a standard EA. The model consists of a population of abstract entities, and two simple functions—*Selection* and *Variation*—that each transform the population. Each run consists of a fixed number of time steps (generations), where at each step these two functions are applied in sequence to the current population to form a replacement population as documented in alg. 2.

```
1 for Generation ∈ [1 ... number of generations] do
2     Population← Selection(Population)
3     Population← Population ∪ Variation(Population)
4     if |Population| < MINIMUMSIZE then break
5 end
```

Algorithm 2: Algorithm for the inheritance and variation model.

Other elements of the model are as follows:

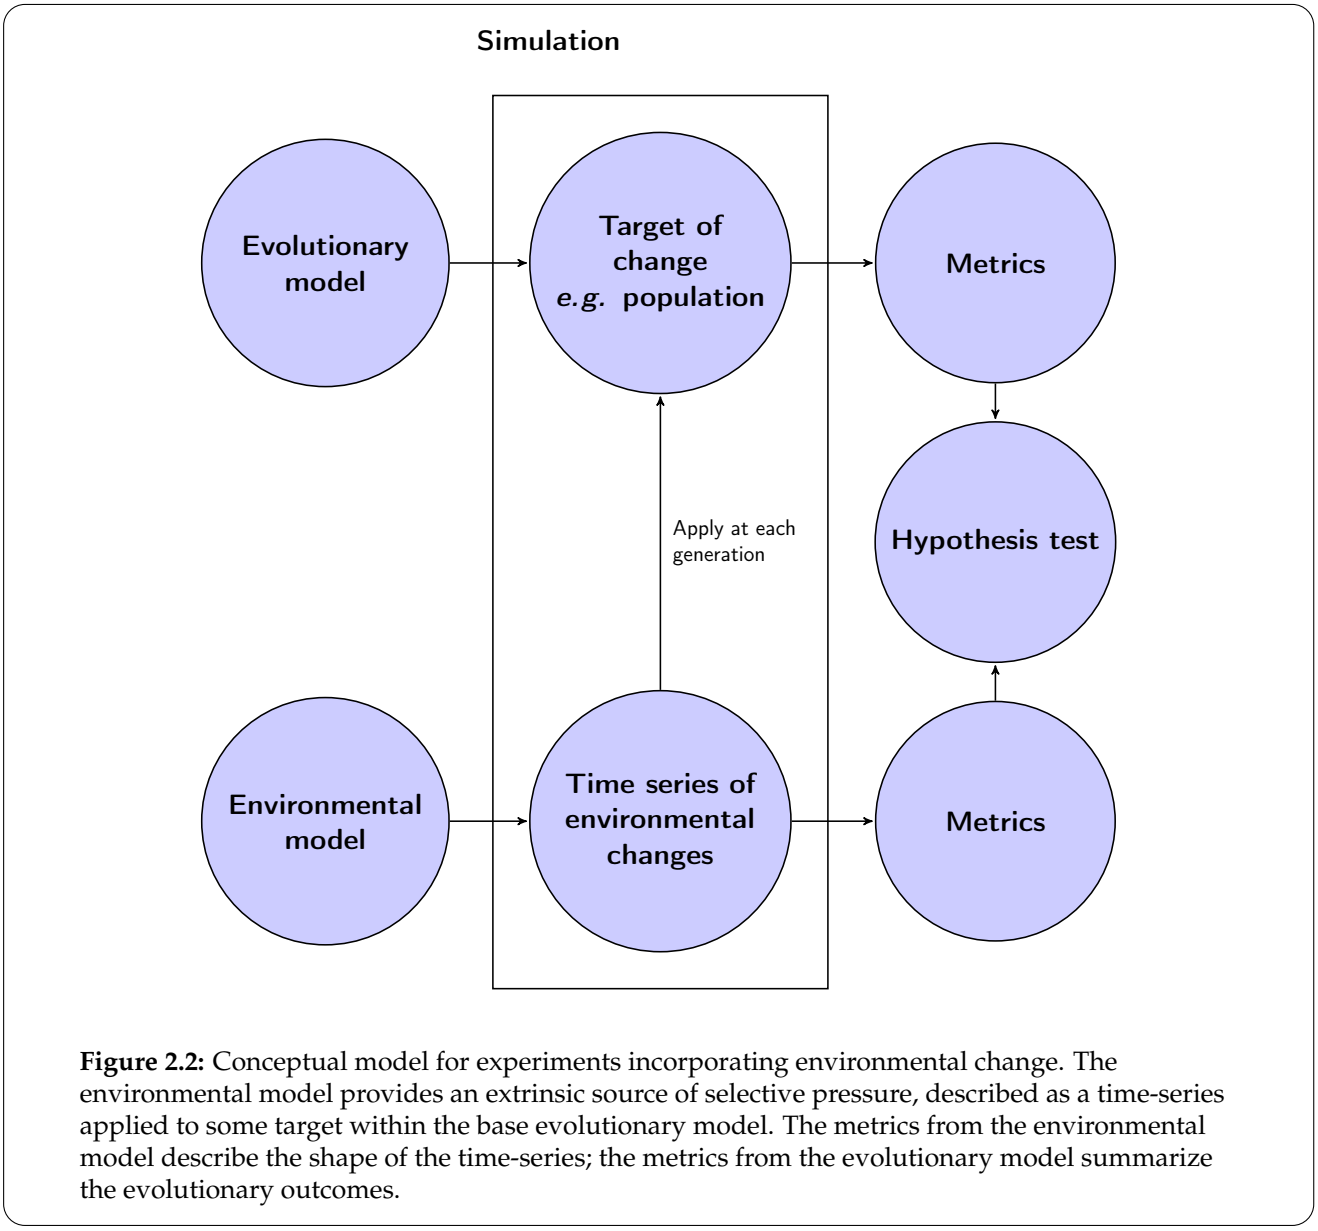

Figure 2.2: Conceptual model for experiments incorporating environmental change. The environmental model provides an extrinsic source of selective pressure, described as a time-series applied to some target within the base evolutionary model. The metrics from the environmental model describe the shape of the time-series; the metrics from the evolutionary model summarize the evolutionary outcomes.

```
1  Function Selection(Population):
2      Population_new ← ∅
3      for Entity ∈ Population do
4          with probability p_selection:
5              Population_new ← Population_new ∪ Entity
6          end
7      end
8      return Population_new
9  Function Variation(Population):
10     Population_new ← ∅
11     for Entity with Fitness and Heritability in Population do
12         with probability p_reproduction:
13             for some number of Children ∈ U[0, n_children] do
14                 Child ← Reproduce(Entity)
15                 Population_new ← Population_new ∪ Child
16             end
17         end
18     end
19     return Population_new
```

Algorithm 3: Definitions for the functions *Selection* and *Variation*.

- The *population* consists of n entities, where $|n| \geq 0$.
- Each *entity* is fully described by two properties—*fitness* and *heritability*.
- *Selection* forms a new population by choosing entities from the current population, as defined in alg. 3. The probability of an element being included in the new population ($p_{selection}$ in the algorithm) may be either a fixed value, or equal to its *fitness*.
- The new population created by *Variation* consists of a number of new entities ("children") for each entity ("parent") in the current population (see alg. 3). Each entity has children with probability $p_{reproduction}$, and if it does, the number of children it has is some random number between 0 and $n_{children}$. The properties of each child are related to the properties of its parent by a mapping, represented in the algorithm by the function *Reproduce*, which maps the parent value to a value in a range with an expected value equal to the parent's value and some degree of correlation captured by the value of the function *Range*.

The model is parameterised to make it easy to describe different specific models within this general structure; these parameters are defined in table 2.1.

Table 2.1: Definitions for all parameters of the model.

Parameter	Value	Description
$p_{reproduction}$	$[0, 1]$	Probability of reproduction
$p_{selection}$	$[0, 1]$	Probability of selection
$n_{children}$	$n_{children} \in \mathbb{Z}_{\geq 0}$	Maximum number of children per parent
Reproduce	$entity \mapsto entity$	Function to create a new entity based on an existing one

Table 2.2: Factors mapped to model parameters, plus factor levels, for the initial investigation into model sensitivity to parameter settings.

Factor	Model parameter	Number of Levels	Levels
$p_{reproduction}$	$p_{reproduction}$	2	0.66 or $fitness$
$p_{selection}$	$p_{selection}$	2	0.66 or $fitness$
$n_{children}$	$n_{children}$	2	2 or 5

2.6 Hypothesis test under stable environment

Returning to the overall goals for these experiments, section 2.4 allows us to make two predictions for stable environments:

1. Average inheritance will tend towards perfect inheritance.

2. The population variance of inheritance will decrease more than would be expected by chance alone.

The first test therefore is to examine if inheritance emerges from low-heritability and low-fitness initial conditions, and then the second test is whether the population variance for inheritance decreases as predicted. Remember that as discussed earlier, inheritance is the outcome of the relationship between parent and child traits, as represented by the measure of heritability, and that heritability ranges between 0, for no correlation between parent and child, and 1.0 for perfect correlation.

2.6.1 Experiment design

The factors and levels along with their mapping to model parameters are given in table 2.2.

Each combination of factor levels has 10 replicates, to give a total of $2 \times 2 \times 2 \times 10$ or 80 experiment runs. The initial population for each run consists of 5,000 entities, each with heritability and fitness chosen independently and with uniform probability from the range [0,0.3]. Each run contains 500 generations.

2.6.2 Sensitivity of the model to parameter values

In the absence of any restrictions on population size there is nothing to prevent a growing population eventually exceeding the capacity of the simulation system. This is unfortunately an unavoidable difficulty in experiments with exponential growth systems rather than a limitation of the theoretical model.

The size of the population is determined by how population entities are introduced and removed. In standard Evolutionary Computation (*e.g.* De Jong (2006, p. 50)) the choice of strategy is important to the performance and outcomes of the algorithm. New entities can be straight replacements, like-for-like, of their parent, or be placed in competition against entities in the parent population, or completely replace the parent population. Elements may be removed as a result of selection, or through fitness-independent sampling to maintain a particular population size, or through some end-of-life calculation. The population size limit may act as both upper and lower bound on population size to maintain a specific size, or solely as upper bound.

In Gaucherel and Jensen (2012) the approach is to remove individuals from the population stochastically, with probability related to e to the negative power of the population size multiplied by a configurable parameter, μ. In the "canonical" Evolutionary Computation algorithm, a population limit results from selection where a set number of entities are extracted from the original population, chosen by one of a wide range of selection algorithms (among many sources, see overviews in De Jong (2006, sect. 4.3.1) and Vose (1999, sect. 4.2).) Here though we separate the selection function from the population size limit in order to qualify the effect of the specific limiting mechanism used. To constrain population size without introducing an associated selection bias, in the experiments in this chapter at each generation we remove with uniform probability the exact number of entities (without regard for fitness) to maintain a fixed population size of 5,000 entities.

Translating model parameters into factors in the experiment design results in the first column of table 2.2. As is usual with exploratory experiments with a number of parameters, where each run of the model has some cost in time or other resources, the key problem is to understand the relationship between parameters and response variables at an acceptable cost. In this case, our main cost is time - each run of an evolutionary model is cheap in resources but takes a little time. Exhaustively sampling the entire parameter space is unrealistic. Therefore, we first reduce the search space by limiting the number of values that each parameter can take. By choosing these values appropriately, we can construct an analysis model from the results that is sufficiently accurate for our exploratory purposes at a greatly reduced cost in time.

From the results of earlier parameter selection experiments, we discovered that the model response is sensitive to the values of $p_{reproduction}$ and $p_{selection}$, but has little sensi-

tivity to $n_{children}$.

2.6.3 Does average inheritance approach perfect inheritance?

From the hypothesis the main property of interest is *heritability*, or the correlation between parent and child property values. *Fidelity* therefore is our response variable. Specifically we require $\overline{heritability}_{start}$ and $\overline{heritability}_{end}$, or the mean value for *heritability* (across all replicates) at the beginning and end of each run.

Our null and alternative hypotheses therefore are as follows:

H_0: heritability does not approach 1.0 during a run, irrespective of factor values, or $|\overline{heritability}_{end} - \overline{heritability}_{start}| = 0$

H_1: heritability increases to near 1.0 during a run, for some factor values, or $\overline{heritability}_{end} - \overline{heritability}_{start} > 0$ and $1.0 - \overline{heritability}_{end} < \delta$ for some δ and for some factor values.

Of the 80 experiment runs, all reached the experiment limit of 500 generations.

As we are primarily interested in the final values for heritability, the data in the following analysis is from the final generation of each run, unless otherwise noted.

The first hypothesis prediction is that average heritability will tend towards exact inheritance, or 1.0. The reasoning is as follows: heredity describes the correlation along a lineage, so a high fitness entity with high heredity is likely to have more high fitness descendants than low fitness descendents. Higher fitness entities will survive longer and reproduce more, and so high fitness/high heredity entities will slowly invade the population.

A simple visual inspection of the results in fig. 2.3 reveals that the heritability measure has a strong peak as predicted at 1.0, with final mean heritability greater than approximately 0.96 for all remaining runs. Fitness is bimodal, with distinct groups around final mean fitness values of approximately 0.45–0.55 and 1.0. From figs. 2.4 and 2.5 it seems that the lower group (where final mean fitness does not approach 1.0) is associated with a particular subset of factor-levels. Closer examination shows that these runs, and only these runs, have the level 0.66 for both factors $p_{reproduction}$ and $p_{selection}$. This supports the earlier observation in section 2.6.2 concerning the importance of the distinction between absolute and relative (*e.g.* fitness-based) values for these two factors.

From inspection it seems clear that heritability does approach 1.0 as predicted by the hypothesis. In conclusion, H_0 can be rejected, and H_1 accepted. Inheritance increases regardless of the model design.

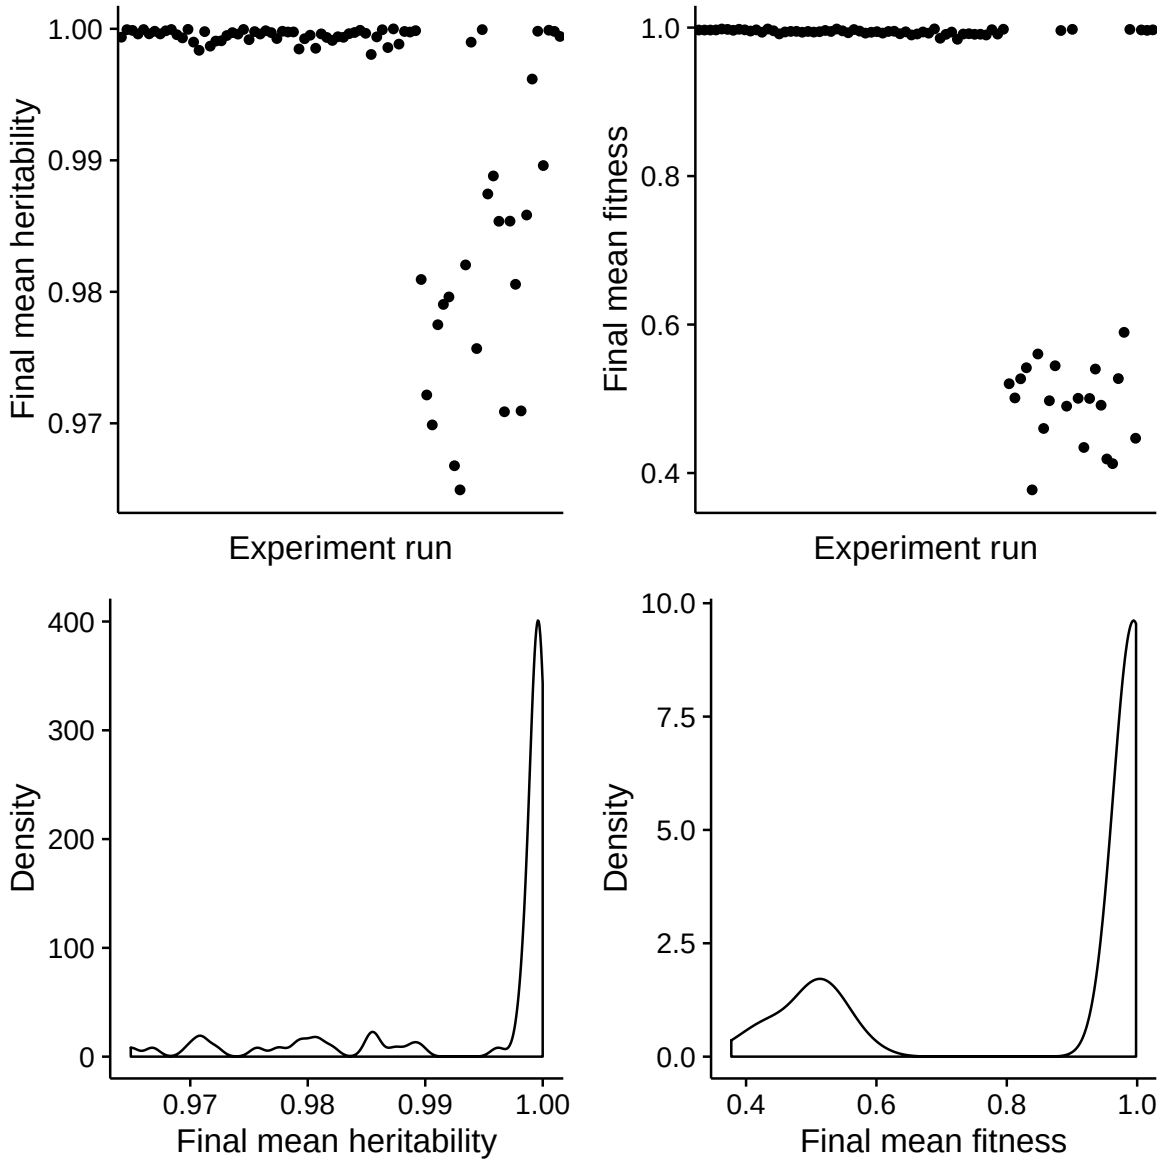

Figure 2.3: Summary results, showing on the left-hand side an overview for the final mean heritability ($\overline{heritability}_{end}$) for each run and a density plot showing the distribution of $\overline{heritability}_{end}$. The right-hand side shows the corresponding plots for final mean fitness ($\overline{fitness}_{end}$)

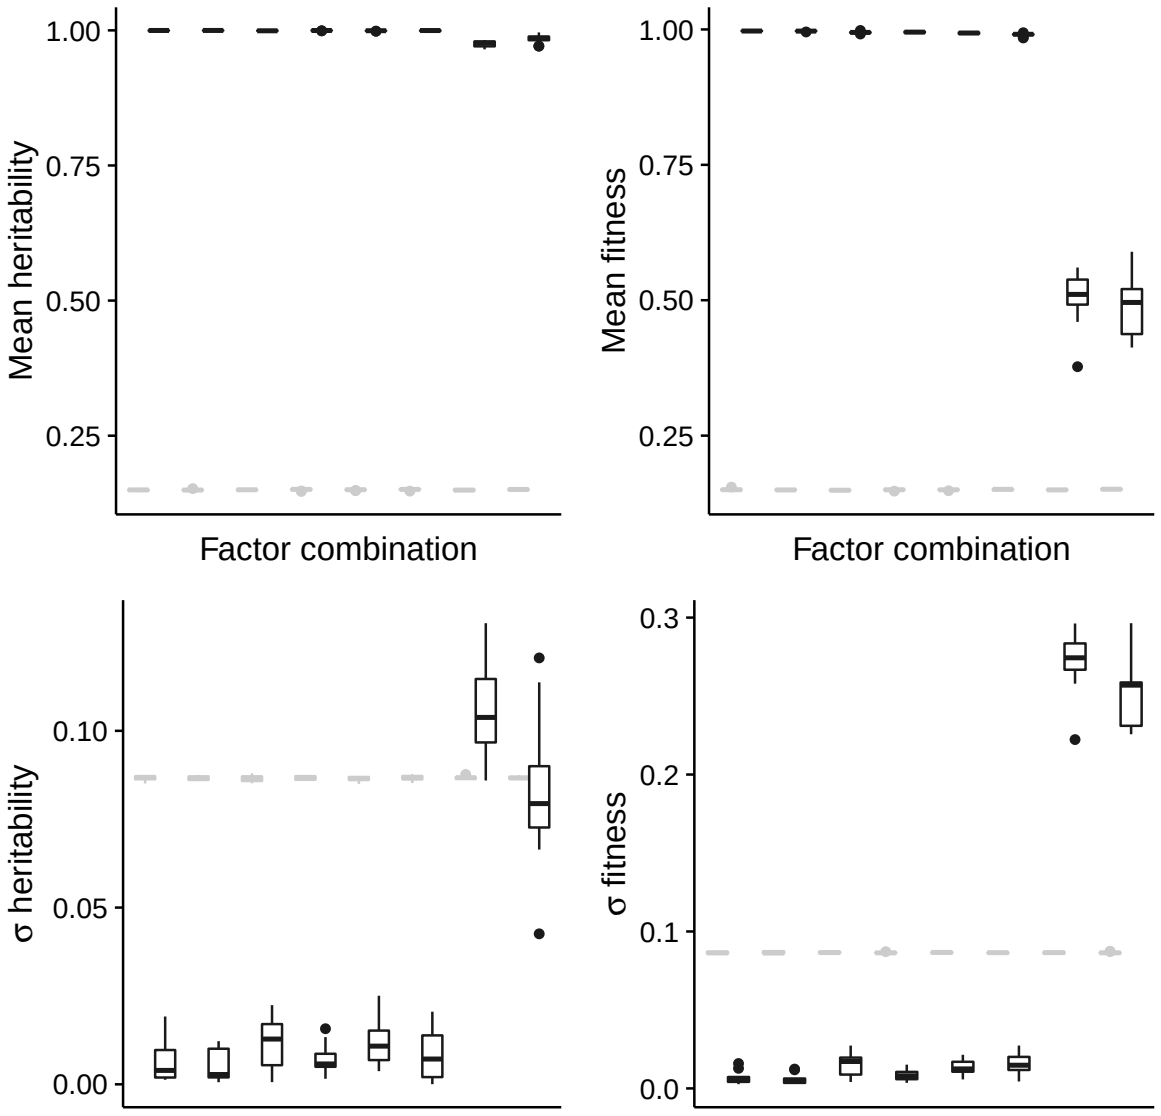

Figure 2.4: Summary results, this time averaged across all replicates of each unique combination of factor levels (that is, grouped by factor level). Dark-coloured points represent values at end of run; light-coloured points in a horizontal line are for initial values, reflecting averaged initial conditions.

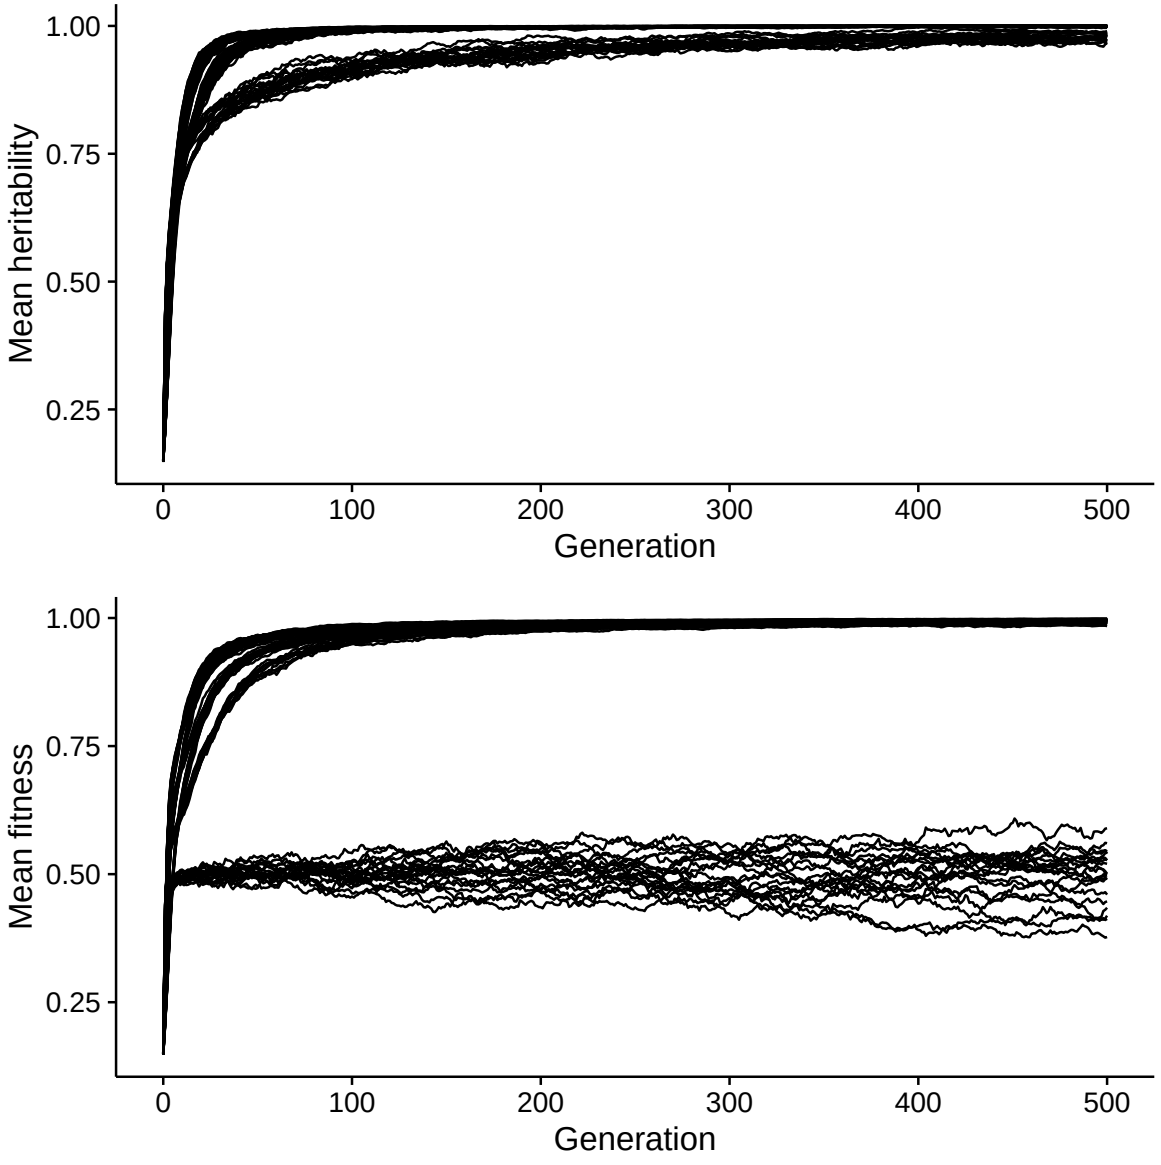

Figure 2.5: Ranges by generation of the mean heritability (top) and mean fitness (bottom) for all levels of all factors. Runs where final mean fitness does not approach 1.0 result from the factor combination $p_{reproduction} = 0.66$ and $p_{selection} = 0.66$ (see text.)

2.6.4 Does the standard deviation of inheritance decrease over time in the population?

The second prediction of the hypothesis is that the population variance for inheritance ($\sigma_{heritability}$) should decrease over time towards a limit of 0 in fixed environments.

H_0: $\sigma_{heritability_{end}} - \sigma_{heritability_{start}} >= 0$, for all factor values.

H_1: $\sigma_{heritability_{end}} - \sigma_{heritability_{start}} < 0$, for some factor values.

Once again, the main property of interest, and so our response variable, is *heritability*. From the hypothesis we require $\sigma_{heritability_{end}}$ and $\sigma_{heritability_{start}}$, the standard deviation of heritability at the beginning of each run, and at the end.

Data is all generations for all 80 fixed-environment runs that reached completion at generation 500.

In a similar fashion to the procedure in section 2.6.3, by visual inspection of fig. 2.6, the variance (square of the standard deviation) does decrease towards zero in all cases. Once again, the runs appear to fall into two distinct groups based on the speed of convergence towards zero, as seen in fig. 2.6. The upper group is exclusively associated with runs where both factors $p_{reproduction}$ and $p_{selection}$ are set to 0.66; these levels are also responsible for the uppermost band of final standard deviation of fitness in the lower section of fig. 2.6, where the standard deviation of fitness remains essentially unchanged throughout the run.

In conclusion, from visual inspection, the variance of heritability does approach 0, for all combinations of factor levels except where both $p_{reproduction}$ and $p_{selection}$ are set to 0.66. Therefore H_0 is rejected and H_1 accepted.

2.7 Hypothesis test under changing environments

In this section we describe an appropriate experimental design and model to describe environmental changes; test the predictions of section 2.4 using the model; and finally discuss the results of the experiment.

2.7.1 Environmental model

Alongside the evolutionary model from section 2.5 we now add an environmental model to describe the changing fitness relationship between entities and the environment at each step of the evolutionary model. This relationship is simply modelled by constructing the entity fitness at timestep $t + 1$ as the entity's fitness at timestep t plus the environmental contribution at timestep t, bound to the range for fitness of $[0, 1]$.

Statistical techniques are commonplace for time-series predictions (Brockwell and Davis 2002), where a *time series* is a set of observations x_i, each one being recorded at a specific

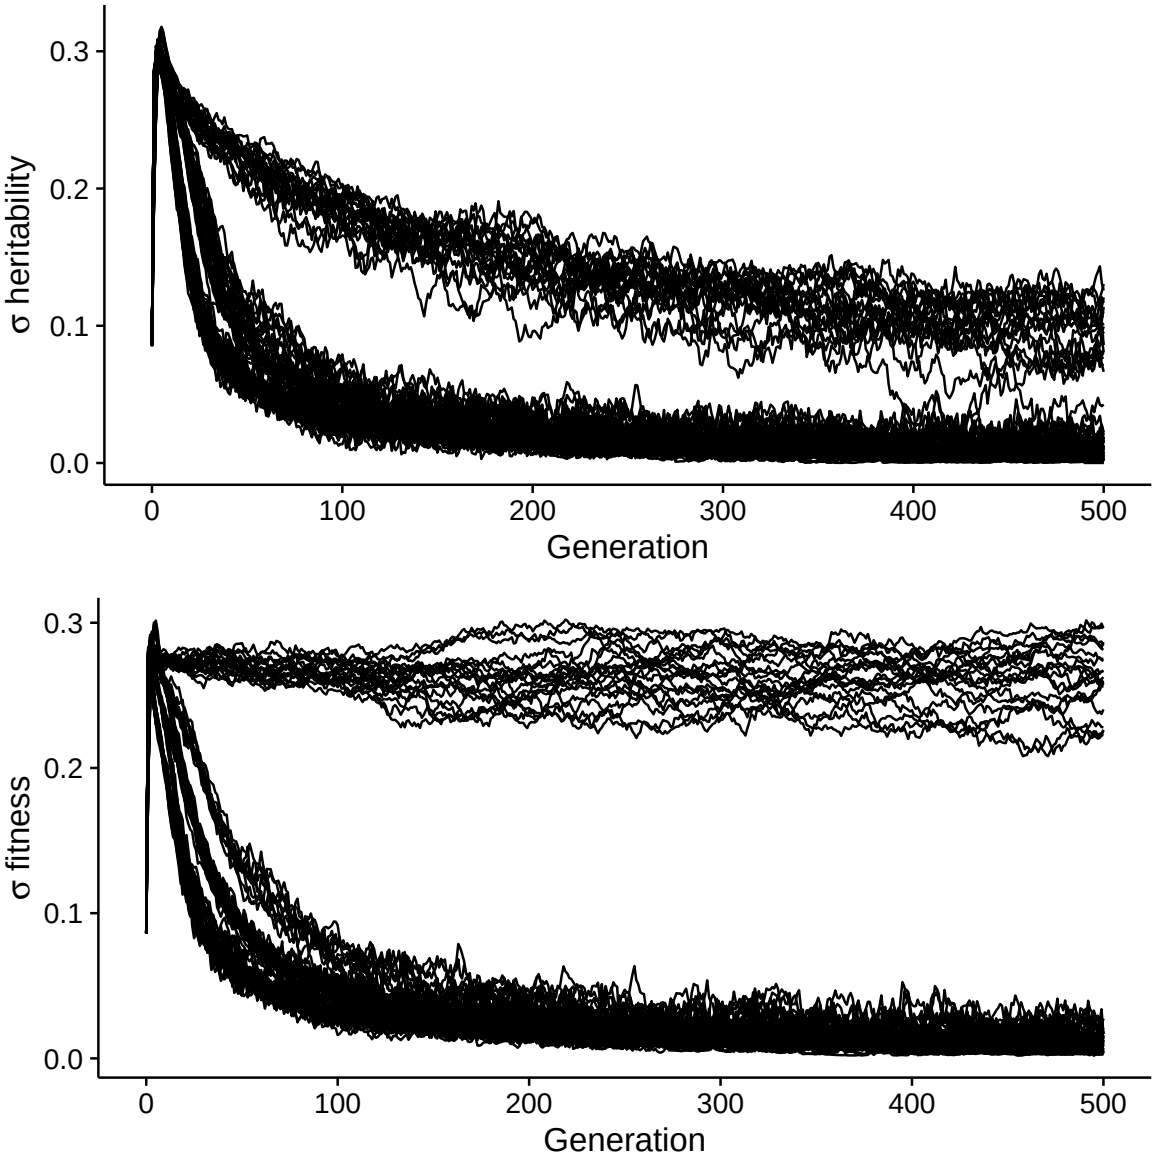

Figure 2.6: Ranges by generation of the standard deviation of heritability on the top row, and standard deviation of fitness below. Runs where standard deviation of fitness remains essentially unchanged (upper band in lower plot) result from the factor combination $p_{reproduction} = 0.66$ and $p_{selection} = 0.66$ (see text.)

time t where an observation $x_i \in$ some set $\{X\}$, assumed to be $\mathbb{R}$ (Brockwell and Davis 2002).

Although time-series modelling provides techniques for describing time-series data in terms of an underlying model, the process can also be reversed to produce a time-series *from* the model; in other words, if the variety of environmental change required to test our hypothesis can be described by a standard time-series model, the parameters that determine that model can also serve as our headline measure for environmental change.

There are various forms of standard time-series models. ARMA models, combining autoregressive (AR) and moving average (MA) components, are often used for *stationary* series, that is time-series whose joint probability distribution (and hence whose statistical properties such as mean and variance) do not change over time. This is equivalent to saying the series does not demonstrate concept drift. By contrast, ARIMA models apply for non-stationary series where the difference between two sequential values of the original series can be shown to produce a stationary series—the I or "Integrated" component of the model provides the differencing.

Our environmental model provides an enhanced AR(1) or first-order autoregressive time-series, with each timestep corresponding to one evolutionary generation. Specifically, we describe the evolutionary change at each timestep as a function of the previous timestep:

$x_t = \Theta x_{t-1} + e_t + \delta$

where x_t is the change at timestep t, Θ is the AR coefficient, e_t is a random, normally distributed, error component around a mean of 0, where $e_t \overset{iid}{\sim} N(0, \sigma_e^2)$, and δ is a fixed bias value.

This series allows us to represent a broad range of environmental changes:

- Each time-series is completely specified by three parameters, Θ, σ_e and δ.

- Θ in an autoregressive time-series can be interpreted as specifying stability or smoothness, while δ is a fixed change. We use δ to model a fitness trend - environments with a positive δ will see the fitness of each entity improved at each generation, with the opposite of course true of negative δ. Note that with this formulation we can model linear trends in fitness from a fixed bias in the environment produced by the δ term. This is not the same as a ARI model where the deltas of the environment time-series itself would follow a trend.

- An AR time-series has the property of stationarity, with the implication that the mean of the series is constant through time. However, as we apply the series values as deltas to element fitness, fitness can be non-stationary, and so may show a long term trend. This allows a simple non-differencing time-series to describe a steady change in fitness.

- As a corollary of stationarity, the range of the series is determined by the initial parameters. This is a useful property as it means that with appropriate parameter choices no scaling of the range is required.
- The time-series is defined by three independent elements, two predictable (driven by Θ and δ) and the other (σ_e) random and unlearnable. By changing the ratio between the predictable and unpredictable we can examine the performance of the evolutionary algorithm on some continuum of predictability.

2.7.2 Experiment design

Section 2.4 allows us to make two specific predictions for changing environments:

1. Heritability is proportional to the predictability of the environment, that is at a minimum in conditions of maximum unpredictability, and at a maximum in stable conditions.
2. Variation in heritability, $\sigma_{heritability}$, is proportional to the degree of environmental variability.. As a corollary, $\sigma_{heritability}$ under changing conditions will be greater than that under stable conditions.

As all three independent variables for our experiment, Θ, σ_e and δ, are continuous in $\mathbb{R}$, and as we wish to test the specific relationship of heritability across a range of these variables without being restricted to the initial choice of fixed levels, we shift from the fixed-effect factorial designs used earlier to a random effects model (Montgomery 2009, chap.13). The independent variables in each run are the parameters to the environmental model, with values taken from a uniform random sample from their range. Or in other words, a series of random samples with uniform probability from a cube formed by a parameter on each of the three axes (see fig. 2.11.)

We create a set of independent datasets by sampling from four different parameter sets, as defined in table 2.3. The extremes of the ranges see a lower proportion of the runs (representing more significant environmental changes or more rapid changes) reach completion than was seen in the centre (where the environment was stable.) The range of each variable is adjusted from dataset to dataset to balance an adequate density of coverage across the overall range against the total run time.

The population consists of 5,000 entities, each with fitness and heritability chosen independently and with uniform probability from the range [0,0.3], and each run is for 500 generations. The other factors and levels along with their mapping to model parameters, are shown below:

- $p_{reproduction}$ = fitness.

Table 2.3: Range of independent variables Θ, σ_e and δ.

Dataset	Θ	σ_e	δ
Dataset no.1	$[-0.4, 0.4]$	$[0, 0.2]$	$[-0.1, 0.1]$
Dataset no.2	$[-0.2, 0.2]$	$[0, 0.1]$	$[-0.05, 0.05]$
Dataset no.3	$[-0.2, 0.2]$	$[0, 0.1]$	$[-0.05, 0.05]$
Dataset no.4	$[-0.4, 0.4]$	$[0, 0.4]$	$[-0.04, 0.04]$

Table 2.4: Summary of results for changing environments.

Dataset	n (total runs)	Completed runs
Dataset no.1	1000	700
Dataset no.2	400	399
Dataset no.3	2000	1972
Dataset no.4	400	151

- $p_{selection}$ = fitness.
- $n_{children} = 2$.

Finally, the dependent or response variables, driven by the hypothesis predictions, are mean population heritability and mean population fitness.

2.7.3 Results

A summary of the results from each of the four datasets is given in table 2.4, and a visualisation of all four datasets combined given in fig. 2.12. Note that the difference in the proportion of completed runs to total runs reflects the varying ranges of the independent variables in each dataset (as seen in table 2.3), shown graphically in fig. 2.13.

2.7.4 Is heritability proportional to the predictability of the environment?

Our null and alternate hypotheses are:

H_0: $\overline{heritability}_i = \overline{heritability}_j \forall i, j$.

H_1: $\overline{heritability}_i \neq \overline{heritability}_j$ for some i, j.

Remembering that a function of Θ, σ_e and δ provides a measure of environmental predictability, and because only Θ and δ are predictable, and hence learnable, we would not expect σ_e to be significant in determining predictability.

By visual inspection of fig. 2.12, $\overline{heritability}_{end}$ is related to $f(\delta)$; this is partially supported by the results of an ANOVA analysis, which finds both δ and Θ to be highly significant (p<0.001). Therefore $\overline{heritability}_{end}$ is proportional to either $f(\delta)$ or $f(\Theta, \delta)$ (and incidentally meets our expectation that σ_e is not a potential parameter to $f()$ as it is not learnable by an evolutionary system.)

Table 2.5: ANOVA analysis for H_1: $\overline{heritability}_{end}$ is proportional to $f(\Theta, \sigma_e, \delta)$.

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
sigma	1	0.00	0.00	0.00	0.9614
delta	1	16.83	16.83	5379.78	0.0000
theta	1	0.05	0.05	16.30	0.0001
Residuals	3214	10.05	0.00		

As a result, we can reject H_0 and accept H_1: heritability is proportional to the predictability of the environment (as measured by some function $f(\Theta, \delta)$.)

2.7.5 Does the standard deviation of heritability vary in proportion to the variability in the environment?

H_0: $\sigma_{heritability_{end}}$ is not proportional to $f(\Theta, \sigma_e, \delta)$ for all functions $f()$.

H_1: $\sigma_{heritability_{end}}$ is proportional to $f(\Theta, \sigma_e, \delta)$ for some function $f()$.

An ANOVA model evaluating a linear relationship between the standard deviation of heritability, and the parameters δ, σ_e and Θ to the environmental model suggests that all three parameters have some effect on the $\sigma_{heritability}$, with the influence of δ and Θ being highly significant (p<0.001) and that of σ_e significant (p<0.01). Unlike in the previous section, here we are interested in the variability, or unpredictability, of the environment. As σ_e is fundamentally unpredictable, it is not unexpected that it contributes to the relationship between the unpredictability of the environment and the variance of heritability.

Table 2.6: ANOVA analysis for H_1: $\sigma_{heritability_{end}}$ is proportional to $f(\Theta, \sigma_e, \delta)$.

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
sigma	1	0.02	0.02	6.64	0.0100
delta	1	21.60	21.60	7963.79	0.0000
theta	1	0.06	0.06	20.49	0.0000
Residuals	3214	8.72	0.00		

The corollary of this prediction is that $\sigma_{heritability}$ will be greater in unpredictable environments than in stable environments (when all three parameters are near zero). As a first approach, we note that in section 2.6.4 we showed that in stable environments the variance

of inheritance tended towards zero. As σ_e is by definition unpredictable (see section 2.7.1) we can provide some support for the corollary if we can show that variance remains above zero for all $\sigma_e > 0$ (incorporating all necessary uncertainties):

H_0: $\sigma_{heritability_{end}}$ is close to 0, for all clearly non-zero combinations of Θ,δ and σ_e.

H_1: $\sigma_{heritability_{end}}$ is not close to 0, for some clearly non-zero combinations of Θ,δ and σ_e.

From a visual inspection of fig. 2.14 we can reject H_0 in favour of H_1: $\sigma_{heritability_{end}}$ is greater under unpredictable conditions than it is under stable environmental conditions.

2.8 Conclusions

In this chapter we have shown by experiment that, as predicted, in stable environments:

- Heritability increases towards an upper limit of 1.0, or perfect inheritance.
- The variance of heritability decreases towards a lower limit of 0.

This is encouraging, and the first point confirms the main result from the exploration by Bourrat (2015). However, from the viewpoint of a creative open-ended evolutionary process, this is not in fact what is desired: perfect inheritance means an absence of novelty.

In section 2.7 we tested the hypothesis predictions from section 2.4 for changing environments. First we defined what we mean by environmental change. Our key assumption is that this change may be modelled as a timeseries, specifically one that is autoregressive (AR). In section 2.7.1 we described an environmental model based on a AR(1) timeseries defined by three independent parameters, Θ, δ and σ_e. Specifically, the change in fitness at each timestep t is given by the function $\Theta x_{t-1} + e_t + \delta$, where e_t is an error term related to σ_e by $e_t \overset{iid}{\sim} N(0, \sigma_e^2)$. From this description it is clear that of these three parameters, two—Θ and δ—are more discoverable by an evolutionary learner than the other, σ_e.

We assume that all entities experience the same relationship to the environment: environmental changes, through the proxy of fitness changes, apply to all equally. Although we experimented with modelling fitness change by lineage—the chain of entities that trace ancestry back to some shared ancestor—the assumption that lineages are homogeneous seems as arbitrary as the assumption that all entities respond alike.

Our model is equally blind to co-evolution, as the only relationship an entity has is with the environment, and not directly with other entities. As co-evolutionary phenomenon in biology though are primarily inter-species, and our model is effectively single-species, co-evolution is outside the current model's scope.

Section 2.4 made these predictions:

1. Heritability is proportional to the predictability of the environment: at a minimum in conditions of maximum unpredictability, and at a maximum in stable conditions.

2. Variation in heritability, $\sigma_{heritability}$, is proportional to the degree of environmental variability. As a corollary, the $\sigma_{heritability}$ under changing conditions will be greater than that under stable conditions.

In section 2.7.4 we constructed a random-effects factorial experiment where the simulation model from section 2.5 was combined with the environmental model to confirm the first prediction: heritability in changing environments is indeed related to environmental predictability. However, this conclusion would be strengthened by an ordering based on Θ, δ and σ_e, as adopted for in the experiments in chapter 6.

The second prediction was also confirmed in similar fashion in section 2.7.5, with a similar caveat. Although we can show a clear difference between stable and changing environments, without an absolute ordering over Θ, δ and σ_e it is difficult to show a proportional relationship. Furthermore, our model is primarily concerned with the connection between predictability and heredity, and is essentially silent on unpredictable change such as experienced during times of evolutionary transitions. However, to go from limited replication to unlimited heredity replicators, it is our belief that we require novelty, rather than steady improvement. Modelling the unpredictable is however beyond the scope of the abstract model described in this chapter, and is better explored by the full simulation model described in later chapters.

Within these bounds though this chapter provides support for our initial working hypothesis from RQ1 that an effective inheritance mechanism can not only emerge from a low-heredity environment through evolution, but that it can be tuned and optimised by evolution. Environmental variation plays a key role, and this insight forms the basis for the work described in the remainder of this thesis.

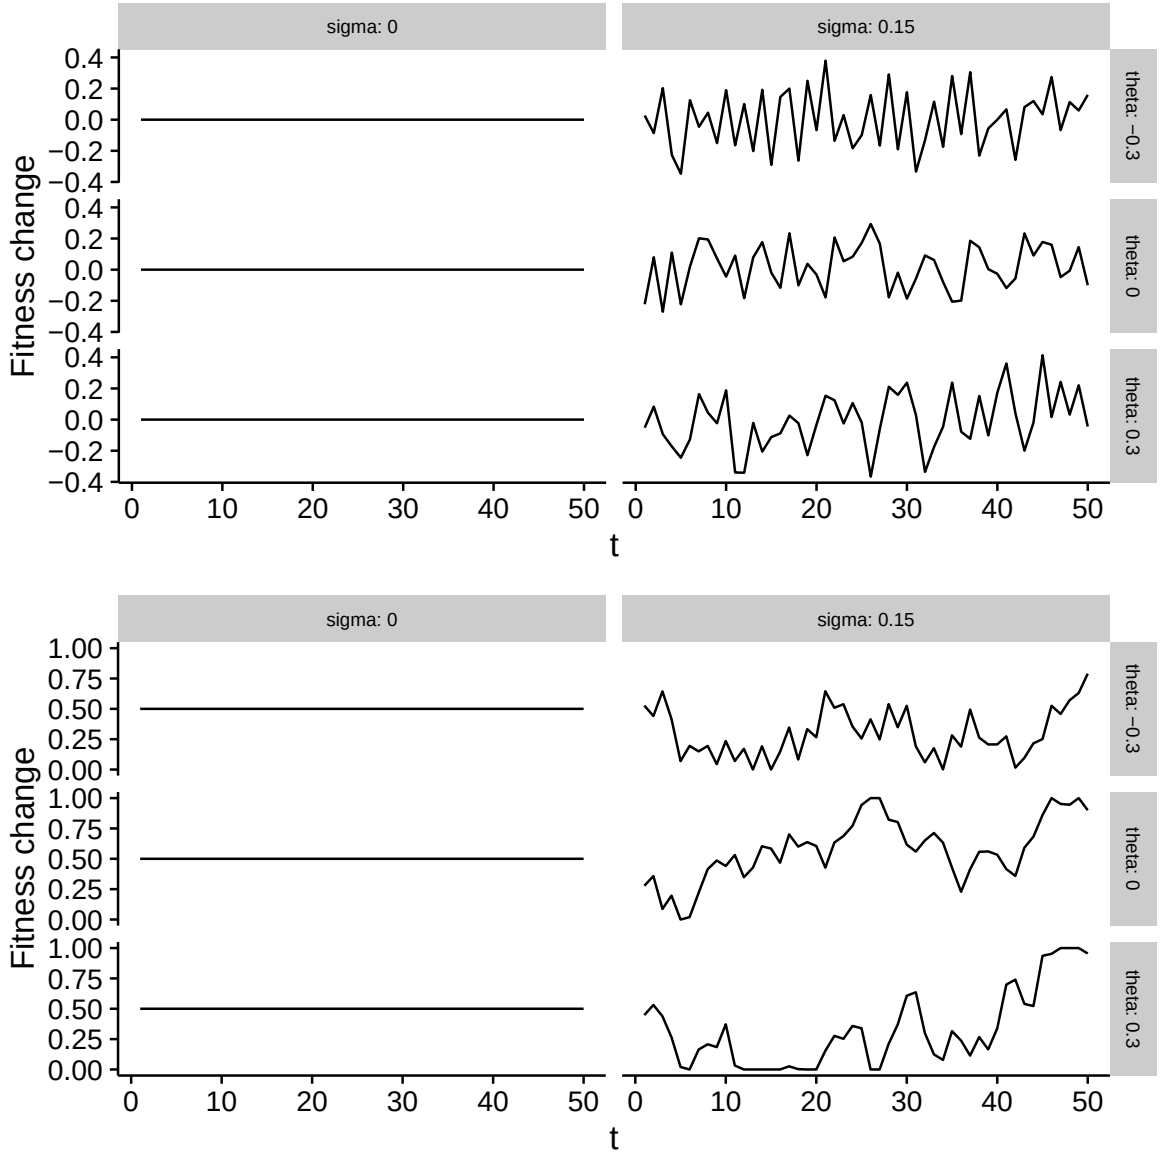

Figure 2.7: Visualisation of the fitness changes at each time interval (top facet) and cumulative fitness change (bottom facet) that result from our environmental model for delta=0 with two values of sigma (left and right), and three values of theta (each facet row.)

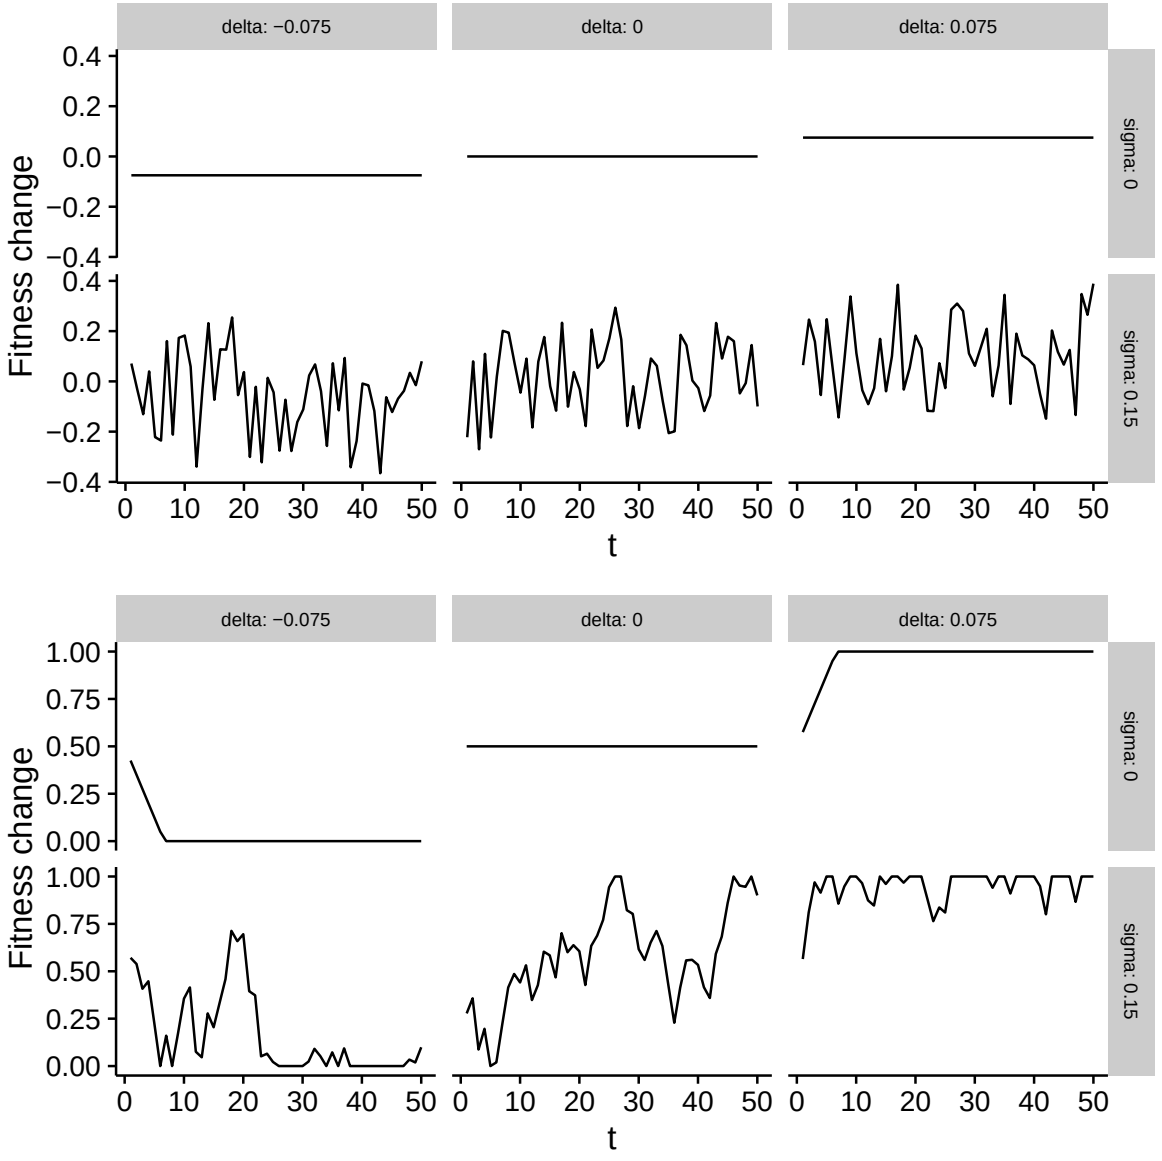

Figure 2.8: Visualisation of the fitness changes at each time interval (top facet) and cumulative fitness change (bottom facet) for theta=0 with three values of delta (columns), and two values of sigma (each facet row.)

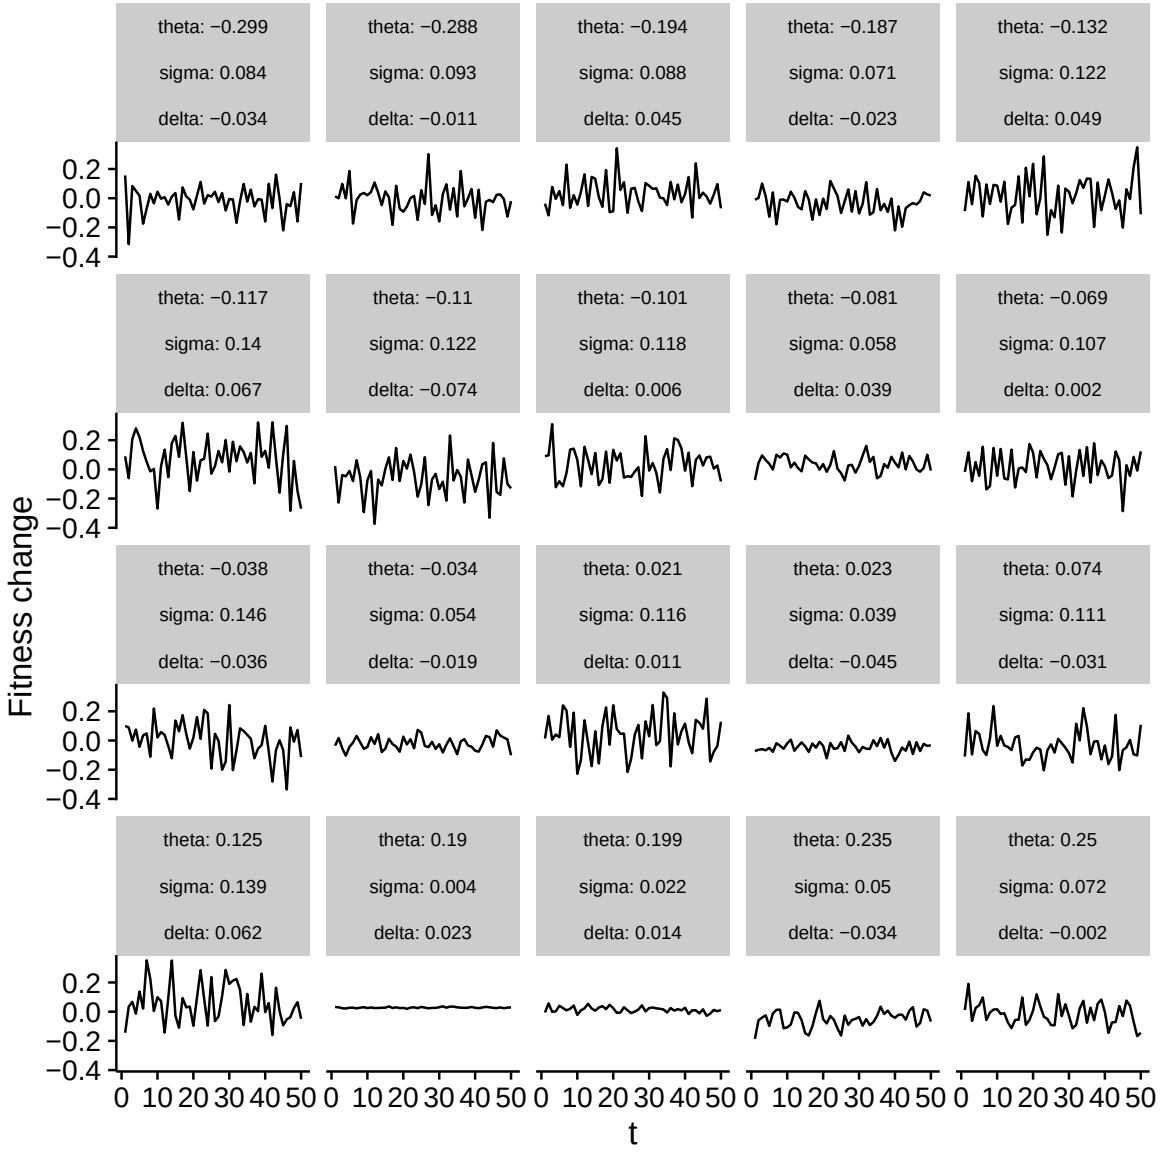

Figure 2.9: Example time-series produced by our environmental model for a sample of theta, sigma and delta values. These are incremental fitness changes, rather than cumulative ones.

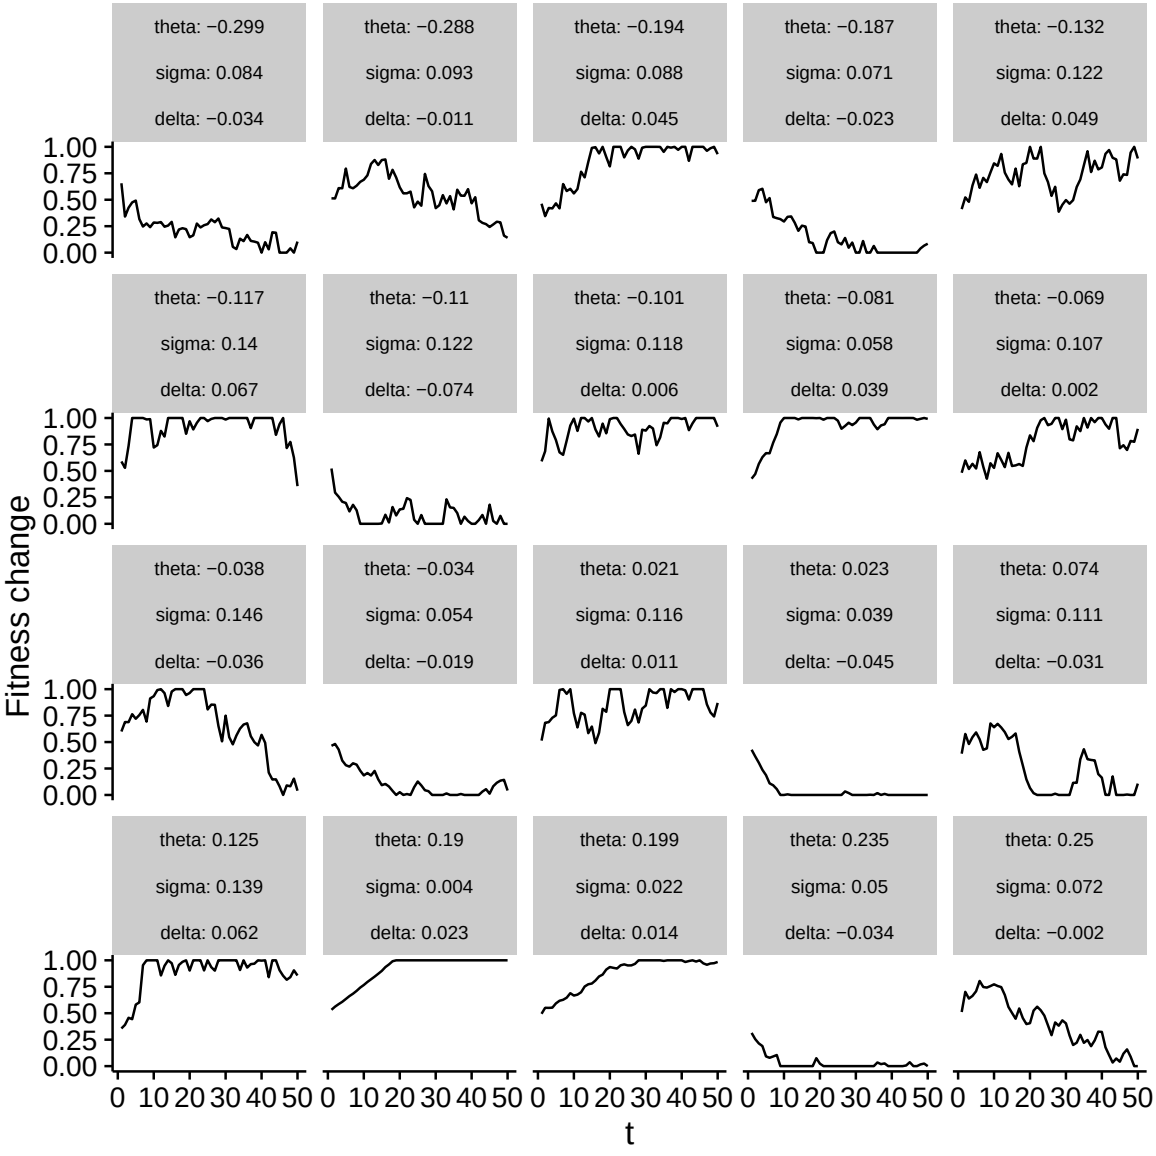

Figure 2.10: Cumulative time-series examples for samples of theta, sigma and delta.

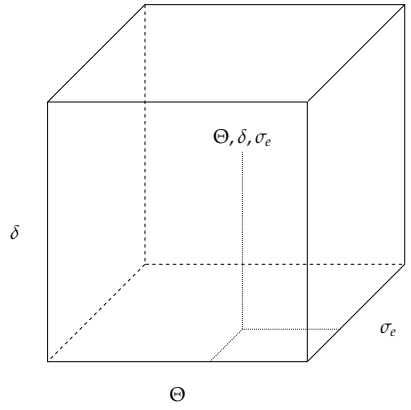

Figure 2.11: The parameter space formed by the ranges of the three parameters to the environmental model, Θ, σ_e and δ.

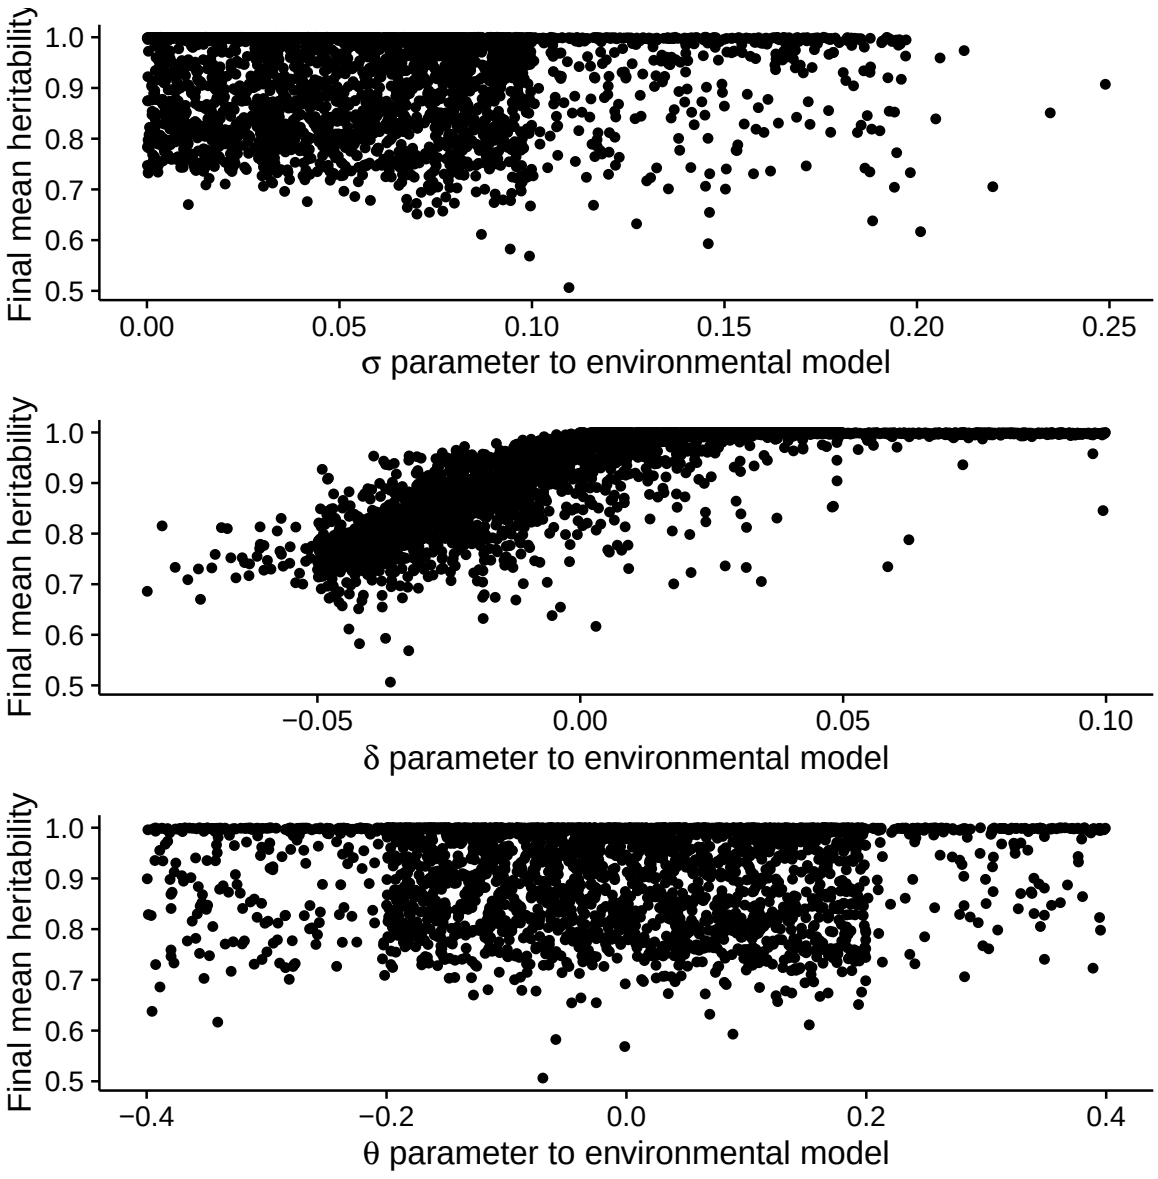

Figure 2.12: Final mean heritability against the environmental model's sigma parameter (top), delta parameter (middle) and theta parameter (bottom) for all experimental runs.

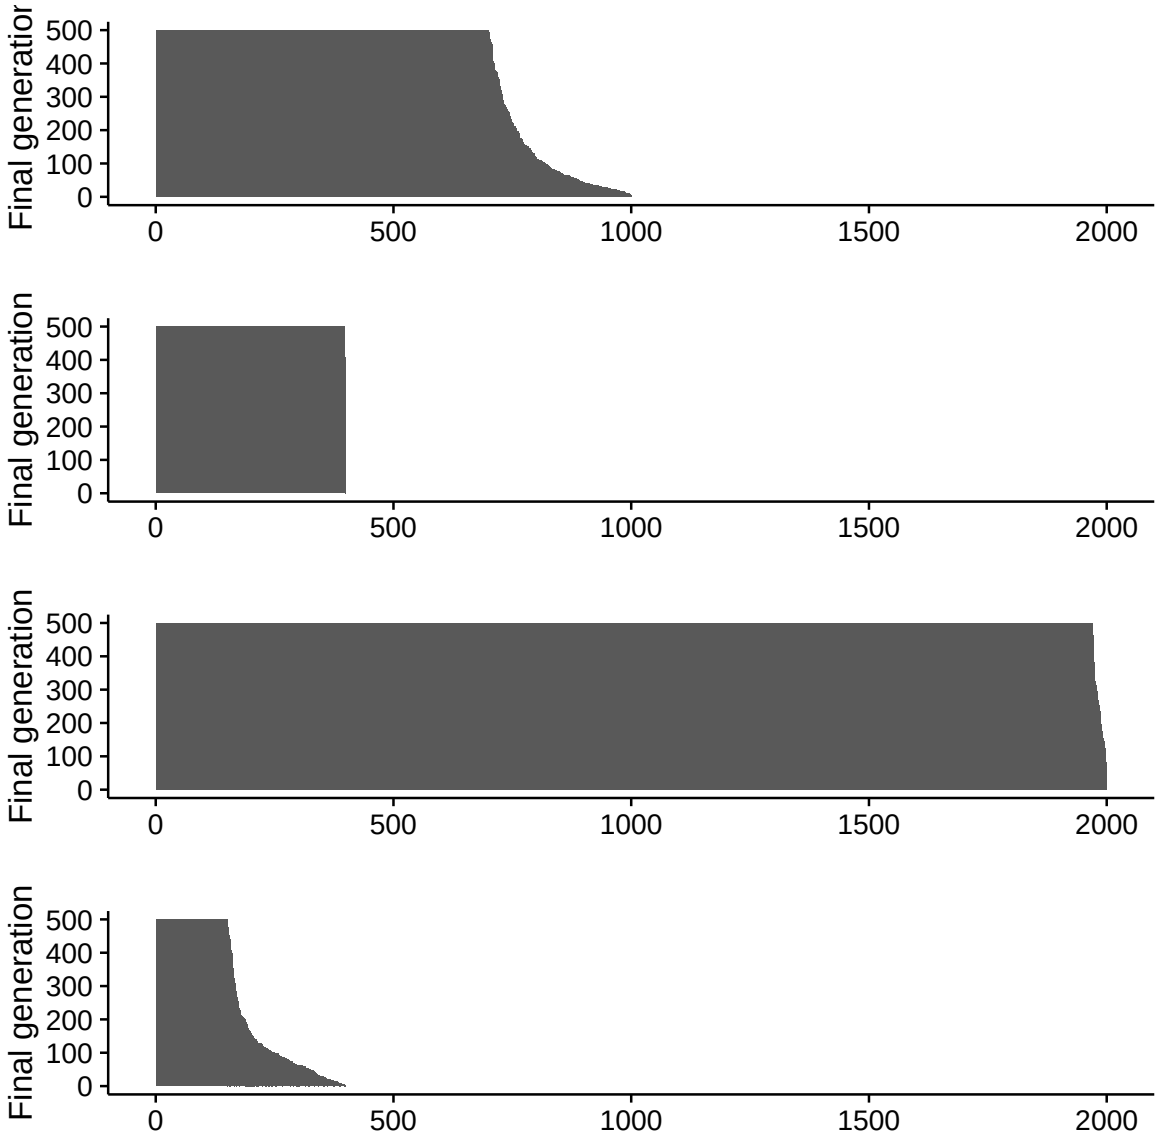

Figure 2.13: Distribution of maximum generation reached for all runs for dataset no.1 (top) through dataset no.4 (bottom). Axes are to the same scale. Runs (along the x-axis) are ordered by final generation reached.

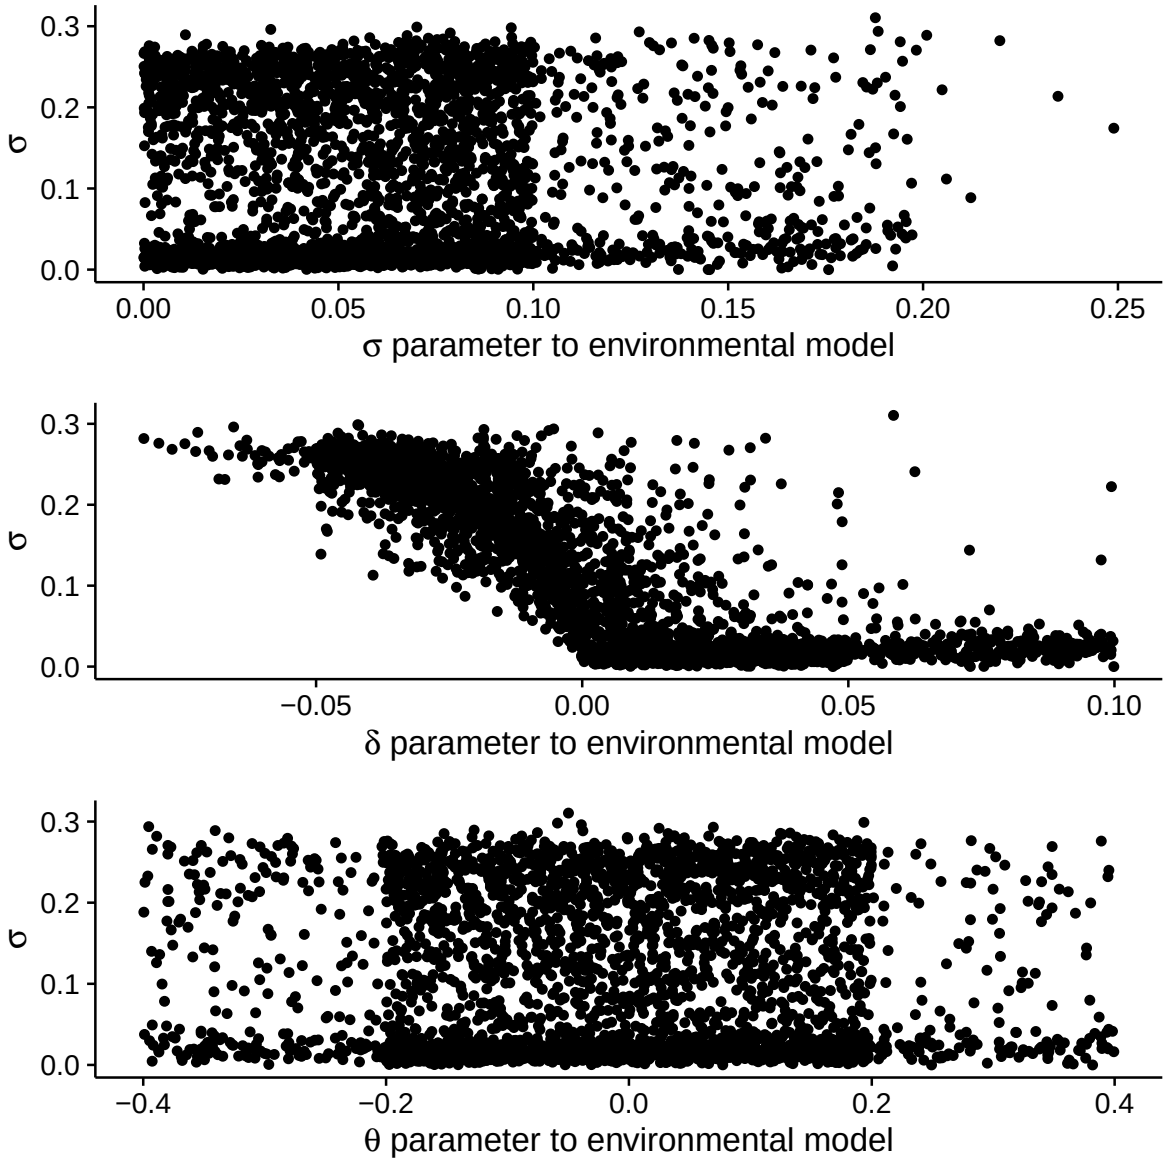

Figure 2.14: Final standard deviation of heritability against the environmental model's sigma parameter (top), delta parameter (middle) and theta parameter (bottom) for all experimental runs.

3 Introduction to Artificial Chemistries

At this point we have essentially completed the investigation of our first research question, and now move on to the second: the emergence of replicators in an artificial chemistry from the interaction of molecules through reactions. Artificial Chemistries provide an interesting testbed for investigating various evolutionary phenomena. This review chapter provides background material to aid in understanding how artificial chemistries are almost uniquely suited to the exploration of our research questions.

Fundamentally, artificial chemistries provide a tunable system, capable of highly complex behaviour, built around familiar metaphors (in many cases real-world chemistry, and potentially biology). In a molecular artificial chemistry, a set of rules describing how atoms interact gives rise to emergent forms: molecules. At a higher level, these molecules, under the same interaction rules, also interact in patterns (reactions.)

Still higher emergent levels emerge under favourable conditions. Reactions may form cycles, where a sequence of reactions eventually returns to an earlier state. Cycles in particular are interesting as many biological processes are cyclical[1]. Replication, resulting in an exact copy of an entity, is a macro-example of a cycle; metabolism is another. Building on the apparent correspondence between higher emergent levels in artificial chemistry evolution and biology, we believe like others (*e.g.* Steel, Hordijk, and Smith (2013)) that cycles, of some form, are a necessary building-block for more complicated structures in artificial chemistries.

Artificial chemistries are regularly employed in three application areas: real-world chemistry simulators; tools for the exploration of artificial life, and models to test various hypotheses of the origin of life. Chemistry emulators and origins-of-life tools aim for fidelity with real-world chemistry, unlike most artificial life models. Real-world fidelity requires either the use of a library of predefined reactions, which conflicts with our goal of unlimited extension, or a chemically plausible method of constructing reactions from first principles. Because of the complexities of real-world chemistry this later method ap-

[1] In the broadest sense life can be seen as an autocatalytic process where an entity catalyses the production of one or more descendant entities.

pears to be quite difficult, and the goal of a realistic, computationally practical, artificial chemistry remains open.

However, the more limited objective of a less realistic, but still unbounded chemistry, has been achieved (see section 3.2 for examples.) One advantage of semi-realistic artificial chemistries, as pointed out by Funes (2001, p. 5), is that it is easier to evaluate solutions in a domain close to the real world as opposed to a purely symbolic or abstract domain such as a lambda-calculus, or a programmatic environment like Tierra. When contemplating difficult problems such as complexity our intuition can be helpful, but only in situations close enough to our normal experience for it to be relevant. Artificial chemistries are a particular type of model, familiar from real world chemistry, for the simulation of reaction-based systems.

A mathematical treatment of artificial chemistries can be found in Benkö et al. (2009). Dittrich, Ziegler, and Banzhaf (2001) and Suzuki and Dittrich (2008) provide excellent reviews of the field, while an influential taxonomy is given in Dittrich, Ziegler, and Banzhaf (2001), described in section 3.2. Every artificial chemistry defines a set of constitutive elements (molecules), a set of transformation rules (reactions) and a mechanism for choosing and ordering the sequence of reactions (a reactor algorithm) (Dittrich, Ziegler, and Banzhaf 2001). These components and taxonomies are described in more detail in the next section, but first we digress to provide a short review of chemistry in the real, rather than artificial, world.

3.1 An introduction to real-world chemistry

Artificial chemistries are not necessarily models of real-world physical chemistry but most artificial chemistries borrow, at least, from the language of real-world chemistry. An understanding of the key concepts in physical chemistry is therefore useful in appreciating artificial chemistries.

The richness of the physical world is built on a substantial foundation of chemical complexity. A *reaction* transforms *reactants* into *products*, and is often described by the quantities (or stoichiometry) of the reactants and resulting product molecules. The reaction can also be characterised by a description of the dynamics, or kinetics, of the reaction to explain how the reaction proceeds in response to temperature changes or to varying concentrations of reactants. Reactions often require an input of energy, such as from molecular collisions, to proceed; the energy required is called the reaction's activation energy (E_a), and is specific to the particular reaction. In general there is no accurate mechanism to predict reaction dynamics without experiment.

The *reaction rate* is the change in concentration of a substance over time: $\frac{-d[A]}{dt} = k[A][B]$ (where $[X]$ means the concentration of X, and k is the reaction rate constant) leading to

Arrenhius' description of the relationship between the activation energy (E_a), the temperature (T) and the rate constant (k): $k = Ae^{E_a/RT}$ (R is the universal gas constant.) The *reaction order* describes how the reaction rate changes with the concentration of the reactants, usually captured as a first-, second-, or third-degree polynomial expression determined empirically. For example, the reaction rate equation for $2NO + Cl_2 \rightarrow 2NOCl$ is $rate = k[NO]^2[Cl_2]$ (experimentally determined), and is second-order with respect to NO, first-order with Cl_2 and overall (the sum of terms), third-order (example taken from Kotz, Treichel, and Weaver (2006).)

Reactions can generally be decomposed to a chain of *elementary steps*, with each step resulting in a single change, such as bond formation or cleavage, to the reacting molecules. Elementary steps are somewhat predictable, practical reactions somewhat less so. Generally in experimental chemistry we know the reactants and products and can sometimes deduce the sequence of elementary steps.

In modelling, we can either represent reactions exactly as an atomic transformation from reactants to products with characteristics that can only be determined experimentally, or we can attempt to construct the reaction from a sequence of elementary steps with the properties of the reaction derived from the properties of the steps involved. This later approach is the only practical one when the reaction is novel, or when we lack experimental data. Several alternate sequences of steps, or *reaction pathways*, may be possible between reactants and products. Each pathway will have a different activation energy, and hence reaction rate.

The kinetics of elementary steps are defined by the *stoichiometry*. In theory you might expect to be able to predict the overall reactions kinetics from the composition of these elementary steps. In practice, the results are close, but not exact, when compared to experiment. However they form a useful abstraction for analysis. The reaction rate of an elementary step is defined by the stoichiometry, where the rate equation is the product of the reactant concentrations and the rate constant. Therefore a step with one reactant has order 1, a bimolecular step ($A + B$) has order 2 and so on. The step $2A + B$ has rate equation $k[A][A][B] = k[A]^2[B]$ and is of order 3.

Elementary steps are an abstraction, an aid to analysis, of the underlying molecular dynamics. At the molecular level, reactions can be modelled as a series of collisions between molecules. The reaction rate is then determined by the percentage of collisions (or in other words, the concentration) that are energetic enough to overcome the inherent stability of the interacting molecules and cause a change in molecular structure or shape (in other words, the activation energy).

For an entity to emerge from a network of reactions, the core problem is how to privilege the reactions that are useful in producing products (ones that can enter into other necessary reactions), over the ones that are not. Given a collection of molecules, the set

of all possible reactions is determined by the chemistry that governs how molecules interact to form other molecules. The reactions that actually take place though are driven by molecular concentrations and kinetics. Molecules increase in concentration as they are produced in reactions, and become scarcer as they are consumed as reactants. More common molecules are more likely to be chosen as reactants, and less common ones less likely. This can also affect the direction of a reaction: for reactions which may run in either direction, the relative concentrations of molecules on one side of the reaction versus those on the other side determines the direction in which the reaction is most likely to run.

Recall that the rate of a reaction is a function of concentration (at the gross level) or collision rate (at the molecular level) and the activation energy for the reaction, which at the molecular level is directly related to the energy required to overcome the stability of the reactants. There are therefore two clear mechanisms to alter the rate for a reaction: either reduce activation energy, or increase concentrations. A catalyst does precisely the first, and autocatalysis is one method for the second.

A *catalyst* is anything that isn't consumed by the reaction and that affects the rate (or the kinetic equation for the reaction) without affecting the reaction's equilibrium constant. Catalysts allow the reaction to proceed by an alternate, lower activation energy, pathway. The reaction equation remains the same, but the dynamics are changed. In the case of biological enzymes the rate can be increased by several orders of magnitude over the uncatalysed reaction, enabling reactions fundamental to life that would be effectively impossible in an uncatalysed form. At the molecular level catalysts often function by providing a substrate that preferentially attacks the bonds critical to the reaction. The platinum within an automotive catalytic converter is a well-known example of a catalytic substrate. Catalysts are often show above the reaction arrow in standard chemical reaction notation, *e.g.* $A \xrightarrow{catalyst} B$.

In *autocatalysis*, introduced by Ostwald (1890), rather than reducing the required activation energy, autocatalysis increases the reactant concentrations: autocatalysis reactions form feedback loops where a compound is both a reactant and a product. In the standard definition, an autocatalytic reaction is one that is catalysed by its own products, resulting in a characteristic rate acceleration over time given by the differential equation

$$\frac{dx_i}{dt} = k(\mathbf{X}) \cdot x_i^n + f(\mathbf{X})$$

where n is the order of the reaction, and $f(\mathbf{X})$ the contribution from all other elements of the system (Plasson et al. 2010). As an example, in the indirect network autocatalysis of glycolysis where the pattern is ATP $\rightarrow$ n.ATP.

Autocatalysis may be realised by 1. either a single reaction *e.g.* $\Sigma x_i + A \rightarrow \Sigma y_j + \Sigma B_k$ (where n of the $B_0, B_1 ... B_k$ products are equivalent to A) or a linear chain of reactions

through intermediate products, called template autocatalysis, branching chain reactions or autocatalytic sets (King 1978), or 2. by a reaction network. Networks may be indirect through a series of intermediary products, or collective where there is no connection between the component cycles other than through catalysis, as seen for example in the replication of viroids where each RNA strand can catalyse the production of the other.

However, in all cases, the reactions reduce to the defining $A \rightarrow A_0...A_n$ pattern described by Ostwald's differential equation (Plasson et al. 2010).

3.2 Classification of Artificial Chemistries

A straightforward scheme for classifying artificial chemistries is given in Faulconbridge (2011), based on the relationship between a molecule's representation and its properties:

- *Symbolic*. Symbols/molecules have no inherent meaning, so no "implicit reactions", only preprogrammed ones.
- *Structured*. One or more atoms arranged in a structure (such as a string, tree, or graph) with bonds within atoms to form molecules. This results in unlimited capability, but is computationally expensive.
- *Sub-symbolic*. Emergence of properties from lower-levels (*e.g.* real chemistry, also neural networks). The symbols (for example, atoms) have an internal structure which determines the macro-properties of the chemistry.

Another system was presented by Nellis (2012, p.132) with artificial chemistries categorised according to a quite different set of three factors:

- *World*. Physical elements such as molecules and atoms.
- *Chemistry*. The interactions between the elements.
- *Constraints*. Restrictions on how the world and chemistry interact. Nellis gives as an example a binding model that describes which interactions are possible between elements.

However, there is an earlier, and more influential, categorisation by Dittrich, Ziegler, and Banzhaf (2001) which also coincidentally categorises artificial chemistries against three factors, and which is commonly used in the literature (*e.g.* Lenaerts and Bersini (2009) and Gardiner, Harland, and Hamilton (2007)). This taxonomy is described in more detail in the section below.

3.3 <S,R,A> classification scheme of Dittrich et al.

In Dittrich, Ziegler, and Banzhaf (2001), a broad variety of artificial chemistries are classified according to the tuple <*S*,*R*,*A*> where *S* describes the form of the component molecules in the artificial chemistry, *R* the rules for the interactions between the molecules, or reaction rules, and *A* the mechanism used to select molecules for reactions[2].

3.3.1 Set of molecules (<*S*>)

This element describes how:

- Molecules are represented in the artificial chemistry, perhaps as labelled graphs (Faulconbridge 2011) or binary strings (Banzhaf 1994), and
- How any implicit properties of molecules may be derived, such as bond energy calculations by Extended Huckel Theory (EHT) (*e.g.* Benkö, Flamm, and Stadler (2003).)

3.3.2 Reaction rules (<*R*>)

Reaction rules in an artificial chemistry describe how reactants are transformed to reaction products. Reactions may be pre-defined (often the case when simulating real-world chemistry where exact pathways are important) or dynamically determined using molecular properties. In artificial chemistries enforcing conservation of energy, elements are neither created or destroyed so reactions can be represented solely by rearrangements or bond changes. Tominaga et al. (2007) showed for a particular artificial chemistry, that it is computationally universal with only uni-molecular and bimolecular reactions; the reaction rules in most artificial chemistries therefore only describe these forms of reaction.

Constructive chemistries

The action of the reaction rules are *constructive* (Fontana, Wagner, and Buss 1994) if new components may be generated through the action of other components—a form of emergence, explicitly linked to the production of novelties: "construction: to understand how the organizations upon which the process of natural selection is based arise, and to understand how mutation can give rise to organizational, that is: phenotypic, novelty."(Fontana, Wagner, and Buss 1994)

A strongly constructive system is one which maintains closure[3], and in which there

[2] An alternative <*S*,*I*> form is also described, where *I* combines the *R* and *A* components, but is less descriptive and less commonly used in the literature.

[3] Or as stated by Fontana, Wagner, and Buss (1994, p. 217),"A strongly constructive system that contains agent *A* must cope with the network of its implications. But, then, it also must cope with the implications of the implications. And so on."

is self-consistency and some form of logical structure. Both implicit laws and implicit molecule definitions are required for constructive chemistries:

- Implicit reaction laws are molecular structure-based, while explicit laws are independent of molecular structure.
- Implicit molecule definitions provide a description for a molecule's construction, while explicit molecule definitions are taken from an fixed set of symbols.

Strongly constructive chemistries provide a mechanism for exploration in the EA sense: new products can be generated, and new products can participate in reactions. A weakly constructive approach might be to pre-specify all possible reactions; the strongly constructive approach is to generate reactions "on-the-fly" from the structures of the interacting molecules.

Pre-specification is well suited to simulating real chemistry as it allows properties of reactions observed in chemical experiments to be attached to their simulated equivalents. However, it does not allow for arbitrary reactions, and it requires reaction properties to be predetermined—rather difficult for novel or artificial reactions. Clearly, although constructive, this is not strongly constructive (in the sense of Fontana, Wagner, and Buss (1994) and Dittrich, Ziegler, and Banzhaf (2001)) as it is not open-ended. However, Hartenfeller et al. (2011) suggests that it is still useful for applications such as drug discovery.

An artificial chemistry with the ability to create reactions "on-the-fly" given a set of possible reactants may discover more than one possible reaction pathway between the same reactants and products. The method used to choose one reaction pathway from the alternatives is an important component of the artificial chemistry, and the mechanism may be tuned or tailored to privilege or preferentially chose particular types of pathways independent to other factors such as temperature or concentration. In Chemistry the choice of reaction pathway is fundamentally linked to those other properties and cannot be treated independently.

3.3.3 Reactor algorithm (<*A*>)

The reactor algorithm provides the mechanism to select (in effect, to order) reactions, or in the words of Faulconbridge (2011, sect. 4.1.3), the "...algorithm which describes the order of and intervals between reactions, starting from an initial collection of molecules..."

If the reaction rules are analogous to chemistry, the reactor algorithm is analogous to physics: the way molecules move and collide in the reaction vessel determines which reactions are possible (for example, by providing enough kinetic energy to overcome a reaction's activation energy), and the order in which the reactions occur.

Faulconbridge (2011) identifies three basic types of *mixing method* or reactor algorithm:

1. *Well-mixed/aspatial*. Either discrete time (uniform probability distribution for selecting reactions) or continuous time (Gillespie 1976) assuming the reactions are known in advance (which of course is not possible for a strongly constructive artificial chemistry.)

2. *N-dimensional*. A grid with reactions in adjacent cells, or a continuous space where molecules have position and velocity. The main advantage is the ability to simulate spatial affects; the primary disadvantage is performance.

3. *Mixed scale*. Hierarchical spaces, such as aspatial cells within bigger grid, mostly for simulating biology (*e.g.* Jeschke et al. (2008)). One possible advantage is the potential for parallelisation.

3.4 Applications in real-world chemistry

Artificial chemistries may be applied to backward-chain from a set of desired products to identify a set of currently-available initial molecules, for example in drug discovery (*e.g.* Hartenfeller et al. (2011)). A second use is in reaction network discovery, where the goal is to describe a closed set of reactions and reactants from some initial reactants and reactions (*e.g.* Faulon and Sault (2001)). Finally, artificial chemistries can be used in the modelling of biological phenomenon such as enzyme function (*e.g.* Flamm et al. (2010).)

The primary requirement in these cases is fidelity with real-world chemistry, which requires either a library of empirically derived reaction definitions and rates, or a model capable of accurately simulating quantum-mechanical processes. The latter approach has been taken by a family of Artificial Chemistries, beginning with Benkö, Flamm, and Stadler (2003), built on Extended Hückel Theory with parameters taken directly from chemical experiments and later extended (for example in Benkö, Flamm, and Stadler (2005)) to a general purpose model with parameters derived from theoretical chemistry. The model was used in Högerl (2010) for the study of the behaviour and topology of chemical reaction networks, specifically Diels-Alder and Formose reaction networks, and in a series of papers (*e.g.* Flamm et al. (2010) and Ullrich et al. (2010)) for the examination of the evolution of metabolic networks in early organisms using a simple model of RNA coding for catalysts.

The core problem in drug-discovery is the selection of a set of reactions to generate a given product. The CASREACT Chemical Reaction database [4] references more than 80.5 million known reactions in published work; clearly it is impractical in a laboratory setting to simply apply each of these in turn to a set of initial molecules in the hope that the goal compound will emerge. Hartenfeller et al. (2012) used simulations in reaction-SMARTS and RDKit (Landrum 2013) to conclude that a much smaller set of 58 reactions, acting on a set of 10,000 to 50,000 "building block" molecules, might instead meet the requirements

[4] `https://www.cas.org/content/reactions`

for *de novo* drug discovery. Of the 58 nominated reactions, 29 were ring-forming; the suggestion is that as so many interesting compounds involve molecular rings, the lack of ring forming reactions was a deficiency of previous approaches. With this small reaction set, chosen for the practicality of transferral to the laboratory benchtop, the combinatorics are such that a search mechanism is needed to identify the most promising pathways from the initial molecule set to the destination compound.

Finally, Artificial chemistries have been used as tools to explore the role of function prediction, for example in the interesting, but necessarily simplified, approach taken by Flamm et al. (2010) and Ullrich et al. (2010). Determining the shape, and hence the function, of an enzyme from its RNA transcript is perhaps the most important problem in current molecular biology: the three-dimensional structure of a molecule—the arrangement of the elements in space—drives many of even the simplest chemical reactions. For example, acid-base reactions result from the shift of charge from one region of a molecule to another, revealing one region while shielding another from activity. Some of the most complex reactions are shape-driven: the function of many enzymes (biological catalysts) derives from their shape, and furthermore this shape is often under regulatory control. Flamm et al. (2010) and Ullrich et al. (2010) attach a catalytic function to a molecule based on its secondary structure (shape) and then investigates the influence of these functions upon the evolution of early metabolic networks.

3.5 Origins of life modelling with artificial chemistries

Real-world chemical processes are also important to modelling scenarios for the origin of life or other related areas such as the formation of metabolic networks in the earliest protocells. In most cases though the specific focus is less on the bottom up model constructed from the most basic elements (although Kauffman's autocatalytic protein sets, and Kaneko's protocell toy mode are counter-examples,) and more on task-based models of processes where the particular base level is predetermined by the researcher, such as Ganti's chemoton (Gánti 2003).

Farmer, Kauffman, and Packard (1986) describe an ODE model of polymers where bidirectional reactions connect each polymer condensate c from monomer or polymer constituents a and b, catalysed by some enzyme e, in the presence of water h. All reactions are catalysed (by some randomly chosen polymer), principally to reduce computational complexity; this is justified by the order of magnitude difference in the reaction rates between catalysed and uncatalysed reactions. The model commences with a food set or initial population of monomers and simple polymers in a well-mixed chemostat, no perturbations, and ends when no further new polymers are generated.

Some other examples of artificial chemistries for the investigation of the origin of life

include:

- *Lattice artificial chemistry* (Madina, Ono, and Ikegami 2003; Ono and Ikegami 2000). Membrane formation and cell division, assumeing five different types of particles (some hydrophilic and some hydrophobic) that together form an autocatalytic cycle similar to those observed in biological cells.

- *SCL*. Three types of particle are employed by the Substrate-Catalyst-Link (or SCL) chemistry of Varela, Maturana, and Uribe (1974) and Suzuki and Ikegami (2008): the eponymous Substrate (S), Link (L) and Catalyst (C). Cells are formed from links around a catalyst, with a single predefined reaction rule $S + S \xrightarrow{C} L$ and some straightforward constraints on movement of the particles in the matrix (for example, bonded Link particles cannot cross each other.)

- *Flamm et al. (2010) and Ullrich et al. (2010)*. The evolution of metabolic networks in early organisms using a simple model of RNA coding for catalysts

- *Högerl (2010)*. The simulation of chemical reaction networks characteristic of the transition to life, specifically the Diels-Alder and Formose reaction networks.

- *Dorin and Korb (2006)*. Concerned, almost uniquely, with a chemical ecosystem, based on a set of atoms interacting in pre-specified ways to represent biological photosynthesis, respiration and biosynthesis (or growth). The goal is to explore the interactions in an ecosystem made up of a set of organisms pre-built to perform various defined roles.

- *Gardiner, Harland, and Hamilton (2007)*. A string-based chemistry to investigate protein metabolism evolution under genetic control. Three types of molecule—protein, gene and service molecule—react in ways determined by the types of interacting molecules. The type and pattern of molecules define the type of interaction.

- *Fernando and Rowe (2008) and Fernando and Rowe (2007)*. A flow-reactor for the evolution of metabolism in lipid aggregates based on predefined molecular types and reactions.

3.6 Artificial chemistries and Alife

Artificial Chemistries have also been used in the exploration of open-ended or creative evolution (*e.g.* table 3.1). Squirm3 (Hutton 2002; Hutton 2009) adopts fixed molecule types, and pre-defined reactions for replication and gene-sequence transcription, and so although capable of interesting behaviour is not capable of unlimited extension. StringMol (Hickinbotham et al. 2011), a bacterial inspired microprogram chemistry, demonstrates a rich

Table 3.1: A sample of Artificial Chemistries for open-ended evolution. Constructive chemistries are capable of ongoing extension.

Chemistry	Constructive? (see section 3.3.2)
Ducharme, Egli, and Legault (2012)	Yes
StringMol (Hickinbotham et al. 2012)	Yes
RBN-World (Faulconbridge 2011)	Yes
Lenaerts and Bersini (2009)	Yes—molecular interactions
NAC (Suzuki 2006)	Yes
GGL/ToyChem (Benkö, Flamm, and Stadler 2005)	Yes
Substrate-Catalyst-Link (SCL) (Varela, Maturana, and Uribe 1974; Suzuki and Ikegami 2008)	Unknown
Fernando and Rowe (2008) and Fernando and Rowe (2007)	No—atomic reactions
Gardiner, Harland, and Hamilton (2007)	No—atomic reactions
Lattice artificial chemistry (Ono and Ikegami 2000; Madina, Ono, and Ikegami 2003)	No
GGL/ToyChem (Benkö, Flamm, and Stadler 2003)	No—pre-defined reactions only
Squirm3 (Hutton 2002; Lucht 2012)	No
ZChem (Tominaga 2004)	No—reactions are atomic with wildcards

inheritance mechanism using string-matching as a model for molecular binding, and RBN-World (Faulconbridge 2011) shows that a form of Random Boolean Network, with the addition of a bonding mechanisms to allow for composition and decomposition of RBNs, can be used to build a chemistry capable of almost limitless extension out of non-traditional components.

3.7 Conclusion

In this chapter, we have:

- Introduced the basic elements of artificial chemistries.
- Established those elements within a commonly-used classification scheme (Dittrich, Ziegler, and Banzhaf 2001) which uses the three factors—S describing the form of the component molecules, R the rules for the interactions between the molecules, and A the mechanism used to select molecules for reactions—to classify artificial chemistries.
- Provided a reminder of the requirement for constructive reaction rules in open-ended chemical systems.
- Reviewed the three dominant areas for the use of artificial chemistries: real-world chemical modeling, artificial life, and the origin of life.

- Suggested, by providing examples of other artificial chemistries in the exploration of Alife and the origin of life, the suitability of an artificial chemistry for addressing our research questions.

From the brief review in section 3.5 of artificial chemistries in the exploration of the origin of life it is clear that there is little consensus around the form of the preferred artificial chemistry to adopt. In the next chapter, we shall move from the general description of artificial chemistries to the introduction of a new and specific artificial chemistry, ToyWorld.

4 Reaction Cycles in a Molecular Artificial Chemistry

The work in an earlier chapter established that heritability increases under selective pressure until either *ideal* replication is achieved, or, when the selective pressure is varying, replication is *real*. In both cases however heritability increases and the variability of selective pressure causes a tuning or adaptation of the heredity mechanism to the environmental conditions. In intuitive terms, heritability fixes at the maximum where risk is low, and adapts to the level of risk otherwise. Entities within the model were described by only two parameters: fitness (implying selection) and heritability (implying multiplication).

The key question that motivates the remaining work in this thesis is: can we replace the simple abstract entities with realized entities with associated and emergent (endogenous rather than exogenous) fitness and heritability? More specifically, are we able to create realized entities that are equivalent to the abstract entities in the model, and does the model from chapter 2 still hold with realized entities?

The approach taken in the following chapters is to begin with a simple molecular artificial chemistry and demonstrate one-by-one the required elements for the realized entities, beginning in this chapter with foundational elements (reactions and cycles), then multiplication (required for heritability) and finally variability, all under various forms of selection.

We begin by introducing a new molecular artificial chemistry. There is no shortage of existing artificial chemistries so it is natural to ask why we introduce yet another. The primary reasons are lack of access to existing code; specific requirements of the artificial chemistry imposed by the planned series of experiments; an easy ability to modify and extend the chemistry to include specific functions, and finally from an engineering viewpoint, the understanding that comes from designing and building a system from the beginning.

The experimental requirements are however the primary driver. Our need is for a chemistry that is parameterised in all important aspects, so that experimental factors may be easily transformed into model parameters.

Other types of chemistry are likely equally suitable. But there are pragmatic reasons driving us towards a semi-realistic chemistry:

1. We can leverage existing chemical concepts–reactions, bonds, energy transformations–

without needing to invent and justify a new model.

2. The most relevant previous work is from the related field of origins of life research, and this of course assumes real world chemistry. It is more likely we shall be successful in using this material to influence and inform if we adopt a closely aligned chemistry rather than something tangential.

3. The overall argument for self-evolving evolutionary systems is postulated on emergence from low-level elements; a clear pathway exists to real-world chemistry from physics, and by analogy we can carry this correspondence through into our artificial chemistry. For example, reactants may be chosen based upon physical locations; possible reactions are driven by the energy of molecular collisions. These aspects follow naturally if we adopt a semi-realistic chemistry.

The potential complexity of a semi-realistic chemistry however introduces a difficulty. No artificial chemistry can accurately model all aspects of real-world chemistry, emergent as that is from the quantum world; therefore, all semi-realistic artificial chemistry face decisions as to which elements of real-world chemistry are to be modelled and which left aside. Increasing fidelity in most cases comes at a price in performance (for example, if a quantum-physics simulator is employed to model atomic interactions) but this fidelity to real-world chemistry may in fact be unnecessary for the purposes of the artificial chemistry. And with choices comes the risk of arbitrariness–if bonds require energy to break and form, then exactly how much energy is required in each specific case? And how is this choice justified?

Therefore, our principle is that only the "core" elements of real-world chemistry are incorporated into our artificial chemistry–atoms, molecules, reactions. Further complexities, such as in the introduction of an energy model for the transformation between bonds and kinetic energy, are included as an optional element, controlled by a model parameter. The risk of arbitrariness is mitigated by adopting as closely as is possible real-world values, such as for bond energies. This strategy not only provides consistency, but is somewhat justifiable on the grounds of maximising where possible the correspondence with real-world chemistry.

4.1 The ToyWorld artificial chemistry

ToyWorld, our artificial chemistry for the exploration of emergent replicators, was first introduced in Young and Neshatian (2013). The main elements of the model–Atoms, Molecules, Reactions, a Reaction Vessel–are recognisable from real-world chemistry, but in highly simplified forms. The most pertinent divergence is that we do not model catalytic reactions–remember that our thesis is that autocatalytic sets based on stoichiometry alone can give

rise to variable replicators.

The ToyWorld model has many degrees-of-freedom, and these must be constrained by parameter choice before the model can be used in simulation. Some values are important to our thesis, and so are considered true independent variables in our investigation. These are examined fully. The remainder however are those to which the simulation is insensitive, but still must be specified. For these we prefer real-world values rather than arbitrary artificial values. In short, where we consider something important, we investigate. Where we think it less important, we use a consistent set of pre-existing values–real-world chemistry.

The name has been chosen as both an acknowledgement of ToyChem (Benkö, Flamm, and Stadler 2003), and as a hint at the simulation's purpose: creating an artificial world for exploration. That world, the ToyWorld model, consists of:

- Analogues of physical elements such as atoms and molecules;
- An overall energy model that describes the transformations that can occur between potential, kinetic and internal energy;
- A physics model to describe how molecules interact within the reaction vessel; and
- A chemical model that details the bond changes that can occur when two molecules collide.

All atoms, and therefore molecules and reactions, are contained within a reaction vessel. ToyWorld provides a basic energy model, where molecules have kinetic energy and bond breaking requires energy input and bond formation releases energy. The reaction vessel, which provides the strategies by which reaction reactants (or input molecules) and products (output molecules) are determined, is described in detail in the following sections.

In our model, all reactions emerge solely from the properties of the reacting molecules. For each reaction between two molecules we generate a list of reaction alternatives by enumerating all possible bond additions, bond subtractions, and changes in bond type between the reactants. For example, the reactants H_2 and O_2 generate three reaction alternatives: breaking of the H-H bond, breaking of the O=O double bond, and a transformation of the O=O double bond to a single bond. The reactants H^+ and OH^- give two alternative reactions: breaking of the O-H bond (giving $H+H^+ + O^-$) and formation of a single bond between H^+ and O to give H_2O.

The ToyWorld "main loop" is described at a high-level in alg. 4: in outline, the sequence is to choose reactants from the molecules in the reactor (later described as reactant selection, or $S_{reactant}$), generate a set of possible reaction products from those reactants and select one set of products from those alternatives (product selection, or $S_{product}$) and then finally replace the reactants in the reactor by the selected products. One time through this sequence constitutes a iteration.

ToyWorld is custom Python code that uses third-party packages and libraries for certain functions. The most significant chunks of external functionality come from RDKit, an open-source toolkit for ChemInformatics, and PyMunk, a physics toolkit. Additionally, we use the NetworkX package for Python for network graph manipulations (mainly in the code that searches for cycles in the reaction graphs), and we use arrays from the Numpy package for Python to record the molecular populations.

RDKit provides a number of convenient functions to the ToyWorld Chemical Model:

1. Format conversions between RDKit molecules and SMILES (Daylight Chemical Information Systems 2011), a standard language for molecular representation in chemistry, and conversions from the representation of a molecule into a canonical form.

2. Reference information: for example, the mass of an atom, and the number of outer electrons for an atom.

3. Molecular structure manipulation: iteration over the atoms or bonds in a molecule, and the addition, modification and deletions of bonds and atoms.

4. Utility functions: combining two representations, each of one molecule, into one representation of two molecules, and vice versa, and sanitising the representation of a molecule by checking for molecular validity.

PyMunk is used extensively in ToyWorld for 2D physics calculations, including:

1. Calculating the future position of a molecule inside a reaction vessel based on the molecules velocity and current position.

2. Collision detection between two (or more) molecules.

3. Summing the forces acting on a molecule and adjusting the molecule's acceleration.

4. Simple visualisations of the molecules within the reaction vessel as shapes on a 2D plane.

4.2 Atoms and molecules

Molecules are modelled as an extension of standard RDKit *Mol* objects, constructed from RDKit *Atoms* connected with *Bonds*. Standard Lewis dot structures built on the inherited atomic properties are used to identify possible bonds, and a formal charge model is used to record the charge changes associated with modifications to the molecular structure caused by reactions.

```
1 while t ≤ number of iterations do
      // Selection of reactants
2     Reactants ← outcome of the Reactant selection strategy
      // Construction of reaction
3     Reactions ← enumerate all possible reaction options
4     Reaction ← choose from Reactions according to the Product selection strategy
      // Execute reaction in the reactor
5     Update the reactor with the results of carrying out the Reaction
6 end
```

Algorithm 4: The main simulation loop in ToyWorld.

The lowest level component in the ToyWorld model is the atom, and atoms can be joined by bonds to form molecules. Reactions between molecules are the only mechanism to modify molecules provided by the model; a reaction is simply the addition or subtraction of a single bond between any two atoms in two molecules.

ToyWorld provides a strongly constructive chemistry (Fontana, Wagner, and Buss 1994) where completely new forms of molecules may be generated by reactions, and where the new molecules may in turn take part in further reactions: the chemistry emerges from the lower level atomic properties.

Some examples of molecules created by ToyWorld are shown in fig. 4.1.

4.3 Mass/energy model and energy transformations in ToyWorld

Energy within ToyWorld is modeled in three forms–kinetic, internal and potential.

- Kinetic energy is the energy associated with the motion of a molecule, and equals $\frac{1}{2}mv^2$ where m is the mass of a molecule and v its velocity.

- Internal energy is an abstraction of vibrational energy or heat energy within a molecule. Internal energy, while incidentally consistent with real-world chemical models, is actually included in ToyWorld as a major mechanism to introduce stochasticity in reactions. Without internal energy, all excess energy following a reaction must be allocated to kinetic energy (following energy conservation), and so there is an iso-map between reaction and product kinetic energies. With internal energy, we can divert an arbitrary proportion into internal energy and therefore we have a stochastic multi-map.

- Potential energy is the energy associated with the bonds between atoms in a molecule. Creating a bond reduces the potential energy of a molecule; breaking a bond increases

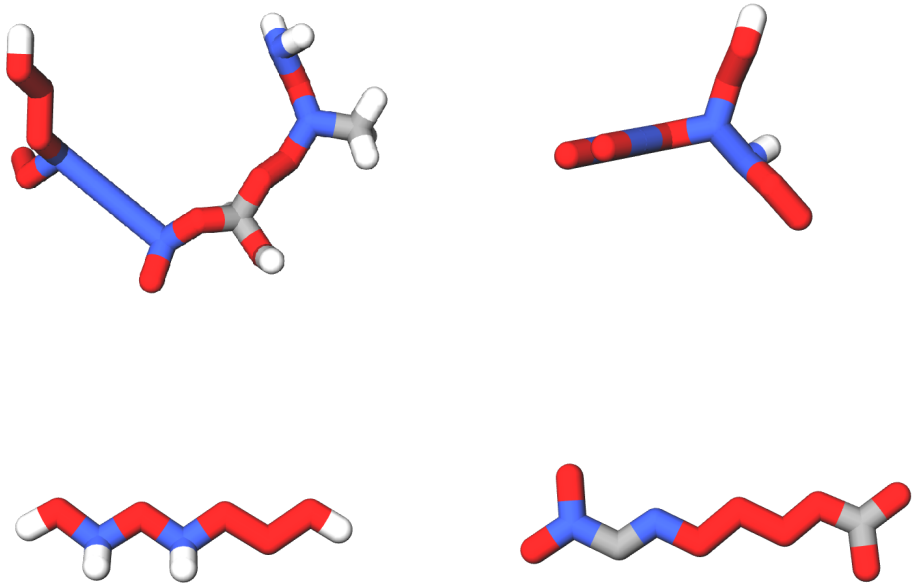

Figure 4.1: Example molecules generated by ToyWorld, rendered as 3-D objects from the original SMILES notation by MolView (`http://molview.org/`) - [H][C]N(O[N])OOC([H])(O[H])ON([O])[N][N]N(O[O])OOO[H] (top left), [H]OON(O[N+]([O-])[N+]([O])[O-])N([H])O[O] (top right), [H]OOON([H])ON([H])O[H] (bottom left) and [O]C(=O)OOOON=[C]N([O])[O] (bottom right).

it. The potential energy of a molecule is the sum of the potential energy of each bond in the molecule. The state without bonds is defined as having zero potential energy, and so molecules with bonds have negative potential energy. The specific value of each bond is absolutely determined by the atomic number of the two atoms at either end of the bond, and the type of bond itself–single, double or triple. ToyWorld provides a table of relative bond values in table 4.2, taken from real-world chemistry [1]. Unspecified bonds are given the average energy of specified bonds of the same bond type. Examples of molecules with associated potential energies are given in table 4.1.

As the model enforces conservation of mass, reactions can be represented solely by bond changes. This follows the approach taken in graph-based chemistries such as GGL/-

[1]Original source: Average Bond Dissociation Enthalpies `http://www.cem.msu.edu/~reusch/OrgPage/bndenrgy.htm`

Table 4.1: Example potential energies calculated by the default Chemistry module using simplified bond energies.

Molecule	SMILES	Potential Energy
H_2O	[H]O[H]	-222.0
H_2	[H][H]	-104.2
O_2	O=O	-119.0
N_2O_4	[O-][N+](=O)[N+]([O-])=O	-434.4
NO_2	N(=O)[O]	-198.0

ToyChem (Benkö, Flamm, and Stadler 2003; Benkö, Flamm, and Stadler 2005) where reactions are modelled as a series of changes to graph edges, or bonds, only.

4.3.1 Energy transformations

The only energy transformations that occur in ToyWorld are those that occur during reactions. The energy at the beginning of the reaction is fully bound in the potential, internal and kinetic energies of the reactants. At the moment of collision, that portion of the kinetic energy due to the collision (kinetic energies of reactants less kinetic energy of centre of mass) is transformed into internal energy in the combined reactants. If a bond forms in this phase the freed potential energy is added into this pool of internal energy; any bond that breaks transforms internal energy into increased potential energy. Post-collision, a portion of the pool of remaining internal energy is transformed back into kinetic energy. The division may be all to kinetic energy, or all to internal, or anywhere in between; as a result, we can model any collision type from fully elastic to fully inelastic collisions.

Reactions are modelled as head-on elastic collisions between two reactants with changes to kinetic energy equalling the increase or decrease in molecular potential energy associated with the creation, destruction or change of order of bonds. Creation of a bond results in a reduction of molecular potential energy and an increase to kinetic energy; destruction results in the reverse. A change in bond type is modelled as the sum of a bond creation and of a bond destruction. Total energy in the system is always constant, and equal to the sum of the initial kinetic energy of all molecules plus the sum of their potential energies.

Reactions conserve energy, and hence the total energy in the reaction vessel should be constant. However, energy can be explicitly added to and removed from the system from outside. Either case is modelled as a uniform change in all molecular internal energies (heat and vibration). This changes the size of the pool from which the kinetic energies of the product molecules are determined after a collision or reaction. Adding energy to the system increases each molecule's internal energy, which increases the size of the pool of merged internal and kinetic energies on collision or reaction, which increases the likely

Table 4.2: Relative energies required to break and/or form a bond. Creating a bond between Atom 1 and Atom 2 releases molecular potential energy in the form of kinetic energy while breaking a bond does the opposite. The values in this table correspond to the -1 level for the E_{Bonds} factor in the experiments of section 4.8.

Bond Type	Atom 1	Atom 2	Energy of both break and formation (in simulation units)
Single	H	H	104.2
Single	C	C	83
Single	N	N	38.4
Single	O	O	35
Single	H	C	99
Single	H	N	93
Single	H	O	111
Single	C	N	73
Single	C	O	85.5
Single	N	O	55
Double	C	O	185
Double	C	C	146
Double	N	N	149
Double	O	O	119
Double	C	N	147
Double	N	O	143
Triple	C	O	258
Triple	C	C	200
Triple	N	N	226
Triple	C	N	213
Quadruple	C	C	200

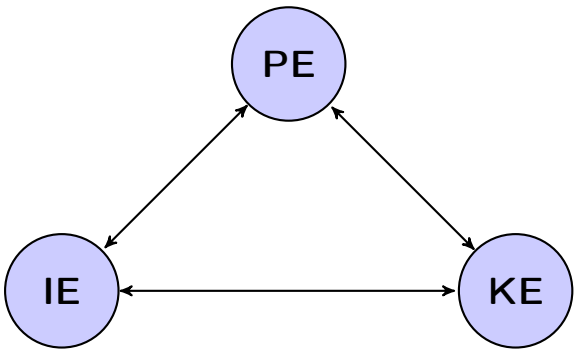

Figure 4.2: Energy transformations in ToyWorld.

kinetic energies of the product molecules. Removing energy from the system has of course the opposite affect, down to the point where effectively all motion and hence all reaction activity ceases.

A reaction may be seen as two stages in sequence: first, the choice of reactants from a population of possible reactant molecules (the Reactant selection strategy, denoted here by S_{Reactant}), and second, the determination of products given that set of reactants (Product selection strategy, denoted S_{Product}). The total energy (that is, potential + kinetic + internal) of the system is maintained, as is the total momentum of the molecules.

In the taxonomy of Dittrich, Ziegler, and Banzhaf (2001), introduced in section 3.3, artificial chemistries are classified along three axis: molecules, reaction rules, and the reactor algorithm. Our terminology, used in the remainder of this work, substitutes Reactant selection strategy (S_{Reactant}) for the Reactor Algorithm of Dittrich, Ziegler, and Banzhaf (2001), and Product selection strategy (S_{Product}) for Reaction rules.

4.4 Reactant selection strategies: selecting reactants for a reaction

Faulconbridge (2011) describes two generic strategies for the selection of reactants–spatial and aspatial–where the primary difference is whether molecular position is a factor in reactant selection. It is possible to further generalise this scheme by considering other differentiating factors. Analogous with real-world chemistry, a cumulative scheme presents itself starting with the pure aspatial, or uniform probability strategy, and then proceeding through a spatial strategy, based on molecular kinetics, to strategies built on kinetics plus intra-molecular and external forces such as electromagnetism.

These strategies are based on a sequence of increasing derivatives of position or location in the reaction vessel; from no position (uniform selection), through fixed position (unin-

teresting as we cannot have a sequence of reactions without motion) to the first derivative (velocity or kinetic selection) and finally to the second derivative (acceleration, or force selection.)

4.4.1 Uniform selection

In a uniform selection strategy ($S_{\text{Reactant}} = \text{Uniform}$), reactants are chosen at random with equal (uniform) probability from the population: no property of a molecule has an effect on the selection. Conceptually we have a well-stirred reaction container with no intra-molecular forces.

4.4.2 Kinetic selection.

By contrast, in a kinetic selection strategy ($S_{\text{Reactant}} = \text{Kinetic}$) molecules have spatial position (and implicitly, velocity) within some assumed reaction vessel, and selection is determined by molecular position–molecules which are spatially co-located (that is, in collision) form a reactant set. Molecules move at constant velocity until they collide with something else (either another molecule or possibly a boundary of an explicit reaction vessel) and then either react, or bounce. Currently in our work we assume that all molecules have a fixed and common size and shape (circular in two-dimensions), irrespective of molecular formula. The algorithm for this strategy is given in alg. 5. At every step the molecules in the reaction vessel move Δ_t units along their velocity vectors, bouncing with no loss of energy off the reaction vessel sidewalls. All molecules that collide during this step are added to the tail of the *ReactantList*.

4.4.3 Intra-molecular selection and external force selection

For completeness, although not considered in the experiments that follow, more complicated forms, where molecular velocities are not constant, can be generated by the introduction of some combination of intra-molecular forces (such as electromagnetism) or external forces (such as gravity or heat.)

4.5 Product selection strategies: determining the products of a reaction

In the ToyWorld chemical model, all reactions arise solely from the properties of the reacting molecules: it therefore defines a *strongly constructive* chemistry in the sense defined earlier. As ToyWorld enforces conservation of mass, reactions can be represented solely by bond changes. This is closely related conceptually to graph-based chemistries such as GGL/ToyChem (Benkö, Flamm, and Stadler 2003), RBN-World (Faulconbridge 2011) or NAC (Suzuki 2006) where reactions are modelled as a series of changes to graph edges.

```
1  Function KineticReactantSelection(Population):
2      if |Population| < 2 then
3          return ∅
       // Step simulation until have a pool of colliding molecules
4      while ReactantList= ∅ do
5          Advance the position of all molecules by Δ_t
6          ReactantList← ReactantList ∪ colliding molecules
           // Adjust Δ_t to maintain pool size
7          if |ReactantList|=0 then
8              Increase Δ_t
9          else if |ReactantList|>10 then
10             Reduce Δ_t
11     end
       // Construct reactant portion of reaction
12     while ReactantList≠ ∅ do
13         Reactants←pop first pair of molecules from ReactantList
14         InitialKE← ∑ KE_i, ∀i ∈ Reactants
15         ReactionEnergy←InitialKE - KE of the Centre of Mass of the Reactants
16         return Reaction between Reactants with energy=ReactionEnergy
17     end
18     return ∅
```

Algorithm 5: *KineticReactantSelection.* Reactant selection strategy where colliding molecules are returned as reactants.

However, in ToyWorld, the graph is implicit rather than explicit as it is in a graph-based chemistry.

For each interaction between two molecules we generate a list of reaction alternatives by enumerating all possible single bond additions, bond subtractions, and changes in bond type between the reactants. Each alternative is the result of a single one of these changes. For example, the reactants H_2 and O_2 generate three reaction alternatives: breaking of the H-H bond, breaking of the O=O double bond, and a transformation of the O=O double bond to a single bond. The reactants H^+ and OH^- give two alternative reactions: breaking of the O-H bond (giving $H+H^+ + O^-$) and formation of a single bond between H^+ and O to give H_2O.

We restrict the options to those that can be generated by a single change to the bond structure of the reactants. This does though mean that long-chain construction by real-world polymerization is not possible in ToyWorld as the modification of a double bond to a single bond (*e.g.* C=C to CC) in each monomer must occur simultaneously with the formation of the replacement CC bond between the two monomers. Long-chain molecules can be constructed from simpler components however in the equivalent to strict polymerization; therefore although this limitation prevents ToyWorld from being used in real-world domains, it is not significant in our domain.

Each reaction alternative can be completely described by the pair of the products of the reaction (that result from the single bond addition, subtraction or change) and the associated change in overall potential energy, which of course will be the same as the potential energy change of the single bond alteration.

Creation of a bond results in a reduction of molecular potential energy, while bond destruction results in an increase. A change in bond type is equivalent to a creation and then a destruction. The magnitude of the change in potential energy, measured in arbitrary energy units, is taken from a table of bond energies for each combination of atoms and bond type (see table 4.2). The standard table is based upon a simplification of real-world chemical bond energies. For example, the creation of a H-H bond releases 104.2 units; the breaking of a C=O double bond takes 185 energy units. The potential energies for a set of example molecules is provided in table 4.1.

How should we choose between alternative sets of possible products for the same reactants? Various product strategies appear plausible: the random choice of an alternative; the most complex alternative; least complex; rarest; most common, and so on, but each strategy requires effort to develop and evaluate.

4.5.1 Least Energy Strategy

When following a Least Energy strategy ($S_{\text{Product}} = \text{LeastEnergy}$) we select a reaction by choosing with uniform probability from a distribution of reaction alternatives weighted

by the total of the energy changes associated with the bond changes. This biases selection towards the Least Energy alternative; the strength of the bias is determined by the degree of the weighting. Examples of the shift in products that occurs as a result as the overall quantity of energy in the system changes can be seen in a later figure (fig. 4.5.)

Let E_i denote the energy required for the bond change in the reaction option i. If $E_i > 0$ the reaction is exothermic, or releases energy; otherwise it is endothermic, requiring energy to proceed.

We calculate a weighted value, e_i, based on the combination of E_i and the available energy for the reaction, $E_{avail} > 0$, as follows:

$$e_i = \begin{cases} |E_i|, & (1) \text{ if } E_i < 0; \\ 0, & (2) \text{ if } E_{avail} < E_i; \\ E_{avail} - -E_i, & (3) \text{ otherwise.} \end{cases}$$

Then, for reaction option i of n options, $p_i = e_i / \sum_{i=1}^{n} e_i$, where p_i is the probability of option i being selected. A number is chosen from the uniform distribution $\mathcal{U}[0, 1]$ and the selected reaction is found from the inverse of the CDF given by p_i for all i by searching for the reaction at that point in the CDF. This method has the property that the probability of a reaction being selected is proportional to its weight.

As a result, for exothermic reactions, highly exothermic reactions are preferred to slightly exothermic ones. For endothermic reactions, the available energy must exceed the energy required by the reaction, and the reaction is preferred according to the degree of the surplus. Note that an option where the energy required exceeds that available ($E_{avail} < E_i$) will have $p_i = 0$, and hence can not be selected. This behaviour is illustrated in fig. 4.3.

4.5.2 Uniform Strategy

We also consider a strategy with minimal bias: a Uniform selection strategy ($S_{\text{Product}} =$ Uniform), where every alternative product set has equal probability of selection.

4.5.3 Energy transformations during a reaction

Molecules have kinetic energy; when they collide the form of the interaction follows from the energy transformations between kinetic and internal and potential energy that are preferred under the chemical model; and finally, the trajectory taken by the resulting products of the interaction is given by their final post-collision kinetic energy.

First, we determine the kinetic energy of the centre of mass of the reactants. The available energy to drive the reaction is the total kinetic energy of the reactants plus the internal

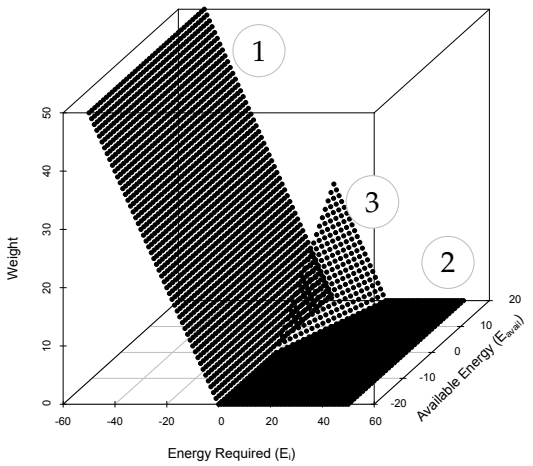

Figure 4.3: Weighting calculation for reaction selection under a Least Energy Strategy. Numerical labels refer to cases in section 4.5.1.

energy of the reactants less the kinetic energy of the centre of mass.

Consider the case where two particles of equal mass but opposite velocity collide. The KE of the centre of mass will be zero (as it is motionless) and the energy liberated by the collision will be the sum of the kinetic and internal energies of the particles.

At the other extreme, consider two equal mass particles, travelling with the same velocities. Intuitively, it is obvious that the only energy released by their infinitesimally gentle collision is from their internal energies, and this is confirmed by the calculation where the kinetic energy of the centre of mass will equal the combined kinetic energies of the reactants, leaving only the combined internal energies.

The available energy for bond modifications is calculated as the sum of the kinetic energies of the reactants less the energy of their centre of mass, plus any energies internal to the reactants. The final kinetic energies of the products equals the sum of the initial kinetic energies of reactants less the change in molecular potential energy from bond changes and the change in internal energies. Total energy in the system is always constant, and equal to the sum of the initial kinetic energy of all molecules plus the sum of their potential energies and internal energies.

The Reactor Algorithm selects one of the reaction alternatives by choosing from a distribution of reaction alternatives weighted by associated energy changes (as discussed in section 4.5.)

4.5.4 Setting product velocities and internal energies

The final step in the reaction mechanism is to determine the velocities and internal energies of the reaction products following the reaction. Although the method is standard physics, there are two complications: the number of products may or may not be the same as the number of reactants, and the pre-reaction energy and post-reaction energy vary as the reaction itself either consumes or liberates energy.

The only constraints are that velocities of the product molecules must conserve momentum and total energy within the reacting system. We recognise that within the frame of reference of the centre of momentum of the reacting system, the vector sum of the momentum of the products must equal zero. Therefore one possible solution to the product velocities is to arrange their vector momentums according to simple geometry: for two products, we arrange their momentums in a line (line 14 in alg. 6), and for three products, an equilateral triangle (line 18). Further extension beyond three products is not required in ToyWorld as, given two reactants, at most three products can result from the single-bond additions, changes or subtractions of the chemical model.

Following a standard method, we first transfer the molecules from the frame of reference of the reaction vessel into the centre of mass (CoM) reference frame by subtracting the velocity of the CoM from each particle (line 3). Correspondingly, we also adjust the energy of the collision by subtracting the KE of the CoM. We recognise that in the CoM frame the vector sum of the momentums will be zero; working in this frame reduces the number of vectors we must sum by one (the momentum of the CoM itself.)

We then choose, from a uniform distribution, the proportion of total available energy to allocate to the product kinetic energies and assign the remainder to internal energy (line 6.) From the kinetic energy allocation we can determine the total scalar momentum of the products using $KE = 0.5 \times velocity \times momentum$, and arrange the vector momentums according to the geometry described earlier.

Finally, we convert from the CoM reference frame to the initial frame by adding back the velocity of the CoM (line 27). This method satisfies our requirements of conservation of momentum and energy for arbitrary numbers of reactants and products while being computationally straightforward. The limitation is that product vectors are arranged in regular and consistent, although reasonably realistic, configurations. An improvement would be to perturb the geometry of the vectors in the CoM frame to remove the regularity.

The outcomes for each reaction alternative from an example collision are shown in fig. 4.4.

kinetic energy of CoM $\leftarrow \frac{1}{2}$(sum of reactant masses)(velocity of CoM)^2
collision energy $\leftarrow$ sum of reactant kinetic energies + sum of reactant internal energies $-$ kinetic energy of CoM
for $i \leftarrow 1$ **to** $|Products|$ **do** `// Transform into CoM frame`
    $v'_i \leftarrow v_i -$ CoM velocity
end
if $|Products|=1$ **then** `// Conservation of momentum implies all excess energy must go into internal energy`
    Internal energy of Products $\leftarrow$ collision energy
else
    KE $\leftarrow$ random($[0,1]$)
    Internal energy of Products $\leftarrow$ collision energy $-$ KE
switch $|Products|$ **do** `// Find a set of momentum vectors that sum to zero...`
    case *1* **do** `// One product`
        $mv' \leftarrow (0,0,0)$
    case *2* **do** `// Two products`
        $mv \leftarrow 2\text{KE} \prod_{i=1}^{n} \text{mass}_i / \sum_{i=1}^{n} \text{mass}_i$
        $mv'_1 \leftarrow (v'_{i\theta} + \frac{\pi}{2}, v'_{i\rho} + \frac{\pi}{2}, mv)$
        $mv'_2 \leftarrow (v'_{i\theta} + \frac{3\pi}{2}, v'_{i\rho} + \frac{3\pi}{2}, mv)$
    case *3* **do** `// Three products`
        $mv \leftarrow 2\text{KE} \prod_{i=1}^{n} \text{mass}_i / \sum_{i=1}^{n} \text{mass}_i$
        $mv'_1 \leftarrow (v'_{i\theta} + \frac{\pi}{3}, 0, mv)$
        $mv'_2 \leftarrow (v'_{i\theta} - \frac{\pi}{3}, 0, mv)$
        $mv'_3 \leftarrow (v'_{i\theta} - \pi, 0, mv)$
end
for $i \leftarrow 1$ **to** $|Products|$ **do** `// Convert momentums to velocities...`
    $v'_i \leftarrow mv'_i / \text{mass}_i$
end
for $i \leftarrow 1$ **to** $|Products|$ **do** `// Transform back to standard frame`
    $v_i \leftarrow v'_i +$ CoM velocity
end

Algorithm 6: Algorithm to set post-collision velocities and internal energies for reactions resulting in one to three products. Extension to more than three products possible by following the same pattern, although this is not necessary in ToyWorld (see text.)

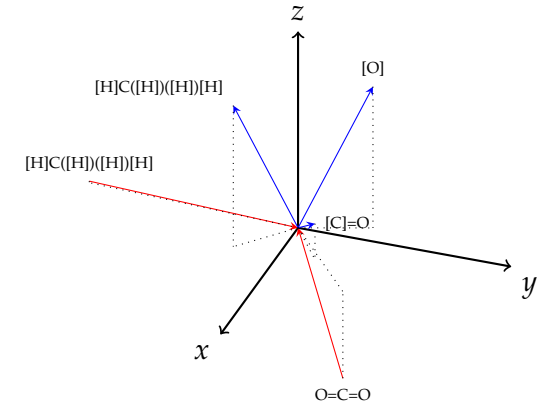

(a) O=C=O + C→ [O] + C + [C]=O

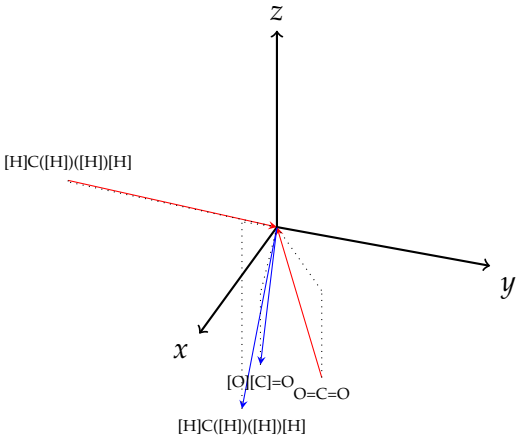

(b) O=C=O + C→ C + [O][C]=O

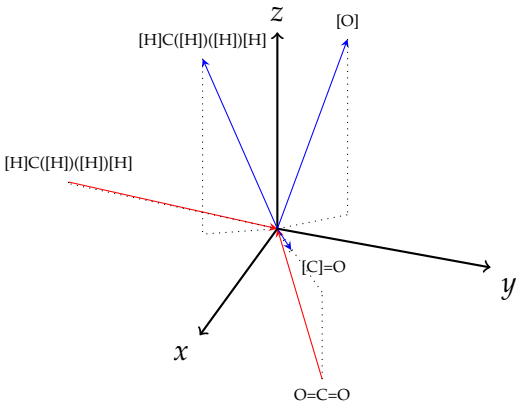

(c) O=C=O + C→ [O] + C + [C]=O

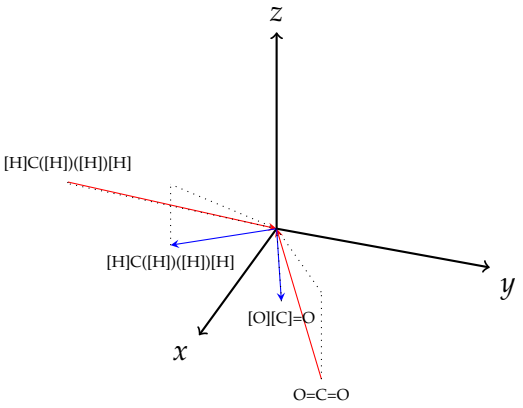

(d) O=C=O + C→ C + [O][C]=O

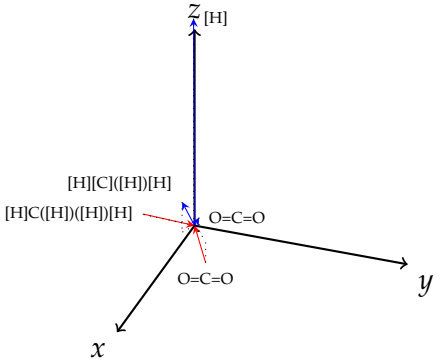

(e) O=C=O + C→ [H] + [CH3] + O=C=O

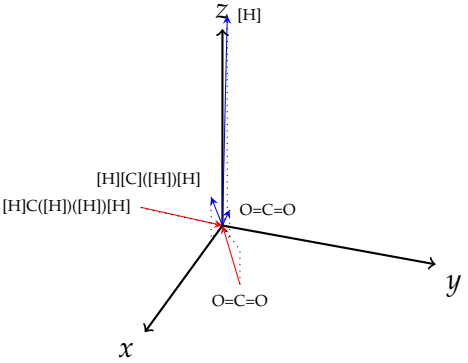

(f) O=C=O + C→ [H] + [CH3] + O=C=O

Figure 4.4: Reaction alternatives generated by ToyWorld from a collision between O=C=O and C molecules. Each alternative generates a unique set of products, with corresponding velocities calculated by alg. 6 so as to preserve energy and momentum.

4.6 Validation of the energy model

The energy-model presented in this chapter leads to some particular expectations for the behaviour of ToyWorld:

1. Given two reactants, changing the reaction energy should result in different sets of reaction products.

2. Molecular quantities should reach equilibrium–that is, the set of interacting molecules is constant, with fluctuations expected in quantities. We expect molecular concentrations to stabilise at non-extreme values (equilibrium rather than driven to an extreme) after some transition period from the initial conditions.

3. The equilibrium point should depend on the energy of the system. Our energy model preferentially forms bonds at low energies, and breaks bonds at high. We expect the average length of molecules in the artificial chemistry to be greater at low energies than at high energies.

These predictions were tested by two experiments: first, we examined the reaction products produced at a range of reaction energies for four sets of reactants. Second, for a given set of reactants, we ran the simulation for 10,000 iterations at four successive initial average kinetic energy levels–0, 67, 133, and 200 units per molecule.

The experiment was run first with a reactant set containing 100xN_2O_4 and 100x2NO_2, and then with a reactant set of 100xH_2, 100xO_2 and 100xH_2O.

4.6.1 Results and discussion

Beyond the initial transition period, both reactant sets showed results essentially consistent with equilibrium. Population variability was high in both cases, but more so for the N_2O_4 and $2NO_2$ reactant set (see fig. 4.5) where some molecules never reached a relatively constant population level (gradient of a best-fit population line remained significantly non-zero.) The fluctuations in the quantities of the other molecules are expected according to our criteria, and result from the inherent variability in reaction selection which causes the quantities to oscillate around a norm.

With both reactant sets the model produced a significant number of molecules which would be considered unstable in real-world chemistry (such as O^- and O.) This is likely an artifact of the method we use to generate reaction options, where a bond-break plus bond-formation reaction–moving through an intermediate unstable ion–occurs in our model as two separate reactions. As all molecules currently react with equal likelihood, significant time can elapse before the intermediate product reacts to form a stable product.

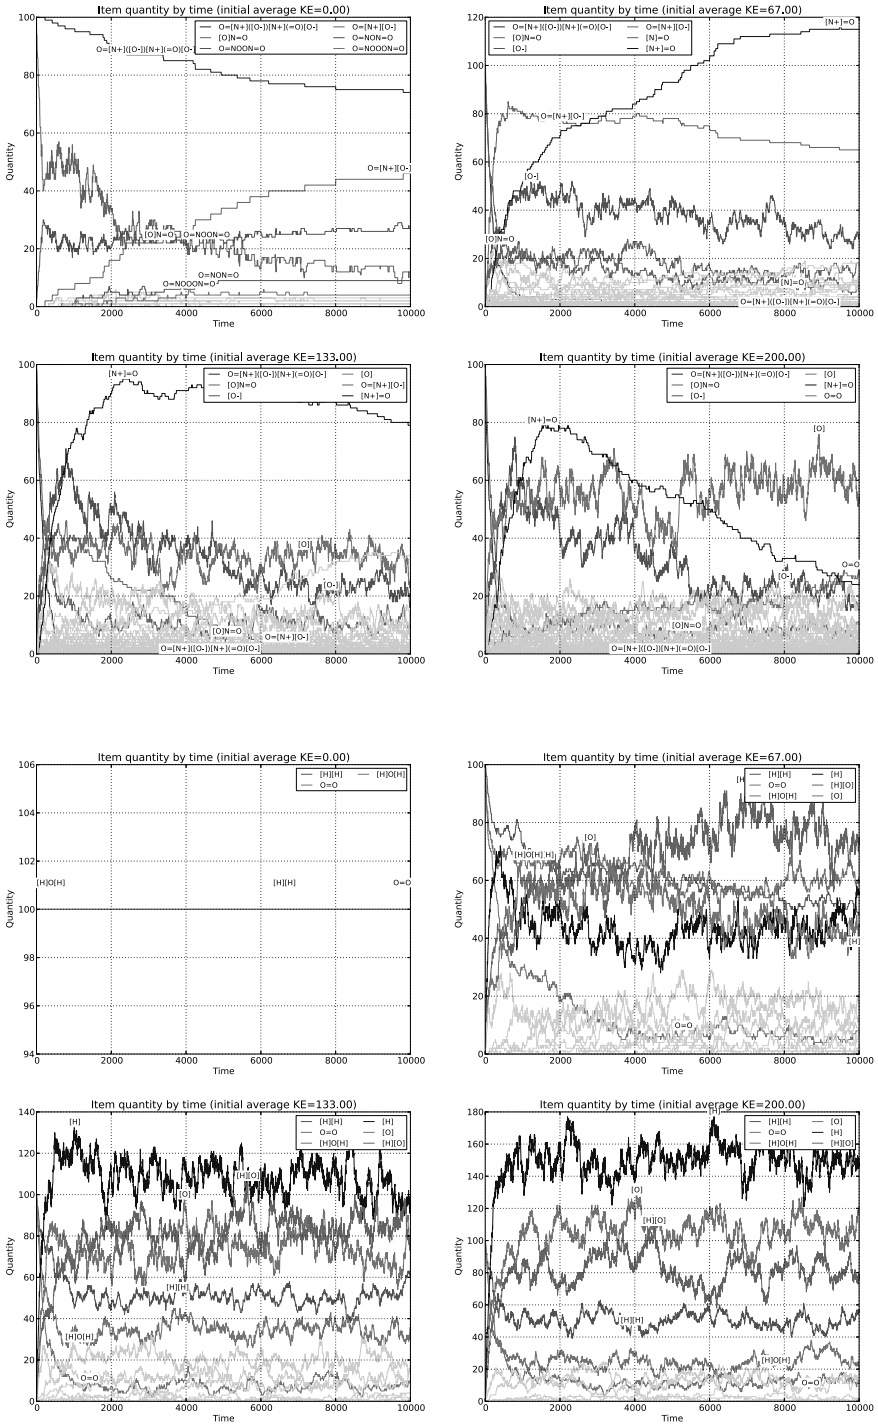

Figure 4.5: Molecular quantities over time for initial population of N_2O_4 and $2NO_2$ (top two rows) and $100xH_2$, $100xO_2$ and $100xH_2O$ (bottom two rows) with initial average KE ranging from 0 to 200 units (only molecules with significant quantities are labelled; remainder appear as light-grey lines). Molecules represented in SMILES notation.

There were clear differences in population composition between the four initial kinetic energy levels in both reactant sets. In the H_2, O_2 and H_2O reactant set, no reactions occurred at the zero energy level. This is expected from our energy model as only bond-formations are possible without free kinetic energy. With reactants of H_2, O_2 and H_2O no bond formations are possible, confirmed by examining the bond options returned by the model for the six possible combinations of initial reactants. By contrast, the reactant set N_2O_4 and $2NO_2$ at energy zero contains one possible bond formation reaction (in SMILES, [O]N=O.[O]N=O to O=N[O][O]N=O) which can proceed without free kinetic energy. This then releases a product which can also react, and so on, thus explaining the different results between the reaction sets.

The energy and reaction models produce results consistent with our predictions for the system's behaviour (with the exception of achieving equilibrium with the N_2O_4 and $2NO_2$ reactant set). An aspatial approach does however come with restrictions. Most obviously, as there is no concept of proximity, there can be no boundaries or membranes or even basic distinctions between *inside* and *outside*. This is critical in biology but it is unclear at this point if this is equally important in non-biological systems. An experimental comparison between the aspatial and spatial approaches in the experiments described in the next section will help to clarify this.

4.7 Response of the model to reactant and product strategies

We now return to the goal of this chapter: to establish the effect of reactant and product selection strategies on cycle formation. Cycles are of fundamental importance in our work as they are the mechanism by which autocatalysis is achieved, and therefore they are themselves candidate replicators.

In this section we explore the following questions:

1. Is there a quantitative difference between different reactant and product selection strategies?
2. Is there a combination of reactant and product selection strategies that leads to the appearance of an increased number of cycles?
3. Is the number of cycles significantly affected by the values of other parameters of an artificial chemistry, such as initial kinetic energy or bond energies?

To the best of our knowledge, this is the first time that reaction and product selection strategies in artificial chemistries have been experimentally compared.

Instead, the general approach of previous work, where there has been a quantitative evaluation, has been to propose a particular strategy, build, and evaluate against the initial goals, rather than against alternatives.

Table 4.3: Factors, or independent variables.

Factor	-1 value	+1 value	Description
S_{Reactant}	Uniform	Kinetic	See Section section 4.4
S_{Product}	Uniform	LeastEnergy	See Section section 4.5
E_{Vessel}	100	300	Initial kinetic energy of each molecule in the reaction vessel
E_{Bonds}	Simplified real-world chemistry. Average values for Single=77.7, Double=148.2, and Triple=224.3 (see table 4.2)	Single=50, Double=100, Triple=200	Energy required to break a bond of the given type

4.8 Experiment design

The experiments follow a full factorial design over four factors (S_{Reactant}, S_{Product}, E_{Vessel} and E_{Bonds}), each at two levels, run in a randomised order, with three (3) replicates of each combination of factors executed in sequence before beginning the next combination. The first replicate of each combination starts with a predefined random seed incremented by one for each successive replicate of the same combination. The factor levels used are given in table 4.3.

Each replicate used the same initial population of 800 molecules, made up of 100 molecules each of [H][H], O=O, [O-][N+](=O)[N+]([O-])=O, and N(=O)[O] and 200 molecules each of O and O=C=O (all represented in SMILES (Daylight Chemical Information Systems 2011).) This initial population is somewhat arbitrary, although reasonable; given that ToyWorld is a strongly constructive chemistry, we would expect that any differences between initial populations would reduce as the simulation proceeds.

4.8.1 Factors

Our two primary factors, or independent variables, are S_{Reactant} and S_{Product}. We also introduce two secondary factors, initial reaction vessel kinetic energy (E_{Vessel}) and bond energy (E_{Bonds}), to assess the sensitivity of the simulation to other parameters.

For simplicity of analysis, all of our factors are two-level, meaning they take one of two possible levels, or values, in each run. The parameter values chosen for each level of E_{Vessel} and E_{Bonds} were chosen as representative from a set of alternatives used in initial exploratory experiments; in each case they allowed the simulation to run for an extended period without running out of possible reactions (from lack of energy for example.)

4.8.2 Response variables

A reaction cycle is formed by a sequence of reactions where one or more of the products of the final reaction in the cycle are of the same species as the same number of reactants in first reaction in the cycle. That is, reaction cycles are partially-closed (as not every molecule must be both reactant and product in the cycle) and capable of autocatalysis.

We concentrate on three related response, or dependent, variables—Number of reaction cycles, Length of longest cycle, and Count of most common cycle. All three are derived from a reconstruction of the network of reactions that occur during each experiment run, where every edge represents a specific reaction connecting a particular set of reactants with a particular set of products. Note that the nodes in the constructed network capture specific molecules, rather than molecular types or species that share the same chemical formula (as would be more usual in the construction of a Reaction Network for real-world chemistry.)

We exclude all unique cycles, and all cycles with three or fewer elements (for example, where a molecule loses, then regains, an atom repeatedly). Unique cycles by nature are unlikely to be representative; very short cycles on the other hand are so common as to dominate other more interesting cycles in any analysis.

4.9 Results

The overall sequence of experiments to run was defined by the experiment design. The experiments were run in the order they were defined, with all replicates of an experiment being completed before moving to the replicates of the next experiment in the sequence. All replicates of an experiment were run with the same set of parameters with the exception of a random seed which varied between replicates in a predictable way. The reactions in the replicate were grouped into fixed-size blocks and run in sequence with results saved incrementally between each block.

All replicates completed a set of 20,000 reactions; given the initial population size of 800 molecules, and from the summary of results below, we believe that this captures a representative set of reactions. This also simplifies the analysis as we can assume a balanced set of treatments in the statistical sense (that is, the sample sizes for all treatments are equal).

A view of the results is given in table 4.4: reaction networks built from the full dataset of 20,000 reactions can be too large for easy analysis. Instead, we choose to partition the reaction data into four equally spaced blocks of 250 reactions each and analyse each block independently.

Table 4.4: Summary of results.

Statistic	Number of cycles	Length of longest cycle	Count of most common cycle
Reactions 4750 to 5000			
Min.	0.00	0.00	0.00
1st Quartile	0.00	0.00	0.00
Median	1.50	3.50	2.50
Mean	219.06	5.04	215.80
3rd Quartile	91.25	7.75	96.00
Max.	5704.00	20.00	2728.00
Reactions 9750 to 10000			
Min.	0.00	0.00	0.00
1st Quartile	0.00	0.00	0.00
Median	6.00	4.00	6.00
Mean	62.10	4.65	169.21
3rd Quartile	68.75	8.25	27.75
Max.	526.00	13.00	6684.00
Reactions 14750 to 15000			
Min.	0.00	0.00	0.00
1st Quartile	1.00	3.00	2.00
Median	5.00	4.50	5.00
Mean	27.17	4.79	42.27
3rd Quartile	34.50	7.00	16.25
Max.	237.00	12.00	862.00
Reactions 19750 to 20000			
Min.	0.00	0.00	0.00
1st Quartile	0.00	0.00	0.00
Median	3.50	4.00	4.00
Mean	20.04	3.90	14.62
3rd Quartile	20.25	6.00	13.25
Max.	199.00	12.00	216.00

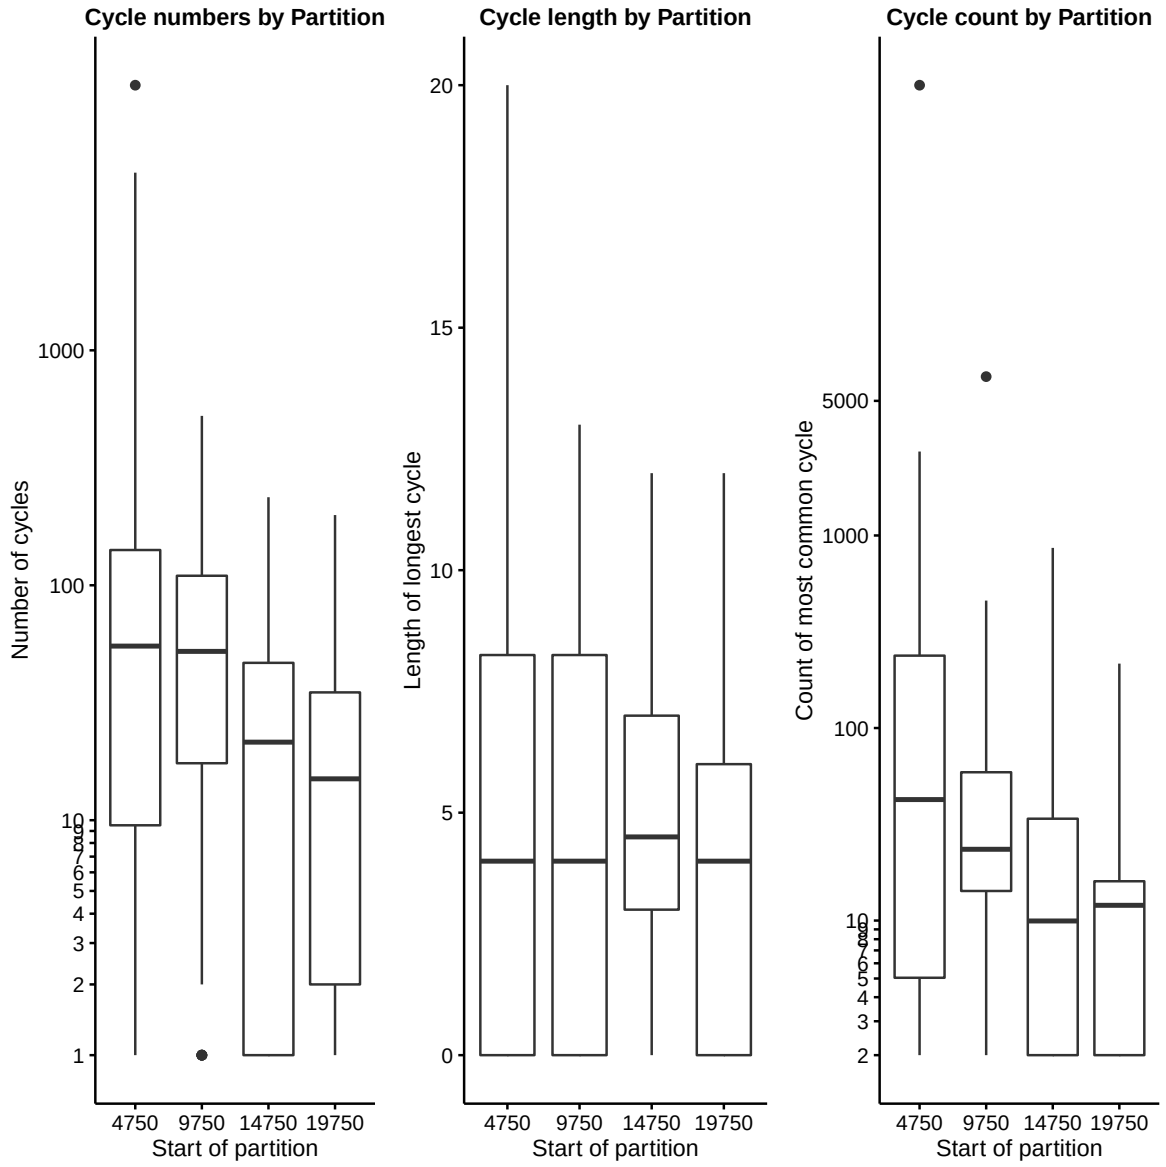

Figure 4.6: Cycles by reaction partition (starting reaction number for each partition along x-axis, y-axis scales vary)

4.10 Analysis and discussion

Figure 4.6 suggests that the first partition, representing the vessel a quarter of the way into its lifespan, is quantitatively different from the other three partitions, with a significantly greater range for all three response variables. Intuitively this corresponds with an initial period where the diversity in the reaction vessel rapidly increases from the limited starting set of molecules, as seen in some (*e.g.* fig. 4.7) but not necessarily all of the replicates. Diversity here is measured by (average molecular quantity)$^{-1}$. All following sections therefore exclude data from the first partition of reaction numbers from 4750 to 5000.

4.10.1 Is there a quantitative difference between the different reactant and product selection strategies?

From visual inspection of fig. 4.9, the number and length of cycles produced by Kinetic reactant selection seems to show a significant increase compared to those produced by Uniform reactant selection. Similarly, from fig. 4.8, there is very little apparent difference between the two product strategies, Uniform selection and Least Energy selection.

We use ANOVA to further examine the relationship of S_{Reactant} and S_{Product} to the response variables using a two-factor with two-levels (2x2) model with interaction effects. There is a highly significant difference ($p<0.001$) between the Uniform and Kinetic reactant selection strategies when comparing the number of cycles (f-value=40.442) and length of cycles (f-value=361.891) (confirming the impression given by fig. 4.9), although again without difference for the count of the most common cycle. The effect of S_{Product} on cycle number and length is also significant (f-value=4.050 and 5.705 respectively, $p<0.05$) and there is a first-order interaction between S_{Reactant} and S_{Product} for number of cycles (f-value=4.011, $p<0.05$).

4.10.2 Is there a combination of reactant and product selection strategies that leads to the appearance of an increased number of cycles?

From fig. 4.11 it is clear that there is no significant relationship between strategy and the number of occurrence of the most common cycle. However, it seems that such a relationship does exist for the number and length of cycles, with the strongest effect as a result of S_{Reactant}, and a lesser effect from the choice of S_{Product}.

We conclude that the greatest levels of emergence are likely to be seen with the combination of $S_{\text{Reactant}} = \text{Kinetic}$ and $S_{\text{Product}} = \text{LeastEnergy}$.

4.10.3 Is the number of cycles significantly affected by the values of other parameters of an artificial chemistry, such as initial kinetic energy or bond energies?

We construct a two-factor with two-levels (2x2) ANOVA model (degrees of freedom=1) with interaction effects to examine the relationship of the independent variables E_{Vessel} and E_{Bonds} to the response variables, and applied it to our dataset (summarised in table 4.4). E_{Bonds} is significant (f-value=4.221, p<0.05) to number of cycles (see fig. 4.10). No other significant relationships exist.

4.11 Conclusions

In this chapter we have introduced the ToyWorld artificial chemistry, a modular molecular artificial chemistry, and explored the effect of four possible combinations of Reactant (S_{Reactant}) and Product (S_{Product}) selection strategies on cycle formation.

By experiment we have shown that the choice of S_{Reactant} is critical to the behaviour of this artificial chemistry; S_{Product} on the other hand appears to have a lesser effect on the emergence of cycles in our experiments. Furthermore, $S_{\text{Reactant}} =$ Kinetic is more effective for cycle emergence than $S_{\text{Reactant}} =$ Uniform. The number of cycles, and length of longest cycle, are both maximized with the combination of a Kinetic Reactant selection strategy and a LeastEnergy Product selection strategy.

The conventional distinction between physics and chemistry is often helpful but not necessarily clear cut in nature. But in our model, the distinction appears both helpful and meaningful. The selection of reactants naturally follows from the spatial relationship of entities in the simulated world. Reaction products are generated by rules based on an artificial chemistry. To be clear, these distinctions aren't the result of the model, as they are in fact choices made in its development, but they do seem to reflect a natural division between domains.

Some theories of life state that it is built on dissipative structures maintained "far from equilibrium" (originating with Prigogine in the 1970s, and as discussed for example in McShea (1998)). Our reaction vessels are closed and so, overall, in equilibrium. Spatial models are not evenly distributed: subareas can be out of thermodynamic equilibrium. On the other hand, the Uniform selection models reflect the global equilibrium, and the differences observed between the two might be explained by this distinction.

Other parameters to the model, such as the initial kinetic energy of the molecular population, and E_{bonds}, the bonding model, had lesser effects on cycle formation, with the only significant relationship being that of E_{bonds} on the number of cycles produced.

The value for E_{Vessel} was not as significant as initially expected, as it only represents the initial energy input, and over time the vessel trends to equilibrium. Models with added energy (modelled by added KE on collision with reaction vessel walls) were trialled but

it proved exceedingly difficult to prevent a runaway effect: the individual molecule KEs generally spiraled out of range.

On the other hand, the bonding model was likely not significant because the alternative factors differed in degree not in kind. Simpler, less real-world chemistry-based, models for bonding are certainly possible however; it would be interesting to compare a variety of bond models in more detail.

However, this leads to an unavoidable limitation of our approach: we must pick a specific instance of the model for experimentation. As we employ simulation, we lack the means to generate a model to produce some desired effect–the relationship between model and effect is, as far as is known, irreversible. This has some interesting implications. Fairly obviously, our results depend on the parameter and model choices made. The corollary, as our models are independent, is that it is difficult to extrapolate our results to other models. This makes it difficult to support broad claims, although we mitigate by attempting to keep the underlying model form itself as general as possible.

Finally, with current technologies it is not possible, even if assumed desirable, to represent the full complexity of the physical world in our models. Abstraction is necessary, but it is still an open question if complexity on the order of real world physics and chemistry is necessary for other forms of life, artificial or biological, or if that just happens to be the case for life on Earth. There is certainly an argument that if life is algorithmic, then any Turing-complete system should be sufficient.

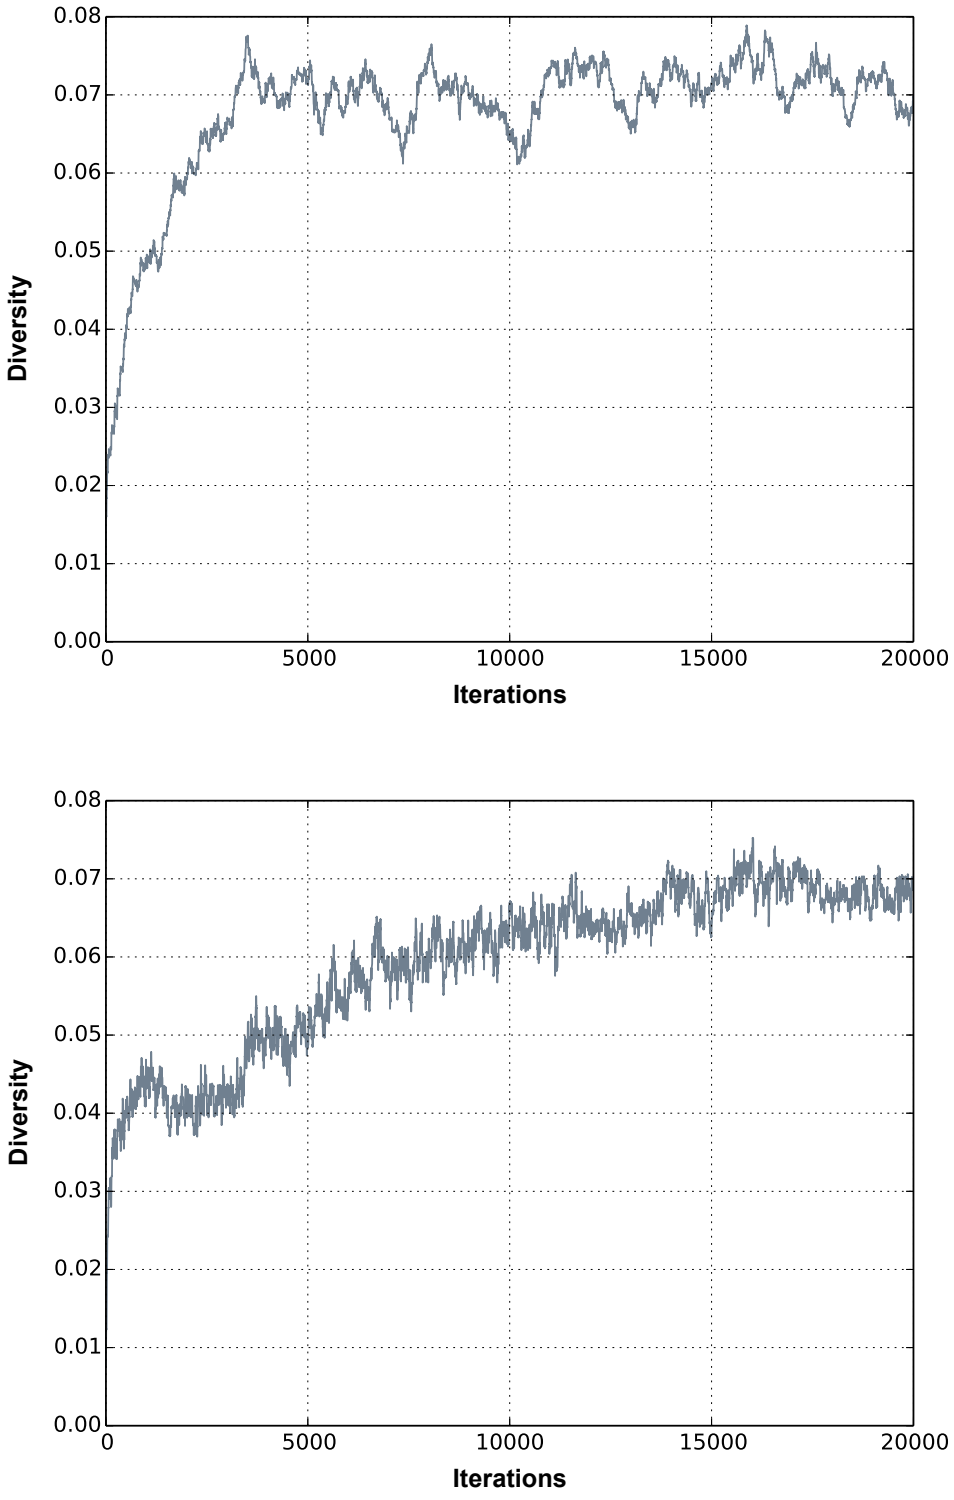

Figure 4.7: Diversity (average molecular quantity^{-1}) for Replicates 12-0 (top) and 16-1 (bottom).

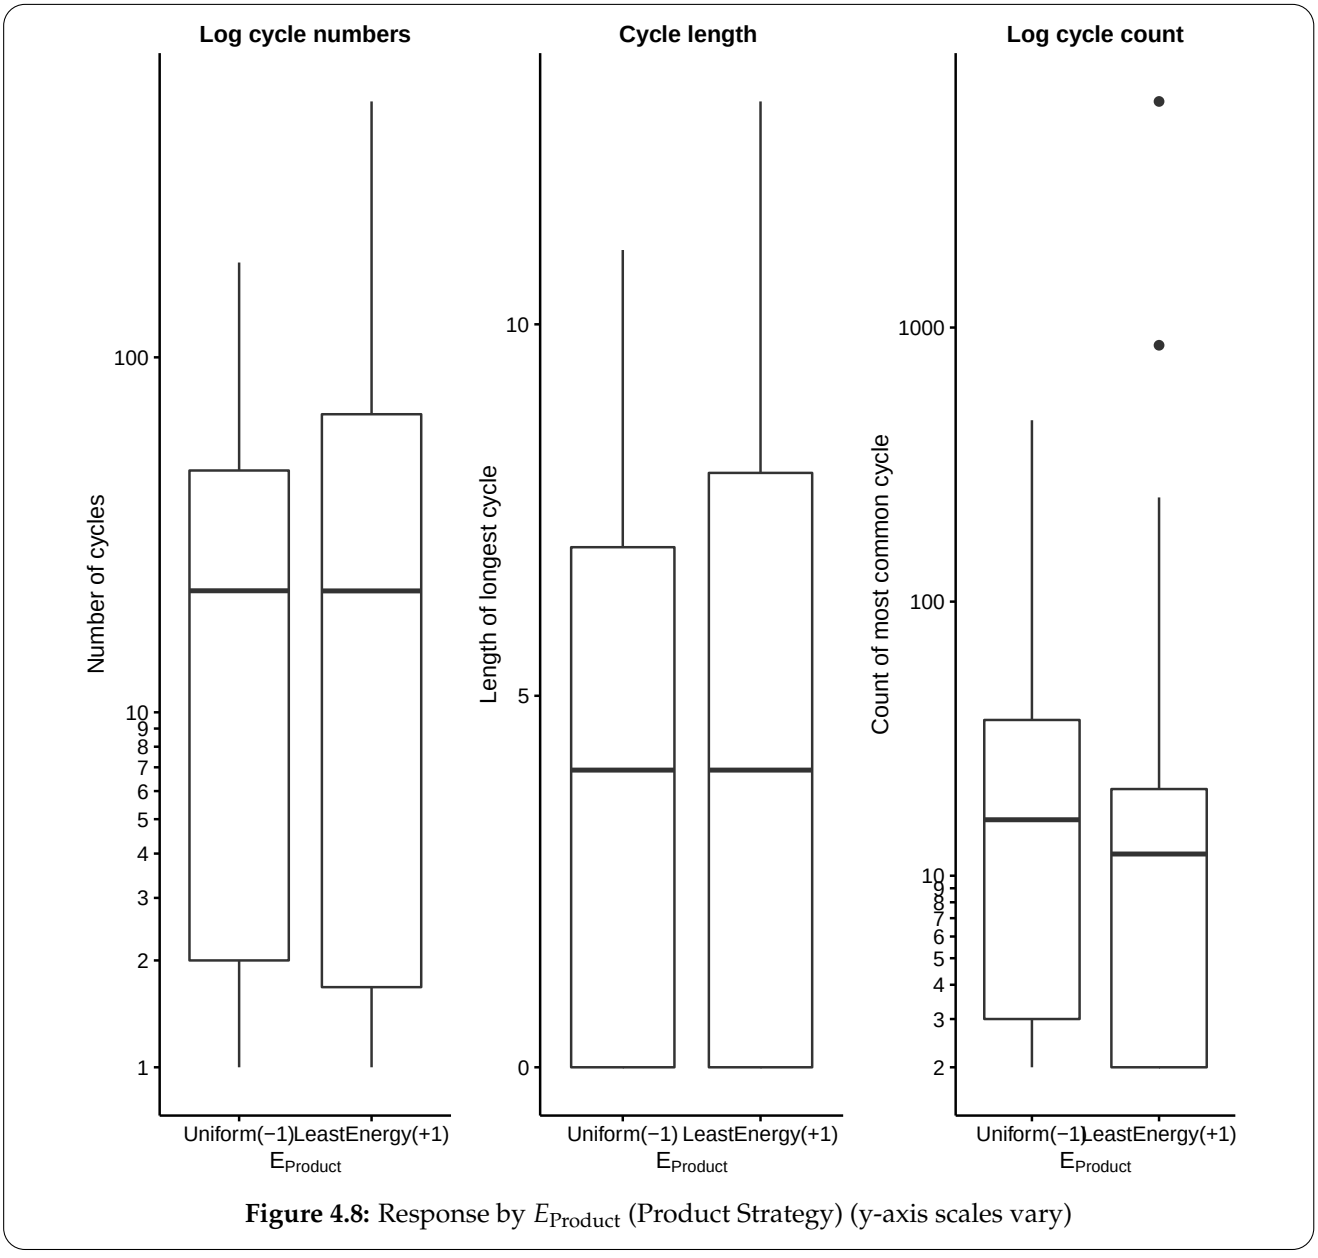

Figure 4.8: Response by E_{Product} (Product Strategy) (y-axis scales vary)

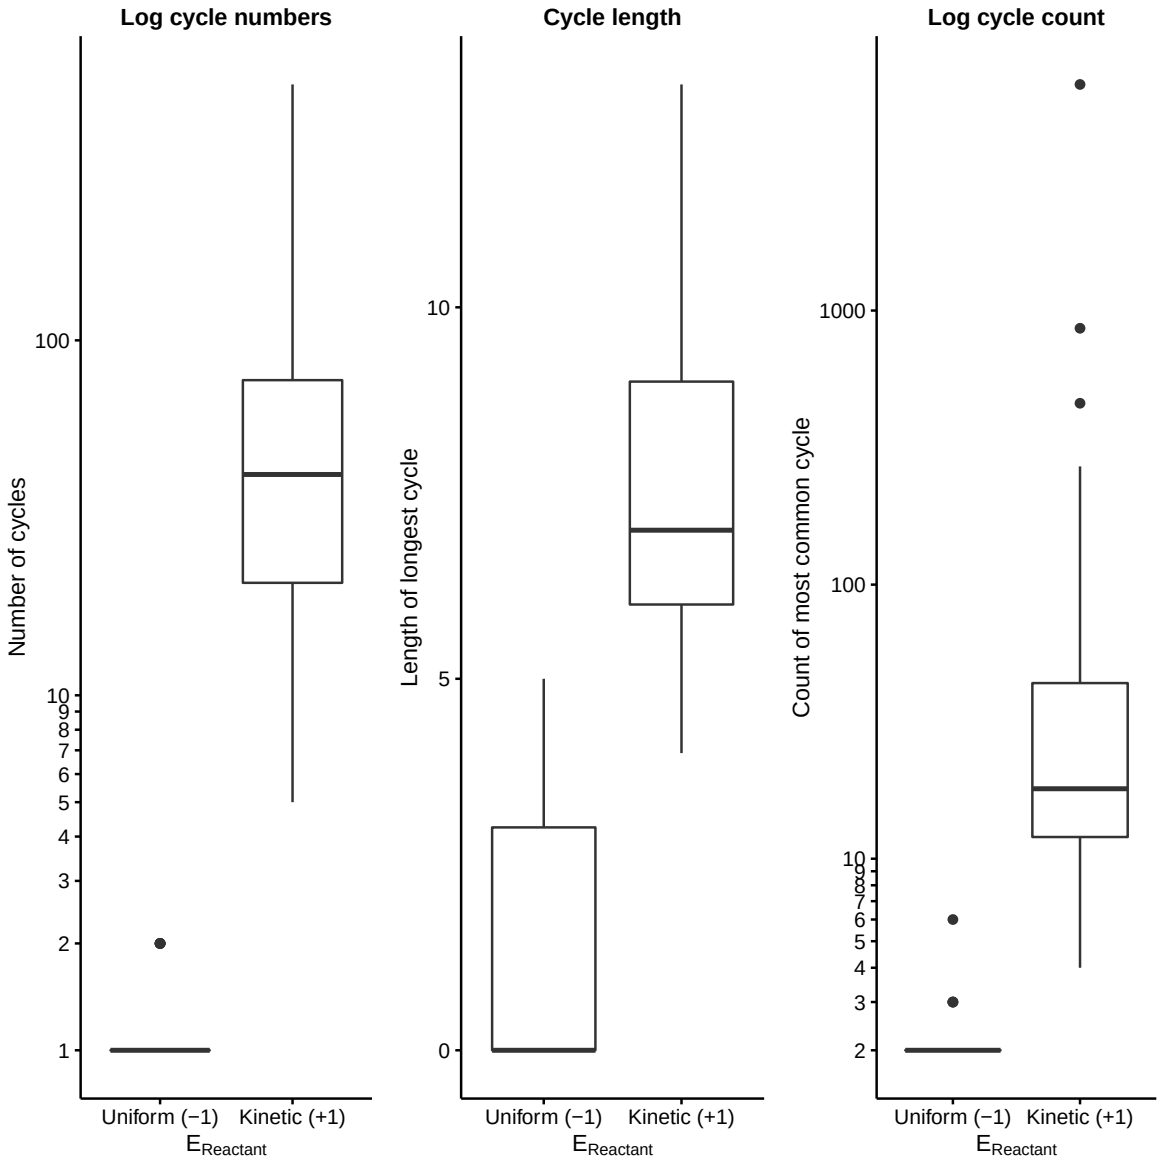

Figure 4.9: Response by E_{Reactant} (Reactant Strategy) (y-axis scales vary)

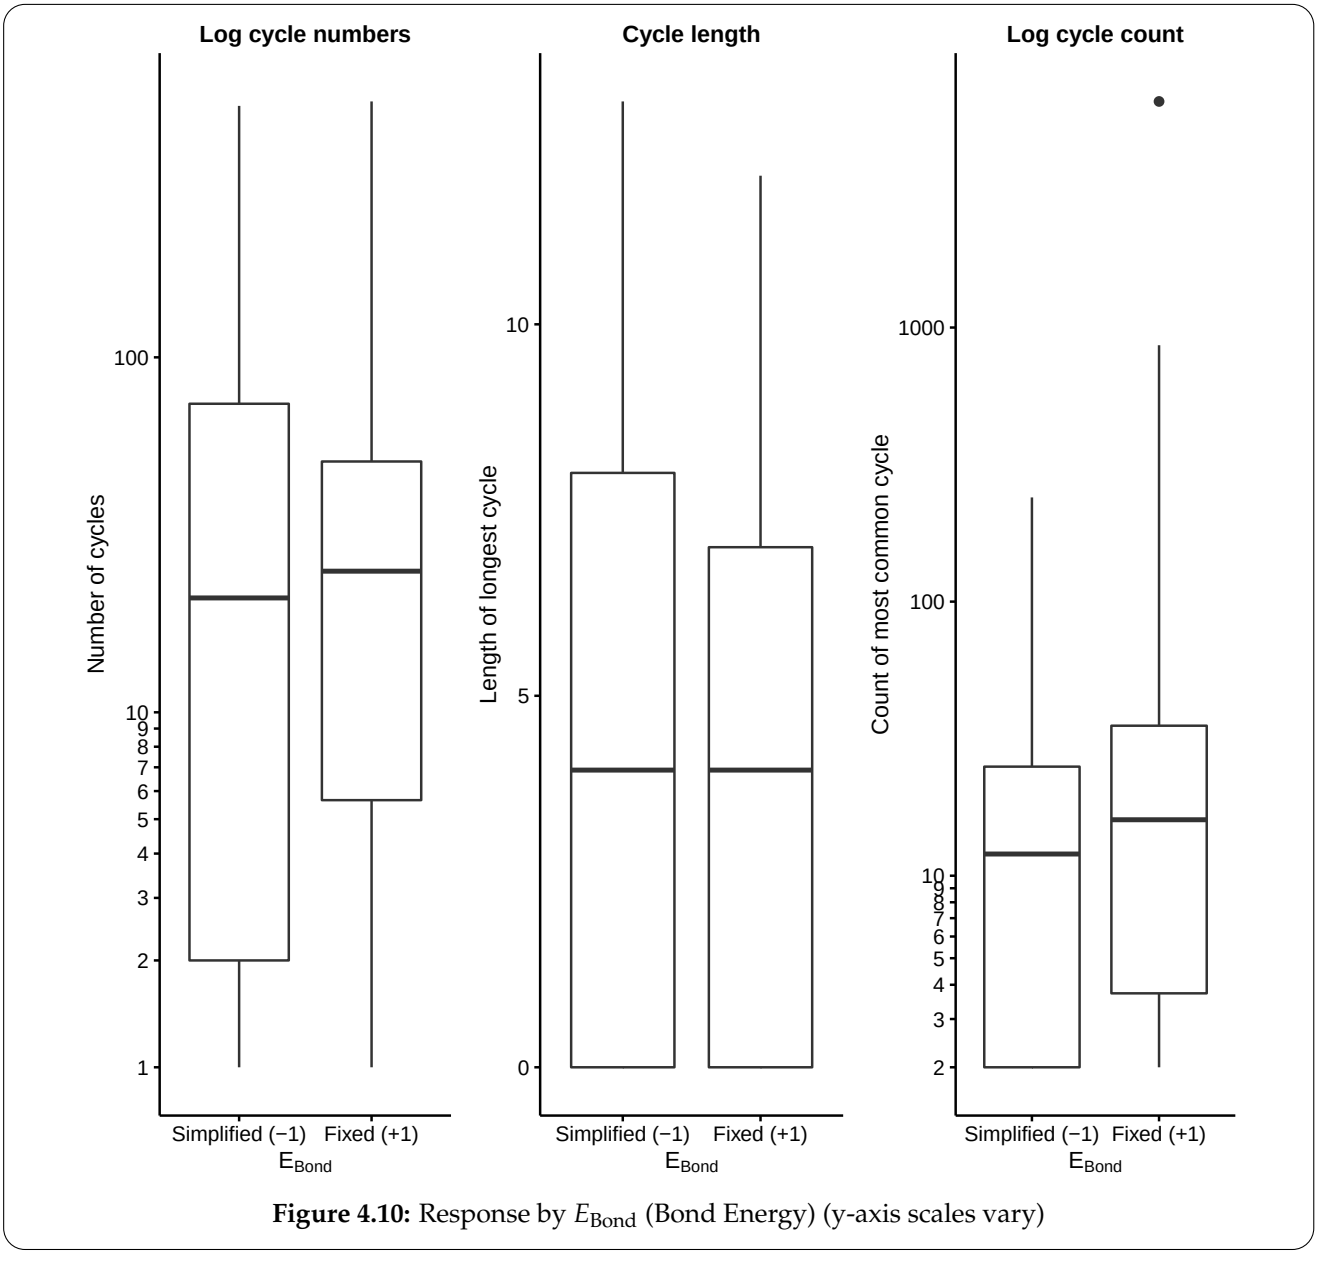

Figure 4.10: Response by E_{Bond} (Bond Energy) (y-axis scales vary)

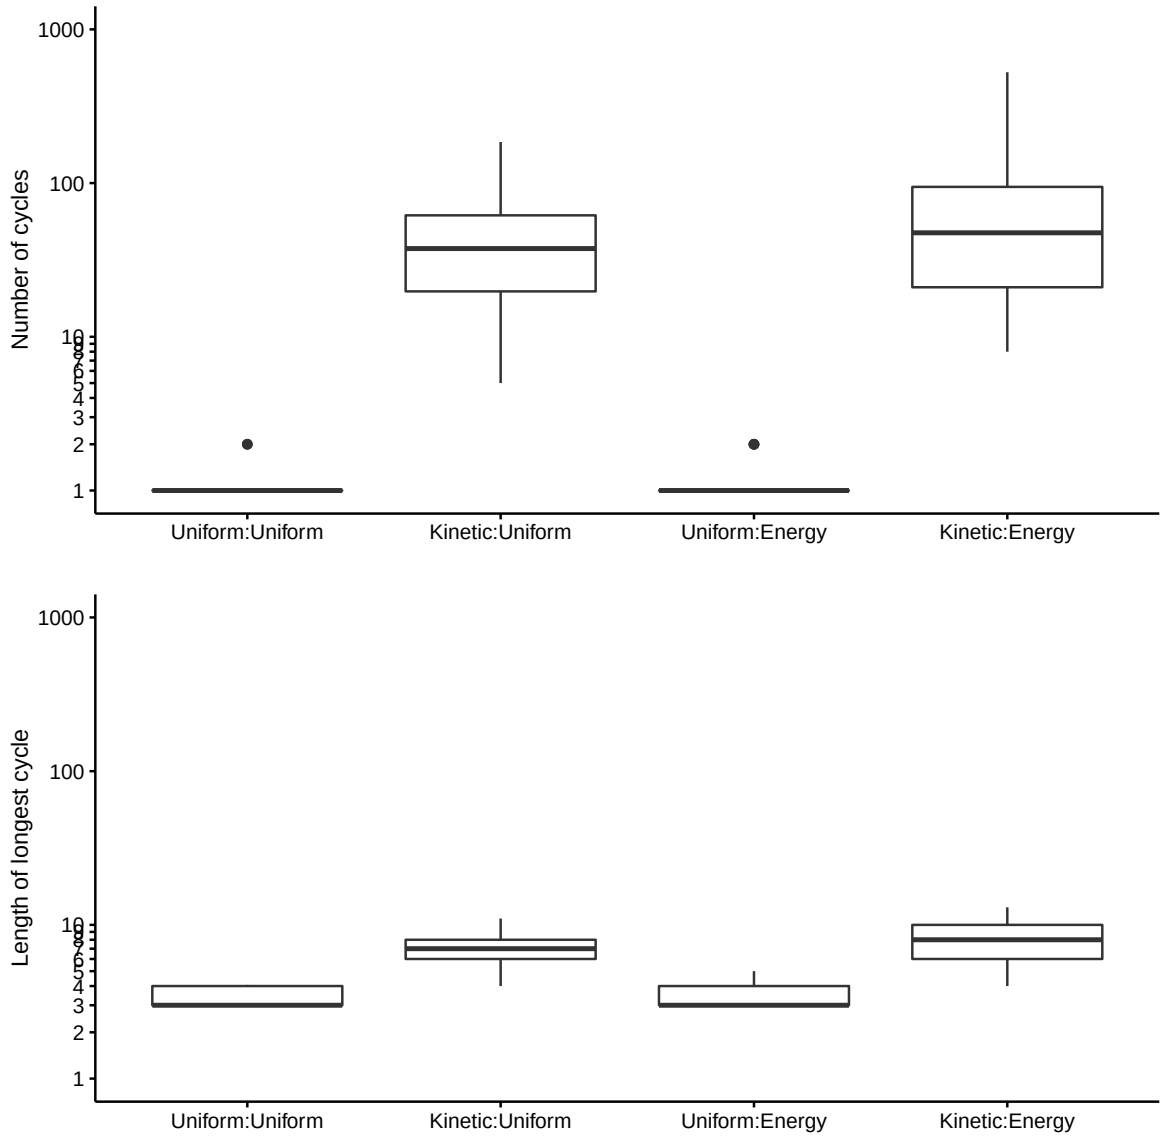

Figure 4.11: Effect of the combination of Reactant and Product Strategies on the number of cycles (top), and length of longest cycle (bottom.)

5 The Emergence of Multipliers

We now move on to the question of multiplication in artificial chemistries. In section 2.2, we saw that the ability to multiply or increase in number is fundamental to all common formulations for evolution. In Zachar and Szathmáry (2010), multiplication is the first of the main discriminators beyond embodiment, or the state of being, (see fig. 1.2) and so in this chapter we continue our investigation of evolutionary emergence in artificial chemistries by identifying multiplying entities.

From the earlier definition of replication (section 1.2.2), given by the autocatalytic reaction $\Sigma x_i + A \rightarrow \Sigma y_j + \Sigma B_k$, A is a multiplying entity if A and some $n < k$ of the B_k are equivalent under selection (Zachar and Szathmáry 2010).

As a secondary goal, we are also interested to learn if our assumption from chapter 4, that simple measures of cycle formation in a chemistry can predict the subsequent emergence of replicators, holds at this next level of investigation.

5.1 Choice of entity

As we are concerned with reaction networks, reaction cycles are a plausible candidate for an entity formed from reactions. In section 4.7 we introduced these cycles as a partially-closed sequence of reactions capable of autocatalysis.

Although most works influenced by origins-of-life objectives both require and assume catalysis as the mechanism by which molecules can rapidly increase in concentration, as has been pointed out before (King 1978 and recently Virgo and Ikegami 2013) if the goal is increased concentrations then simple stoichiometry (reaction dynamics) can also suffice. The major advantage of stoichiometry over catalysis for artificial evolution is that no additional mechanism is required other than existing reaction dynamics; catalytic systems however must somehow determine the catalytic effect of a molecule upon a particular reaction, and without a full quantum chemical model this is inherently prone to assumptions. For these reasons we shall not base our entity definition upon catalysis, but instead on the stoichiometry of sequences of reactions.

A reaction cycle is not defined by the sequence of individual molecules linked by reactions, as individual molecules are unique and transitory and so cannot form cycles, but rather by a sequence of molecular species. Molecular species define disjoint groups to which individual molecules of the same chemical composition belong: for example, the 'OH' species that includes all individual 'OH' molecules. In this way the cycle itself may be quite long-lived even though constituent molecules are produced and consumed.

Other candidate entities are available though, such as molecular species, or interrelated groups (or sets) of reaction cycles. Why then should we choose reaction cycles instead of any of these other alternative levels of entity? The definition of multiplication (given above) is of little direct assistance: multiplying entities may be of any form.

But if we consider the context of multipliers within evolutionary systems (from section 2.2), entities that evolve are multipliers that can vary under selection. This establishes a useful lower bound on the level of any candidate type of entity. For example, our entity must be capable of variation without changing its inherent nature. This is necessarily less definite for chemical elements, but it seems clear that molecules at least are not capable of variation without becoming a different type of molecule (adding a carbon atom to a molecule does not go unnoticed.)

At the other extreme, requiring multipliers to be full informational replicators is clearly overly restrictive. Holistic replicators are not informational replicators, yet no one disputes that they are multipliers.

For our purposes, the most useful or valuable class of entity when investigating multiplication will be the simplest type that possesses the capacity for variation and selection: single reaction cycles.

5.2 Measures

For exact multiplication, we seek reaction cycles where a cycle results in two or more cycles of the same species. Although one product cycle would be sufficient for replacement, as has been argued elsewhere (*e.g.* Zachar and Szathmáry (2010)), replacement alone without excess production exposes an entity to depletion from side-reactions and other forms of degradation.

For template autocatalysis (section 3.1) of cycles, where the autocatalysis results from stoichiometry, the connecting molecules must also have a stoichiometry $n > 1$. For exact multiplication we can derive the cycle stoichiometry (k) for exact multiplication by substituting A for B in the above autocatalytic expression to obtain $\Sigma x_i + A \rightarrow \Sigma y_j + \Sigma A_k$. Now as molecules are unique, each connecting molecule will appear once as product of one A cycle and once as a reactant of one other A cycle. Therefore $k = n$ and as $k > 1$ for template autocatalysis, $n > 1$.

What exactly is meant though by this stoichiometric relationship? When discussing non-physical replicators, such as memes, this relationship must also necessarily be non-physical and hence the general definition of replication cannot be restricted to a physical linkage between left- and right-hand sides of the relationship arrow. Our molecules and reaction cycles are simulations, not embodied in the real-world, but we can assume a simulated physical linkage between parent and offspring entities: we require that they are connected by one or more shared molecules (individual molecules, not species.)

Combining these two elements allows us to define an exact multiplier. An exact multiplier consists of:

- *Two or more* copies of the *same reaction cycle species*, where
- The reaction cycle species has *stoichiometry greater than one*, and
- Where *each cycle in the multiplier is connected to at least one other multiplier cycle* by a molecule that is a product of one cycle and a reactant in the other.

The algorithm for detecting these multipliers is given in alg. 7. The functions $IdentifyCycles$, to detect cycles where at least one product molecule is produced in excess quantities ($n > 1$ from above), and $IdentifyClusters$, for clustering cycles of the same form that are interconnected by shared molecules, are defined in alg. 8 and alg. 10 respectively.

5.3 Cycle detection in a reaction network

Choosing single reaction-cycles as entities in the hypothesis means that cycle detection is central to our analysis. The time complexity of standard cycle detection algorithms is generally a factor of the number of nodes and the number of edges. For example, Johnson's algorithm for elementary circuits has time complexity O((n+e)(c+1)) for n nodes, e edges and c elementary circuits (Johnson 1975). This is problematic when applied to each of our reaction graphs, where $n = 100,000$, $e = 200,000$, and a typical value for c might be somewhere between 10,000 and 60,000 (as a molecule can be part of multiple cycles.) Detecting all cycles is not only slow, but generates very large files. Tests against some representative reaction networks results in lists of cycles that are larger than a workstation can hold in memory without swapping; this has a commensurate effect on analysis times.

An algorithm to detect all cycles (where the sampling proportion is 1.0 in alg. 8) is expensive, and analysis of datasets containing 50,000 or more runs can take one or more days using this algorithm on a modern workstation. Furthermore, standard cycle detection algorithms require adaptation as our cycles are completed not by returning to the same node, but by visiting any node of the same species as a cycle consists of molecule species while the graph is by molecule instances and therefore (by definition) acyclic. The approach

we take instead is to relax the requirement to identify all cycles in favour of a sampling or lower-bounds approach, suggested by an analogous strategy in Hordijk, Smith, and Steel (2015, p.7). We sample a proportion of all reactant molecules and search the reaction graph only for cycles incorporating one or more of the sampled molecules.

There is a significant implication to this approach. Because we are sampling "seed" molecules, the cycles identified will be clustered around the seed molecules in the reaction graph, while intervening sections of the reaction graph will remain unexplored; any chains of cycles may be broken at these unexplored regions. This can be mitigated, although not prevented, by careful consideration of the sample size. Additionally, as the number of cycles detected is a factor of the proportion of seed/non-seed molecules, if that proportion is constant between graphs then the relative numbers of cycles detected in each graph will also be constant (as shown on the left-hand of fig. 5.1.) This means that relative comparisons will hold true regardless of sampling proportion.

When it comes to the number of multiplier species detected by our sampling algorithm, the sampling proportion p has a non-intuitive effect (right-hand of fig. 5.1.) As the proportion increases beyond approximately 0.75 (from the plot) the number of species detected drops, instead of continuing to increase as expected. The reason for this lies in the inclusion of stoichiometry in the definition of a multiplier—multipliers must stochiometrically increase, or in other words, the number of component cycles that produce an excess number of products must exceed the number that do not. As most of the multiplier species detected consist of a low number of component cycles, the detection by sampling of even a single additional cycle to the multiplier can change the stoichiometry of the overall collection of cycles in the multiplier. This is biased towards changing the stoichiometry from increasing to neutral or decreasing simply because of the possible combinations of cycles in the multiplier, where the multiplier consists of a small number of cycles. This explains the drop in fig. 5.1.

The experiments in this chapter all use a sampling proportion $p = 0.2$, as a balance between performance and predictability (as very small values for p produce noisy results) although this will likely undercount the true number of multipliers due to the sampling effect previously mentioned.

5.4 Experiment design

We can now state the working hypothesis of this chapter as:

Hypothesis 5.1. *Exact multipliers, in the form of entities composed of a repeated reaction cycle, can emerge in an artificial chemistry.*

Multipliers are structures that, in our context of artificial chemistries, emerge from the

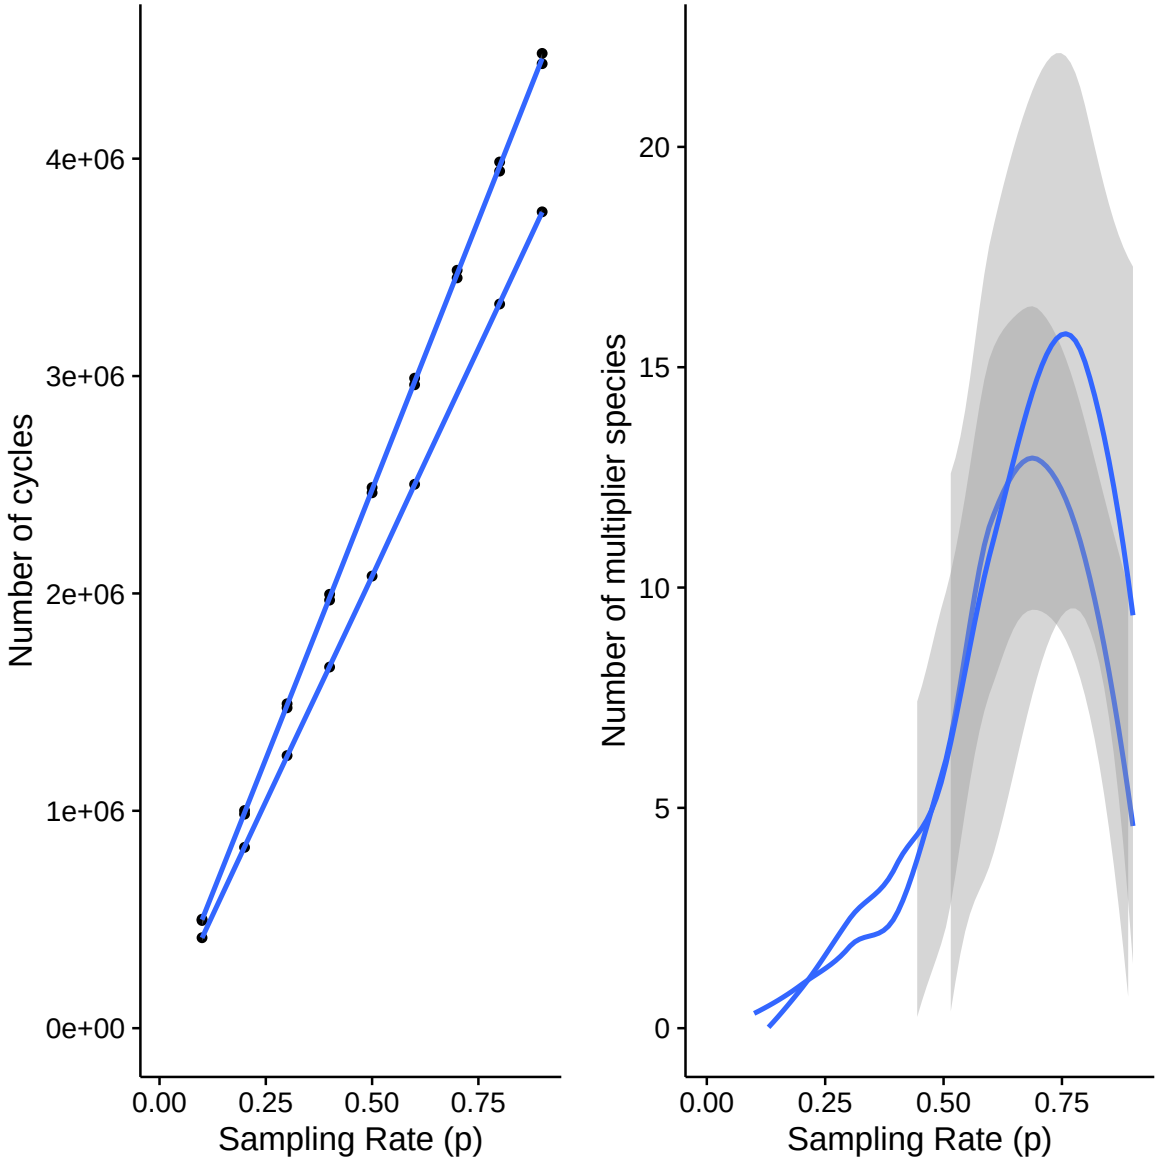

Figure 5.1: Comparison of number of cycles (left) and number of multipliers (right) for different sampling rates (p, in text). Note the highly linear relationship between sampling rate and number of cycles. Values are grouped by source dataset; two datasets (1489554358 and 1489565574) result in two grouped sets of data in each plot. Three replicates chosen with uniform probability for analysis. See text for discussion of drop in multiplier species at high p.

```
1  Function DiscoverMultipliers(Reactions):
2      Multipliers ← ∅
3      MolecularCycles ← IdentifyCycles(Reactions)
4      MolecularCyclesBySpecies ← MolecularCycles grouped by cycle species
5      for CyclesForSpecies in MolecularCyclesBySpecies do
6          Clusters ←IdentifyClusters(CyclesForSpecies)
7          for Cluster in Clusters do
8              AllProducts ← ⋃_{x∈Cluster} Products(x)\Reactants(x)
9              AllReactants ← ⋃_{x∈Cluster} Reactants(x)\Products(x)
10             LinkingMolecules ←AllProducts ∩ AllReactants
11             if LinkingMolecules≠ ∅ then
12                 ProductCycles ← {x∈Cluster | Products(x) ∩ LinkingMolecules ≠ ∅}
13                 ReactantCycles ← {x∈Cluster | Reactants(x) ∩ LinkingMolecules ≠ ∅}
                   // Check overall stoichiometry is greater than 1.0
14                 if |ReactantCycles| > |ProductCycles| then
15                     for Cycle in Cluster do
16                         if (products of cycle ∩ LinkingMolecules ≠ ∅) ∨ (reactants of cycle
                             ∩ LinkingMolecules ≠ ∅) then
17                             Multipliers ← Multipliers ∪ Cycle
18                         end
19                     end
20                 end
21         end
22     end
23     return Multipliers
```

Algorithm 7: *DiscoverMultipliers*. Detect multiplying exact replicators within a reaction network.

```
1  Function IdentifyCycles(Reactions):
       // Construct reaction graph from a list of Reactions
2      for Reaction in Reactions do
3          Add a node for the reactant side of Reaction
4          Add a node for the product side of Reaction
5          Add an edge from reactant node to product node
6          for Reactant in Reaction do
7              Add node for Reactant and edge from Reactant to node for reactant side
                of Reaction
8              Reactants ← Reactants ∪ Reactant
9          end
10         for Product in Reaction do
11             Add node for Product, and edge from Product to node for product side of
                Reaction
12         end
13     end
       // Find all cycles for a sample of reactants in the graph
14     Cycles ← ∅
15     Seeds ← sample with uniform probability some proportion p of Reactants
16     for Seed in Seeds do
17         ShortestPaths ← FindCyclesFromNode(Seed)
18         for Path in ShortestPaths do
19             if the stoichiometry for Seed in the cycle given by this Path≥ 2 then
20                 Cycles ← Cycles ∪ Path
21         end
22     end
23     return Cycles
```

Algorithm 8: *IdentifyCycles*. Sampling algorithm to identify autocatalytic cycles of molecules within a reaction network. A cycle is a sequence of reactions where at least one product molecule in the final reaction is of the same species as a reactant molecule in the first reaction of the cycle. The cycle is considered autocatalytic for a given product molecule if the stoichiometry for the molecule is greater than one, or in other words, if the cycle produces more of the product than is consumed. The sampling proportion p determines the proportion of the total set of reactants to be used as seeds when searching for cycles.

```
1  Function FindCyclesFromNode(Source):
2      Target ← species of Source
3      Stack ← [(Source, [Source])]
4      while Stack do
5          (Vertex, Path) ← Stack.pop()
6          for NextNode in predecessors(vertex)\path do
7              if NextNode is a molecule ∧ species of NextNode = Target then
8                  Yield reversed(Path + [NextNode])
9              end
10             else
11                 if |Path| < MAXDEPTH then
12                     Stack.append((NextNode, Path + [NextNode]))
13                 end
14             end
15         end
16     end
```

Algorithm 9: *FindCyclesFromNode*. Recursive breadth-first algorithm to identify all shortest paths (and hence cycles) from *Source* to any predecessor molecule of the same molecular species as *Source* in the directed reaction graph.$MAXDEPTH$ is a constant to bound the cost of finding cycles.

reactions in a reaction network; given the nature of emergence, to determine if an artificial chemistry can result in multipliers, we turn to experiment. The hypothesis test is as follows:

H_0: No multiplying single reaction-cycle entities exist in any reaction network generated by an artificial chemistry.

H_1: Multiplying single reaction-cycle entities exist in an artificial chemistry reaction network.

We again use the ToyWorld parameterized evolutionary model introduced in chapter 4. The experiments follow a full factorial design although with only a single factor (S_{Product}) at two levels (Uniform and LeastEnergy), while S_{Reactant} is set to Kinetic (alg. 5). The experiments run in a randomised order, with two (2) replicates of each combination of factors executed in sequence before beginning the next combination. The factors and levels used are given in table 5.1.

Each replicate runs for 100,000 iterations. The initial population for each replicate contains 10×[H][H] molecules, 10×FO, 20×O, 10×[O-][N+](=O)[N+]([O-])=O, 10×N(=O)[O], and 20×O=C=O (all in SMILES notation), and the initial reactor kinetic energy (E_{Vessel}) is set to 100 units.

```
1  Function IdentifyClusters(MolecularCyclesForSpecies):
2      Clusters ← ∅
3      Unclustereds ← ∅
4      for CycleMolecules in MolecularCyclesForSpecies do
5          CanCluster ←False
           // First check if part of any existing cluster
6          for Cluster in Clusters do
7              if Cluster ∩ CycleMolecules then
8                  Cluster ←Cluster ∪ CycleMolecules
9                  CanCluster ←True
10                 break
11         end
           // Otherwise see if can form a new cluster with a previously
              unclustered cycle
12         if ¬CanCluster then
13             for Unclustered in Unclustereds do
14                 if Unclustered ∩ CycleMolecules then
15                     Clusters ←Clusters ∪ new cluster of [Unclustered ∪ Cycle]
16                     Unclustereds ←Unclustereds\Unclustered
17                     CanCluster ←True
18             end
           // If still can't cluster, then add to Unclustereds
19         if ¬CanCluster then
20             Unclustereds ←Unclustereds ∪ Cycle
21     end
22     return Clusters
```

Algorithm 10: *IdentifyClusters*. Identification of clusters: reaction cycles of the same form interconnected by shared molecules where each such molecule is produced in one cycle in the cluster, and consumed in another. The parameter *MolecularCyclesForSpecies* contains all cycles of the same cycle species, that is, all cycles that consist of the same sequence of reactions when written in SMILES (species) form.

Table 5.1: Factors, or independent variables, for experiments in chapter 5.

Factor	-1 value	+1 value	Description
S_{Product}	Uniform	LeastEnergy	See section 4.5

5.5 Results and discussion

Table 5.2: Summary results for numbers of multipliers measured by alg. 7 in reaction networks generated by LeastEnergy and Uniform S_{Product} strategies, under stable environments (no external selective pressure). Sampling proportion p in alg. 8 shown in column labels.

Dataset	Experiment	Environment	Replicate	S_Product	No. of Species (p=0.05)	Multipliers(p=0.05)	Average Multiplier Size.x	No. of Species (p=0.2)	Multipliers(p=0.2)	Average Multiplier Size.y
1489554358	0	0	0	LeastEnergy	0	0	0.00	2	2	3.00
1489554358	0	0	1	LeastEnergy	0	0	0.00	0	0	0.00
1489554358	1	0	0	LeastEnergy	0	0	0.00	1	1	3.00
1489554358	1	0	1	LeastEnergy	0	0	0.00	0	0	0.00
1489554358	2	0	0	LeastEnergy	0	0	0.00	1	1	3.00
1489554358	2	0	1	LeastEnergy	0	0	0.00	0	0	0.00
1489554358	3	0	0	LeastEnergy	0	0	0.00	0	0	0.00
1489554358	3	0	1	LeastEnergy	0	0	0.00	0	0	0.00
1489565574	0	0	0	Uniform	0	0	0.00	1	1	3.00
1489565574	0	0	1	Uniform	0	0	0.00	2	2	3.00
1489565574	1	0	0	Uniform	0	0	0.00	1	1	3.00
1489565574	1	0	1	Uniform	0	0	0.00	1	1	3.00
1489565574	2	0	0	Uniform	0	0	0.00	0	0	0.00
1489565574	2	0	1	Uniform	0	0	0.00	0	0	0.00
1489565574	3	0	0	Uniform	0	0	0.00	1	1	3.00
1489565574	3	0	1	Uniform	0	0	0.00	1	1	3.00

Identification of cycles within the reaction network (our hypothesised reaction-cycle entities) is by alg. 8, with a sampling proportion $p = 0.20$ ($p = 0.05$ also shown for comparison), and alg. 7 applied to identify any multipliers within those cycles.

From table 5.2 we can see that exact multipliers arise in approximately half of the experiment runs. Most of those runs produced only one type (species) of multiplier; the remaining runs with multipliers produced two types. Each type existed only once in a run (so the number of multiplier entities equals the number of species), and the number of cycles in each multiplier was again quite low: 3 cycles only. One of the resulting multipliers is shown in fig. 5.2, constructed from three instances of the same cycle species ('Y', see table A.3 for other species descriptions.) Exact multipliers do occur, but infrequently, and they are not long-lived.

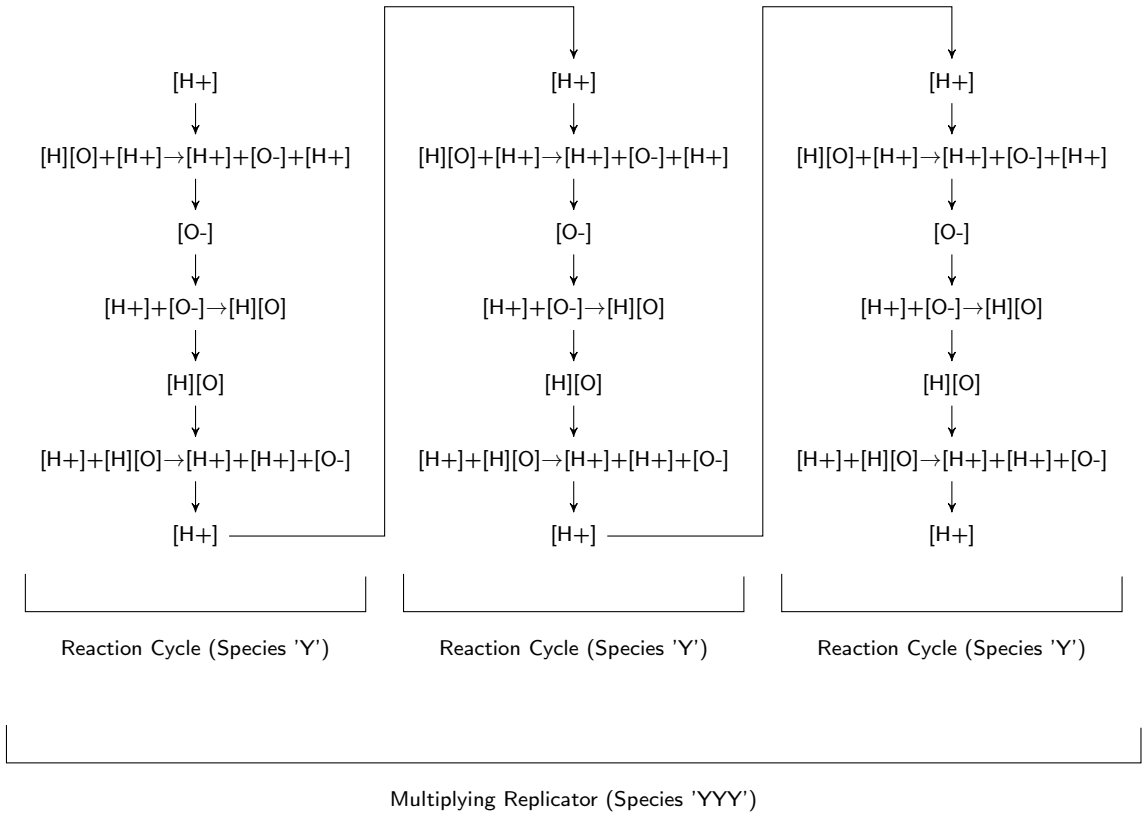

Figure 5.2: Example multiplier from experimental results (of type or species 'YYY', see table A.3). A multiplier is formed from two or more of the same species of reaction cycle, connected by shared molecules. For clarity, the only molecules shown are those that form part of the multiplier cycle (all others can be easily discovered from the reaction equations.)

5.5.1 Do multipliers arise simply by chance?

Perhaps multiplier emergence is simply a product of chance, and that given enough iterations, probabilities will favour two or more of the cycles that form being linked by molecules in a way that meets our definition of a multiplier. In this view, multipliers are inevitable in an artificial chemistry because the definition is not sufficiently stringent to exclude unsurprising events.

This alternative explanation is not incompatible with our results. The average multiplier size—the number of cycles in each instance of a multiplier—is quite low at 3 cycles per lineage. That would suggest that the multipliers are not particularly stable; that they form, and then after a short period, dissipate.

This can be tested by comparing the number of replicators in these experiments with

those generated by a control experiment with a neutral bias.

The control uses the Uniform S_{Product} and S_{Reactant} strategies to chose reaction products and reactants (respectively) with uniform probability, and hence any multipliers will be those that result from chance alone. The statistical hypothesis is as follows:

H_0: $x_i = x_j \forall i, j$, where x indicates the presence of multipliers: $x_i = |multipliers_i| > 0$.

H_1: $x_i \neq x_j$ for some i, j.

Table 5.3: Numbers of multipliers from 10 replicates of control experiment. Multipliers identified by alg. 7, in reaction networks generated by Uniform S_{Reactant} and S_{Product} strategies, under stable environments (no external selective pressure). Sampling proportion p in alg. 8 set to 0.2.

Dataset	Experiment	Environment	Replicate	S_Product	No. of Species	Lineages	Average Multiplier Size
1489951262	0	0	0	Uniform	0	0	0.00
1489951262	0	0	1	Uniform	0	0	0.00
1489951262	0	0	2	Uniform	0	0	0.00
1489951262	0	0	3	Uniform	0	0	0.00
1489951262	0	0	4	Uniform	0	0	0.00
1489951262	0	0	5	Uniform	0	0	0.00
1489951262	0	0	6	Uniform	0	0	0.00
1489951262	0	0	7	Uniform	0	0	0.00
1489951262	0	0	8	Uniform	0	0	0.00
1489951262	0	0	9	Uniform	0	0	0.00

The results are shown in table 5.3; no replicate of the control experiments generated multipliers. Testing the statistical hypothesis with a two-sample, two-tailed, Student's t-test gives a t-statistic of 7.3508989 with $p = 9.1293962 \times 10^{-11}$, and so we reject H_0 (multipliers result from chance alone) in favour of the original hypothesis, that multipliers emerge from biases in the artificial chemistry. Although the sampling proportion p undoubtedly affects the total number of multipliers identified in each experiment, the relative numbers between experiments are maintained and hence we do not believe that this result is solely an artifact of sampling.

5.5.2 Is the number of reaction cycles a good predictor for the number of multipliers?

Returning to the underlying assumption of chapter 4, is there a relationship between the number of cycles and the number of exact multipliers? If there is, then the parameters for

an artificial chemistry designed to investigate replicators (of which exact multipliers are the simplest form) could be established based simply on a count of cycles calculated from some calibration runs. If not, then cycle counts cannot be used as a guide for parameter choices, and in the absence of any other predictor, all parameter combinations must be considered in any replicator experiment.

The test hypothesis is as follows:

H_0: $|Multipliers| \lesssim |Cycles|$ for all parameter combinations.

H_1: $|Multipliers|$ cannot be predicted from $|Cycles|$.

A scatterplot of the relationship (fig. 5.3) shows only a weak relationship between cycles and multipliers. The association between the two variables is positive, and the form appears to be essentially linear. However, the correspondence with the simple linear regression is poor, and the predictive value of cycles appears weak.

We examine this further by testing the predictive value of the relationship for two specific datasets.

From the previous section we have the example of a dataset (1489951262) in which the number of multipliers is zero in all runs, and the number of cycles is non-zero (4.17165×10^5.) However, in another dataset (1489554358) both the number of multipliers (4.7027027) and the number of cycles (2.5188076×10^6) are non-zero.

We reject H_0 as the relationship between cycles and multipliers in these two datasets is complex, and conclude that the number of multipliers cannot easily be predicted from the number of cycles. Our original supposition, that we could use a simple count of cycles to select the parameters, and parameter values, for an artificial chemistry to investigate replication is not supported by these results.

5.6 Conclusions

We began this chapter with an examination of the most appropriate form of entity for an examination into multiplication in artificial chemistries. The same reaction network contains a variety of candidate entities arranged in a compositional hierarchy. At the base level of the artificial chemistry are atoms, formed by reactions into molecules. Reactions and molecules can be grouped into various other forms: reaction cycles, where the products of the final reaction in the cycle go on to become reactants in the first reaction of the next iteration of the cycle, and the various forms of reaction sets, such as the autocatalytic sets introduced earlier in this work. And of course groups of reactions and molecules can likely form higher level entities in turn. We selected reaction cycles as the appropriate entity form as the simplest of the candidate forms capable of both variation and selection.

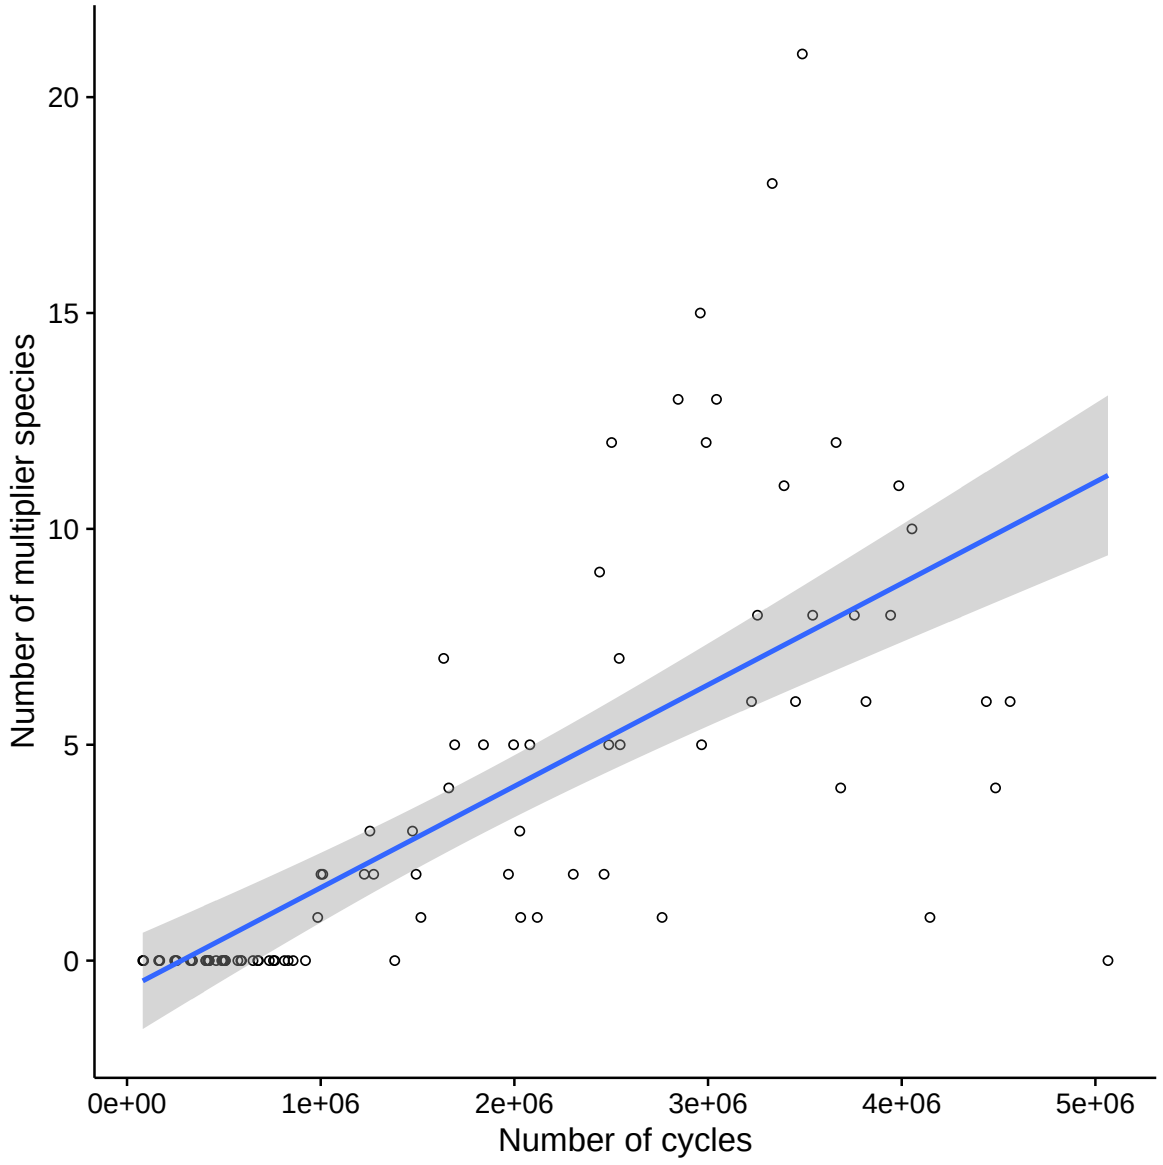

Figure 5.3: Relationship between the number of cycles and the number of multipliers for a sample of replicates from datasets 1489951262 (horizontal line at bottom-left), 1489554358 and 1489565574 including linear regression line and 95% confidence region.

Next we considered the definition of exact multiplication, and concluded that an exact multiplier consisting of reaction cycles must meet three conditions:

1. The multiplier must consist of two or more copies of the same reaction cycle species.
2. The reaction cycle species must have stoichiometry greater than one.
3. Each cycle in the multiplier must be connected to at least one other multiplier cycle by a molecule that is a product of one cycle and a reactant in the other.

However, to identify multipliers using this definition we must first identify reaction cycles in the reaction graph and exhaustive enumeration of these cycles is impractical. Instead we introduced a sampling algorithm that discovers cycles by searching from a subset of "seed" molecules, chosen with uniform probability from the full set of reactant and product molecules in the reaction graph. The sampling algorithm discovers cycle numbers in rough proportion to the sampling proportion p, and we showed that this is sufficient for multiplier discovery.

Finally, we tested the hypothesis that multipliers of the form described can arise in an artificial chemistry by experiment using the ToyWorld artificial chemistry from chapter 4. In summary, exact multipliers do arise in the ToyWorld artificial chemistry, but not in any great numbers, and when they do, they do not persist for long. We have shown though that they do occur as the result of a non-neutral combination of Product and Reactant selection strategies, and not purely by chance. It is also apparent that our earlier supposition, that the number of reaction cycles would be a good proxy or predictor for multiplier or replicator activity, is not supported by the evidence.

6 From Multiplication to Variability

Having shown the existence of exact multipliers in the previous chapter, we now move on to considering the implications of selective pressure and variability on multiplication. First, we replace the stable environment of the previous chapter by varying selective pressure and examine the effect that that has upon the emergence of multipliers. Second, we investigate the emergence of variable replicators.

A variable replicator is a multiplier that can change its structure while remaining the same under selection (Zachar and Szathmáry 2010). The selection requirement provides an objective way of discriminating between changes that maintain continuity between different forms or states of the same fundamental entity from those changes that result in something intrinsically different. Although there is no obvious limit to the number of possible states that a variable replicator may take, for pre-template replicators the maximum number is constrained by the ability of the entity to sustain non-interfering chemical networks (see introduction in section 1.2.3.) In other words, pre-template variable replicators are multipliers that can occupy any of a limited set of states without losing their underlying identity.

For the studies in this chapter we require some extensions to our earlier models and algorithms. Given the association between variable replicators and explicit selection, we need to adopt or develop a model for environmental change. We also need to extend our method for detecting multipliers to detect variable replicators in the reactions generated by our evolutionary model. These two extensions are covered in section 6.2 and section 6.5 respectively. The conceptual basis for these studies is the same as introduced in fig. 2.2: a base evolutionary model under extrinsic selective pressure provided by an environmental model.

Finally, we conduct experimental investigations of multipliers under varying selective pressure, and then variable replicators in our artificial chemistry. A common experiment design is introduced in section 6.3, followed by separate result and discussion sections for each study.

6.1 Previous work

Perturbations are applied to a well-mixed reaction vessel by Fontana and Buss (1994) with objects taken from λ-calculus. Organisations (entities, hypercycles) in all cases are observed to be self-maintaining but not self-reproducing (excluding cases that included an explicit replication expression.) The model is without locality or spatial effects; perturbations take the form of adding a small number of copies (10 or 50) of a random object every 30,000 or 50,000 collisions. The purpose of the perturbations however is simply to assess the sensitivity of the model.

A series of successively enhanced models—Polymer-GARD (P-GARD) (Shenhav et al. 2005), Environment Exchange Polymer-GARD (EE-GARD) (Shenhav, Oz, and Lancet 2007), and Universe-GARD, proposed in Markovitch et al. (2012)—address extending the range of composomal inheritance. P-GARD lacks a mechanism for environmental variation, although EE-GARD, which builds upon P-GARD, moves towards such a mechanism by returning half of the products of composome split or fission events to the environment. This return mechanism is the main justification for the claim in Shenhav, Oz, and Lancet (2007, p1819) that the composomes and environment co-evolve. Universe-GARD further extends the importance of variation of environment; the population of the main GARD environment is continually exchanged with that of an enveloping "universe" that contains significantly more molecular species. No results have yet been reported however to the best of our knowledge as Markovitch et al. (2012) describes the model only.

Regular perturbations are at the core of the liposome model of Fernando and Rowe (2007), where periodic "avalanches" of novel side-reactions spawn new and rare reaction products to inject variation into the simulation, either by making new reactions possible, or by shifting the balance between competing reactions, or by perhaps altering the reaction rates through changes to catalysis. Avalanches are described as an exploration mechanism, extending the range of the simulation into otherwise unlikely regions, although no experimental assessment of their effectiveness was performed. The focus of the work primarily lies in the hypothesis that liposomes (significantly more complex than our reaction cycles) form a plausible candidate for early forms of life, kick-started towards the evolution of novelty by chemical avalanches.

6.2 Sources of selection pressure

Although selective pressure can be purely intrinsic, for example in the competition between reaction cycles for molecules (Eigen 1971), this is inherently more difficult to model and to control than explicit selective pressure, supplied by an external agency. In these experiments we concentrate solely on explicit selective pressure generated by a changing

environment.

A *changing* environment is one in which the relative phenotypical fitness is time-dependent, while, conversely, a *stable* environment is one where the relative fitness of a phenotype with respect to the environment is independent of time[1]. In the models in this chapter, environmental change only occurs between generations, and the phenotype of an entity remains constant over its lifespan. All fitness changes are caused by environmental changes; we do not consider fitness changes that are the result of processes that operate on a sub-generational timescale such as learning or development.

Any model of changing environments must capture two particular dimensions of change: the target of the change, and the shape of the change. The *target* of change refers to the aspect of the evolutionary model that is affected by the change. Any model parameter or combination of parameters might be a target. For example, in the evolutionary model in chapter 6, any or all of the parameters S_{Reactant}, S_{Product}, E_{Vessel} and E_{Bonds} could be targets of environmental change. The *shape* of change describes the way change varies over time, as a series of δ at timestep t, where $\delta_t \in R$.

In previous work, Jablonka et al. (1995) and Paenke et al. (2007) share a simple model for environmental change, with instantaneous switches between two defined environments according to some schedule. Gaucherel and Jensen (2012) also describes a single model for environmental change—a repeated module of a period of "smooth" change (either according to a form of *sine* curve or unchanging) followed by an abrupt change—with variation in the length and degree of each period. By contrast, Schuster (2011, p. 79), when describing the relationship between fitness landscapes and error thresholds, details five distinct models of change: "(i) the single-peak landscape corresponding to a mean field approximation, (ii) the hyperbolic landscape, (iii) the step-linear landscape, (iv) the multiplicative landscape, and (v) the additive or linear landscape."

Although the shape of change can clearly take an infinite number of different forms, we concentrate on two representative shapes, one adopted from Jablonka et al., and one novel:

- *Bistate* selective pressure, where the shape of change follows a pattern that switches between two alternate states at regular intervals.

- *AR-timeseries*, a new approach adopting the environmental model introduced in section 2.7.1 and described by three parameters Θ, σ_e and δ. Θ is the AR coefficient, e_t is a random, normally distributed, error component around a mean of 0, where $e_t \overset{iid}{\sim} N(0, \sigma_e^2)$, and δ is a fixed bias value. An AR-timeseries combines a regular component (described by Θ and δ) and an unpredictable component (σ_e.) Unlike the bistate case, the AR-timeseries is neither constant nor completely regular.

[1] This is similar to statistical stationarity: the phenotype's fitness is independent of when it is measured.

Table 6.1: Factors mapped to model parameters, plus factor levels, for all experiments in chapter 6.

Factor	Model	-1 value	+1 value	Description
S_{Product}	Evolutionary	Uniform	LeastEnergy	See section 4.5
E_{Shape}	Environmental	Bistate (alternate every 10,000 iterations between -20 and +20 values.)	AR-timeseries (applied every 1000 iterations.)	See section 6.2
E_{Target}	Environmental	KE	Population	See section 6.2

As for the target of change, we begin with two defined alternatives:

- *KE*, where each molecule's kinetic energy is increased by the change value of δ: $KE_{new} = KE + \delta$. As seen in chapter 4, the population composition is influenced by the kinetic energy available for reactions. Changing the KE of the reaction vessel should indirectly result in a change in reactor population, and hence influence the types of reactions and reaction cycles that develop.

- *Population*, in which the δ value is used to adjust the population size. If δ is negative, then $|\delta|$ molecules chosen with uniform probability are removed from the population; if δ is positive, then δ molecules chosen again with uniform probability from the initial population are added to the current population. In summary, environmental changes will directly affect the composition of the reactor population, once again influencing the relative proportions of reaction cycle species.

6.3 Experiment design

The experiment design for all experiments in this chapter is the same as the design used in chapter 5, with the addition of factors for the environmental model. All factors and levels are given in table 6.1. Environmental changes are applied according to the schedule in the same table.

6.4 Exact multipliers under varying selective pressure

First we explore the nature of any relationship between simple multiplication and selective pressure. We start by extending our investigation of exact multipliers in stable environments into the effects of selective pressure from varying environments, using the number of multipliers as our response variable for the effect.

We use two proxy metrics calculated from E_{Shape} as our measure of selective pressure: DFA and Sample Entropy. DFA is the calculated Hurst parameter (H) using Detrended

Fluctuation Analysis, while sample entropy is assessed by the method provided in Richman and Moorman (2000). Both are calculated from the timeseries as applied to the target, with changes at intervals of 1000 iterations for the AR-timeseries, and every 10,000 iterations for the bistate case. As a result the sample entropy in particular is significantly lower than the sample entropy of the source timeseries.

Our hypothesis test is:

H_0: There is no difference in the number of exact multipliers between different values of selective pressure.

H_1: The number of exact multipliers depends on the degree of external selective pressure.

6.4.1 Results and discussion

Raw results are given in table 6.2. Approximately 38% of runs produced multipliers (34 of 90 runs.) Almost all runs with multipliers contained 1, 2 or rarely 3 species, with the notable outliers of one run with 6 species and another with 13 (with the only parameter in common being E_{Target}=KE.)

If these two outliers are excluded, then the distribution of the number of exact multipliers against extrinsic pressure (DFA and Sample Entropy) is shown in fig. 6.1. From this it seems unlikely that a simple relationship between multiplication and selective pressure exists, confirmed by the results of an attempted linear regression. For the relationship between the number of multipliers and Sample Entropy the F-statistic=2.0001963 and the p-value=0.1608072; the fit is no better for DFA with the F-statistic=1.5314888 and the p-value=0.2195091.

From this we conclude that we cannot reject H_0.

Table 6.2: Multipliers measured by alg. 7 in reaction networks generated by LeastEnergy and Uniform S_{Product} strategies under forms of environmental change given by the columns *Target* and *Shape*. Sampling proportion p in alg. 8 set to 0.2. Sample Entropy scaled by a factor of 1000 (see section 6.3).

Dataset	Experiment	Environment	Replicate	$S_{Product}$	Target	Shape	DFA	Sample Entropy	No. of Species	Lineages	Average Lineage Size
1489554358	0	0	0	LeastEnergy	KE	BISTATE	0.00	0.00	2	2	3.00
1489554358	0	0	1	LeastEnergy	KE	BISTATE	0.00	0.00	0	0	0.00
1489554358	0	1	0	LeastEnergy	KE	BISTATE	0.00	0.00	0	0	0.00
1489554358	0	1	1	LeastEnergy	KE	BISTATE	0.00	0.00	0	0	0.00

Continued on next page

	Dataset	Experiment	Environment	Replicate	$S_{Product}$	Target	Shape	DFA	Sample Entropy	No. of Species	Lineages	Average Lineage Size
	1489554358	0	2	0	LeastEnergy	KE	BISTATE	0.00	0.00	1	1	3.00
	1489554358	0	2	1	LeastEnergy	KE	BISTATE	0.00	0.00	3	3	3.00
	1489554358	0	3	0	LeastEnergy	KE	BISTATE	0.00	0.00	3	3	3.00
	1489554358	0	3	1	LeastEnergy	KE	BISTATE	0.00	0.00	0	0	0.00
	1489554358	0	4	0	LeastEnergy	KE	BISTATE	0.00	0.00	1	1	3.00
	1489554358	0	4	1	LeastEnergy	KE	BISTATE	0.00	0.00	2	2	3.00
	1489554358	1	0	0	LeastEnergy	KE	AR		0.00	1	1	3.00
	1489554358	1	0	1	LeastEnergy	KE	AR		0.00	0	0	0.00
	1489554358	1	1	0	LeastEnergy	KE	AR	1.92	1.69	2	2	3.00
	1489554358	1	1	1	LeastEnergy	KE	AR	1.92	1.69	0	0	0.00
	1489554358	1	2	0	LeastEnergy	KE	AR	1.90	1.67	2	2	5.00
	1489554358	1	2	1	LeastEnergy	KE	AR	1.90	1.67	1	1	3.00
	1489554358	1	3	0	LeastEnergy	KE	AR	1.91	1.60	6	6	3.00
	1489554358	1	3	1	LeastEnergy	KE	AR	1.91	1.60	1	1	3.00
	1489554358	1	4	0	LeastEnergy	KE	AR	1.92	1.70	0	0	0.00
	1489554358	1	4	1	LeastEnergy	KE	AR	1.92	1.70	1	1	3.00
	1489554358	2	0	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	1	1	3.00
	1489554358	2	0	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	1	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	1	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	2	2	3.00
	1489554358	2	2	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	2	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	1	1	3.00
	1489554358	2	3	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	3	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	4	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	4	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	2	2	3.00
	1489554358	3	0	0	LeastEnergy	POPULATION	AR		0.00	0	0	0.00
	1489554358	3	0	1	LeastEnergy	POPULATION	AR		0.00	0	0	0.00
	1489554358	3	1	0	LeastEnergy	POPULATION	AR	1.92	1.66	0	0	0.00
	1489554358	3	1	1	LeastEnergy	POPULATION	AR	1.92	1.66	0	0	0.00
	1489554358	3	2	0	LeastEnergy	POPULATION	AR	1.90	1.85	0	0	0.00
	1489554358	3	2	1	LeastEnergy	POPULATION	AR	1.90	1.85	0	0	0.00
	1489554358	3	3	0	LeastEnergy	POPULATION	AR	1.90	1.79	0	0	0.00
	1489554358	3	3	1	LeastEnergy	POPULATION	AR	1.90	1.79	0	0	0.00
	1489554358	3	4	0	LeastEnergy	POPULATION	AR	1.94	1.44	0	0	0.00

Continued on next page

	Dataset	Experiment	Environment	Replicate	$S_{Product}$	Target	Shape	DFA	Sample Entropy	No. of Species	Lineages	Average Lineage Size
	1489554358	3	4	1	LeastEnergy	POPULATION	AR	1.94	1.44	0	0	0.00
	1489565574	0	0	0	Uniform	KE	BISTATE	0.00	0.00	1	1	3.00
	1489565574	0	0	1	Uniform	KE	BISTATE	0.00	0.00	2	2	3.00
	1489565574	0	1	0	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	1	1	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	2	0	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	2	1	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	3	0	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	3	1	Uniform	KE	BISTATE	0.00	0.00	3	3	3.00
	1489565574	0	4	0	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	4	1	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	1	0	0	Uniform	KE	AR		0.00	1	1	3.00
	1489565574	1	0	1	Uniform	KE	AR		0.00	1	1	3.00
	1489565574	1	1	0	Uniform	KE	AR	1.94	1.20	1	1	3.00
	1489565574	1	1	1	Uniform	KE	AR	1.94	1.20	0	0	0.00
	1489565574	1	2	0	Uniform	KE	AR	1.90	1.72	0	0	0.00
	1489565574	1	2	1	Uniform	KE	AR	1.90	1.72	2	2	3.00
	1489565574	1	3	0	Uniform	KE	AR	1.93	1.71	0	0	0.00
	1489565574	1	3	1	Uniform	KE	AR	1.93	1.71	0	0	0.00
	1489565574	1	4	0	Uniform	KE	AR	1.93	1.67	0	0	0.00
	1489565574	1	4	1	Uniform	KE	AR	1.93	1.67	13	13	3.00
	1489565574	2	0	0	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	0	1	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	1	0	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	1	1	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	2	0	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	2	1	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	3	0	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	3	1	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	4	0	Uniform	POPULATION	BISTATE	0.00	0.00	3	3	3.00
	1489565574	2	4	1	Uniform	POPULATION	BISTATE	0.00	0.00	2	2	3.00
	1489565574	3	0	0	Uniform	POPULATION	AR		0.00	1	1	3.00
	1489565574	3	0	1	Uniform	POPULATION	AR		0.00	1	1	3.00
	1489565574	3	1	0	Uniform	POPULATION	AR	1.94	1.28	0	0	0.00
	1489565574	3	1	1	Uniform	POPULATION	AR	1.94	1.28	0	0	0.00

Continued on next page

	Dataset	Experiment	Environment	Replicate	$S_{Product}$	Target	Shape	DFA	Sample Entropy	No. of Species	Lineages	Average Lineage Size
	1489565574	3	2	0	Uniform	POPULATION	AR	1.93	1.12	1	1	3.00
	1489565574	3	2	1	Uniform	POPULATION	AR	1.93	1.12	0	0	0.00
	1489565574	3	3	0	Uniform	POPULATION	AR	1.93	1.45	1	1	3.00
	1489565574	3	3	1	Uniform	POPULATION	AR	1.93	1.45	1	1	3.00
	1489565574	3	4	0	Uniform	POPULATION	AR	1.91	1.65	1	1	3.00
	1489565574	3	4	1	Uniform	POPULATION	AR	1.91	1.65	1	1	3.00

6.5 The emergence of variable replicators

We now move on to the question of the emergence of variable replicators. From the introduction to this chapter, our working definition of a variable replicator is "a multiplier that can change its structure while remaining the same under selection." In alg. 11 we present an algorithm for the first part of that definition—a multiplier that can change its structure—based on the earlier algorithm for the identification of multipliers.

The second part of the definition, that the candidate remains the same under selection, is equivalent to the statement that changes to the entity are neutral with respect to selection, or in other words, selection and the changes of entity state are independent. From this we can construct a hypothesis test:

H_0: The sequence of state changes of any candidate variable replicator is only in response to external selective pressure.

H_1: The sequence of state changes of all candidate variable replicators is independent of external selective pressure.

6.5.1 Results and discussion

As might be expected, because the reaction cycles that make up a variable replicator do not have to be identical (unlike those in an exact multiplier), the number of lineages and the average length of lineages for candidate variable replicators in table 6.3 significantly exceed those seen for exact multipliers.

Now, we can develop a notation for replicator species by denoting each species of reaction cycle by a single digit or character; a replicator composed of a particular sequence

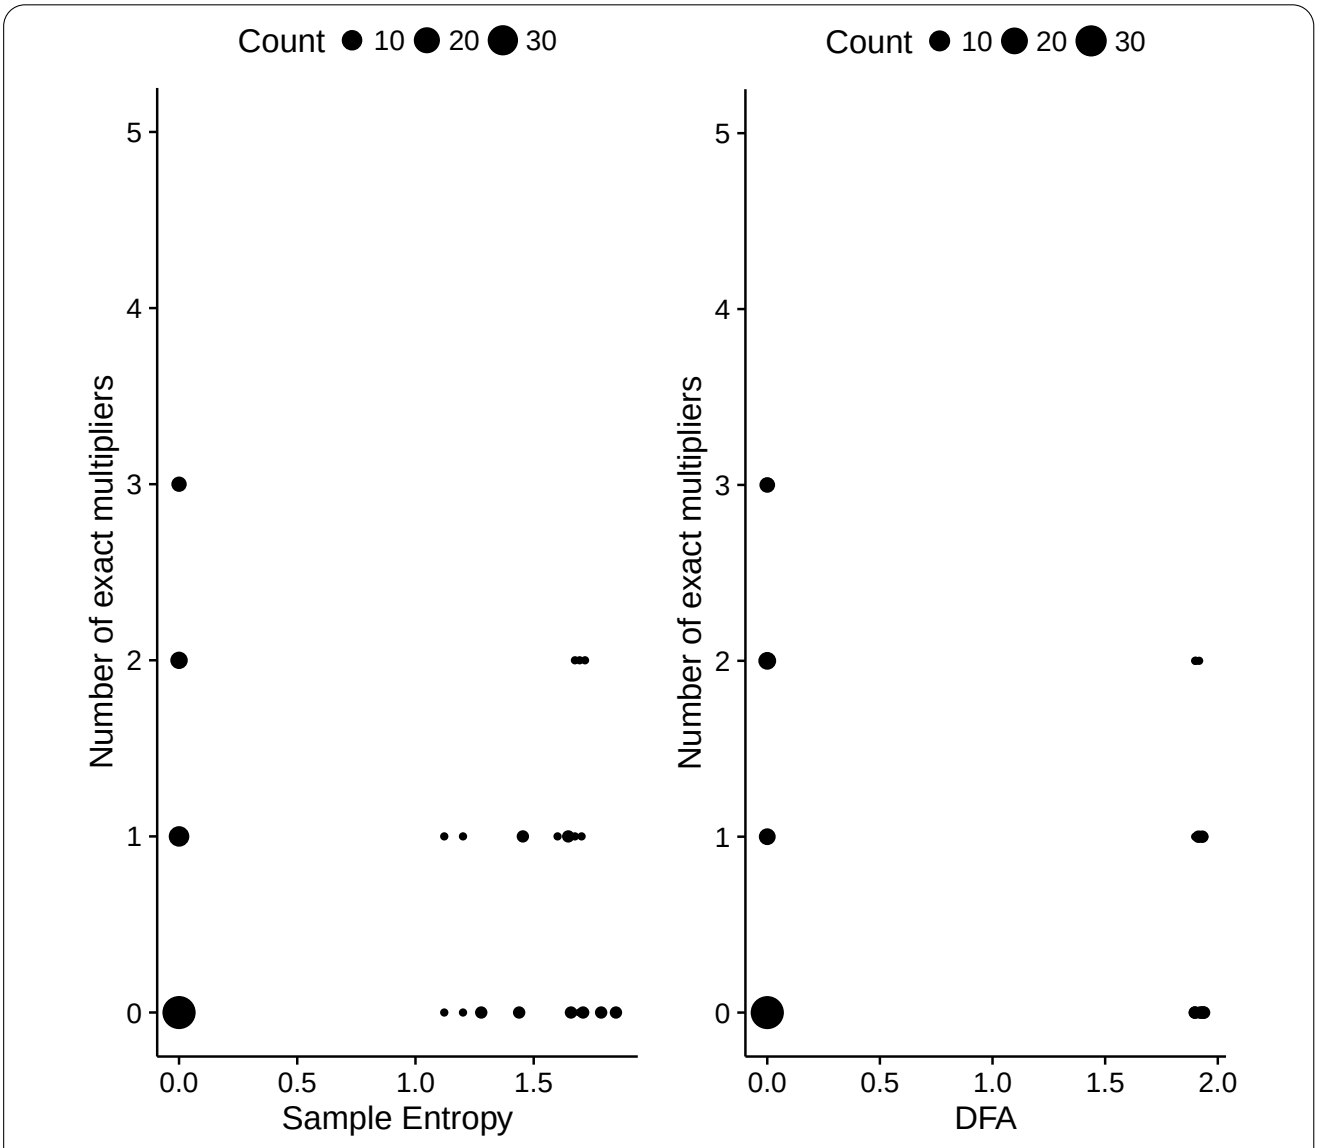

Figure 6.1: Number of exact multipliers against measures of extrinsic selective pressure Sample Entropy (left) and DFA (right.) The two outliers mentioned in the text have been excluded from each graph.

```
1  Function DiscoverVariableMultipliers(MultipliersBySpecies):
2      Candidates ← ∅
3      AllProducts ← ⋃_{∀x∈MultipliersBySpecies} Products(x)\Reactants(x)
4      AllReactants ← ⋃_{∀x∈MultipliersBySpecies} Reactants(x)\Products(x)
5      LinkingMolecules ←AllProducts ∩ AllReactants
6      if LinkingMolecules then
7          for Cycle in Cluster do
8              if (Products(cycle) ∩ LinkingMolecules ≠ ∅) ∨ (Reactants(cycle) ∩
                LinkingMolecules ≠ ∅) then
9                  MultipliersInSpecies ← MultipliersInSpecies ∪ Cycle
10             end
11         end
12     return Candidates
```

Algorithm 11: *DiscoverVariableMultipliers*. Detect candidate variable replicators within a reaction network, where a candidate is a continuous sequence of multipliers linked by shared molecules. The algorithm is similar to that for detecting multipliers (alg. 7) without the need to check stoichiometry.

of cycles can then be categorized as belonging to the replicator species represented by the set of single digit or character cycle species. The sequence of cycle states for the replicators identified in this experiment is shown in table 6.4.

Our null hypothesis is that changes in variable replicator state only arise as the result of external selection. That is, that replicator states behave differently under selection. This can be tested by comparing the sequence of states of the same replicator across a variety of environments, each providing a different selective pressure.

Note though that only four cycle species ('M', 'G', 'P' and 'i') appear in more than one run, with the great majority of species endemic to a single run. As a result, only three candidate replicators ('MMM', 'GGG' and 'iii') appear in more than one run in our results, and as these are all single-state replicators none can be considered variable by definition. The cycle 'P' appears in two candidate replicators, but as in one instance it appears with 'Q' and in the other by itself, these also cannot be considered as instances of the same replicator as the two candidates have a different number of states. Unfortunately therefore we cannot test the hypothesis by examining the behaviour of the same replicator across different environments.

Furthermore, if the states in a variable replicator are more or less equally likely, as seems a reasonable approximation, then we would expect a replicator to revisit states over time. However, none of our identified replicators do this; they either occupy only one state (in which case they cannot be claimed as a variable replicator), or move through a

non-repeating sequence of unique states.

In summary, we lack unequivocal evidence of variable replication, and so cannot reject H_0: we conclude that there is no evidence for variable replication in the ToyWorld artificial chemistry under the experimental conditions described.

Table 6.3: Variable replicators measured by alg. 11 in reaction networks generated by LeastEnergy and Uniform S_{Product} strategies under forms of environmental change given by the columns *Target* and *Shape*. Sampling proportion p in alg. 8 set to 0.2. Sample Entropy scaled by a factor of 1000 (see section 6.3).

	Dataset	Experiment	Environment	Replicate	$S_{Product}$	Target	Shape	DFA	Sample Entropy	No. of Species	Lineages	Average Lineage Size
	1489554358	0	0	0	LeastEnergy	KE	BISTATE	0.00	0.00	1	6	11.00
	1489554358	0	0	1	LeastEnergy	KE	BISTATE	0.00	0.00	0	0	0.00
	1489554358	0	1	0	LeastEnergy	KE	BISTATE	0.00	0.00	0	0	0.00
	1489554358	0	1	1	LeastEnergy	KE	BISTATE	0.00	0.00	0	0	0.00
	1489554358	0	2	0	LeastEnergy	KE	BISTATE	0.00	0.00	1	3	13.00
	1489554358	0	2	1	LeastEnergy	KE	BISTATE	0.00	0.00	1	9	19.00
	1489554358	0	3	0	LeastEnergy	KE	BISTATE	0.00	0.00	3	9	11.00
	1489554358	0	3	1	LeastEnergy	KE	BISTATE	0.00	0.00	0	0	0.00
	1489554358	0	4	0	LeastEnergy	KE	BISTATE	0.00	0.00	1	3	10.00
	1489554358	0	4	1	LeastEnergy	KE	BISTATE	0.00	0.00	1	6	11.00
	1489554358	1	0	0	LeastEnergy	KE	AR		0.00	1	3	10.00
	1489554358	1	0	1	LeastEnergy	KE	AR		0.00	0	0	0.00
	1489554358	1	1	0	LeastEnergy	KE	AR	1.92	1.69	1	6	11.00
	1489554358	1	1	1	LeastEnergy	KE	AR	1.92	1.69	0	0	0.00
	1489554358	1	2	0	LeastEnergy	KE	AR	1.90	1.67	1	10	11.00
	1489554358	1	2	1	LeastEnergy	KE	AR	1.90	1.67	1	3	10.00
	1489554358	1	3	0	LeastEnergy	KE	AR	1.91	1.60	3	18	11.00
	1489554358	1	3	1	LeastEnergy	KE	AR	1.91	1.60	1	3	13.00
	1489554358	1	4	0	LeastEnergy	KE	AR	1.92	1.70	0	0	0.00
	1489554358	1	4	1	LeastEnergy	KE	AR	1.92	1.70	1	3	10.00
	1489554358	2	0	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	1	3	10.00
	1489554358	2	0	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	1	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	1	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	1	6	11.00
	1489554358	2	2	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	2	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	1	3	10.00
	1489554358	2	3	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	3	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00

Continued on next page

	Dataset	Experiment	Environment	Replicate	$S_{Product}$	Target	Shape	DFA	Sample Entropy	No. of Species	Lineages	Average Lineage Size
	1489554358	2	4	0	LeastEnergy	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489554358	2	4	1	LeastEnergy	POPULATION	BISTATE	0.00	0.00	1	6	11.00
	1489554358	3	0	0	LeastEnergy	POPULATION	AR		0.00	0	0	0.00
	1489554358	3	0	1	LeastEnergy	POPULATION	AR		0.00	0	0	0.00
	1489554358	3	1	0	LeastEnergy	POPULATION	AR	1.92	1.66	0	0	0.00
	1489554358	3	1	1	LeastEnergy	POPULATION	AR	1.92	1.66	0	0	0.00
	1489554358	3	2	0	LeastEnergy	POPULATION	AR	1.90	1.85	0	0	0.00
	1489554358	3	2	1	LeastEnergy	POPULATION	AR	1.90	1.85	0	0	0.00
	1489554358	3	3	0	LeastEnergy	POPULATION	AR	1.90	1.79	0	0	0.00
	1489554358	3	3	1	LeastEnergy	POPULATION	AR	1.90	1.79	0	0	0.00
	1489554358	3	4	0	LeastEnergy	POPULATION	AR	1.94	1.44	0	0	0.00
	1489554358	3	4	1	LeastEnergy	POPULATION	AR	1.94	1.44	0	0	0.00
	1489565574	0	0	0	Uniform	KE	BISTATE	0.00	0.00	1	3	10.00
	1489565574	0	0	1	Uniform	KE	BISTATE	0.00	0.00	1	6	11.00
	1489565574	0	1	0	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	1	1	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	2	0	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	2	1	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	3	0	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	3	1	Uniform	KE	BISTATE	0.00	0.00	3	9	10.00
	1489565574	0	4	0	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	0	4	1	Uniform	KE	BISTATE	0.00	0.00	0	0	0.00
	1489565574	1	0	0	Uniform	KE	AR		0.00	1	3	10.00
	1489565574	1	0	1	Uniform	KE	AR		0.00	1	3	10.00
	1489565574	1	1	0	Uniform	KE	AR	1.94	1.20	1	3	10.00
	1489565574	1	1	1	Uniform	KE	AR	1.94	1.20	0	0	0.00
	1489565574	1	2	0	Uniform	KE	AR	1.90	1.72	0	0	0.00
	1489565574	1	2	1	Uniform	KE	AR	1.90	1.72	1	6	11.00
	1489565574	1	3	0	Uniform	KE	AR	1.93	1.71	0	0	0.00
	1489565574	1	3	1	Uniform	KE	AR	1.93	1.71	0	0	0.00
	1489565574	1	4	0	Uniform	KE	AR	1.93	1.67	0	0	0.00
	1489565574	1	4	1	Uniform	KE	AR	1.93	1.67	1	39	23.00
	1489565574	2	0	0	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	0	1	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	1	0	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00

Continued on next page

	Dataset	Experiment	Environment	Replicate	$S_{Product}$	Target	Shape	DFA	Sample Entropy	No. of Species	Lineages	Average Lineage Size
	1489565574	2	1	1	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	2	0	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	2	1	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	3	0	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	3	1	Uniform	POPULATION	BISTATE	0.00	0.00	0	0	0.00
	1489565574	2	4	0	Uniform	POPULATION	BISTATE	0.00	0.00	1	9	13.00
	1489565574	2	4	1	Uniform	POPULATION	BISTATE	0.00	0.00	2	6	10.00
	1489565574	3	0	0	Uniform	POPULATION	AR		0.00	1	3	10.00
	1489565574	3	0	1	Uniform	POPULATION	AR		0.00	1	3	13.00
	1489565574	3	1	0	Uniform	POPULATION	AR	1.94	1.28	0	0	0.00
	1489565574	3	1	1	Uniform	POPULATION	AR	1.94	1.28	0	0	0.00
	1489565574	3	2	0	Uniform	POPULATION	AR	1.93	1.12	1	3	10.00
	1489565574	3	2	1	Uniform	POPULATION	AR	1.93	1.12	0	0	0.00
	1489565574	3	3	0	Uniform	POPULATION	AR	1.93	1.45	1	3	10.00
	1489565574	3	3	1	Uniform	POPULATION	AR	1.93	1.45	1	3	10.00
	1489565574	3	4	0	Uniform	POPULATION	AR	1.91	1.65	1	3	10.00
	1489565574	3	4	1	Uniform	POPULATION	AR	1.91	1.65	1	3	10.00

6.6 Conclusions

This concludes our investigation into replication in an artificial chemistry. The previous chapter introduced exact multipliers in stable environments; in this chapter we introduced a changing environment and examined the effect of varying external selective pressure upon multiplication. The effects of selection were also fundamental to the final portion of the chapter: a search for variable replicators, or replicators that can take several states, each state equivalent under selection.

The chapter began by constructing a model for environmental change, made up of two dimensions: shape of change, and the target of change. The target of change defines the area of the evolutionary model to be affected by change, while the shape of change (essentially a time-series) captures the magnitude of the change to be applied at each time-step to the target. For the experiments described in the chapter 6, each of shape and target is associated with two alternative forms or types.

Shape in our experiments follows either a straightforward alternation between two stable values ("bistate"), or a more complex form described by an AR-timeseries. In order that the forms may be simply described, we summarize each time-series by two metrics: DFA, a type of Hurst parameter, and Sample Entropy, a measure of complexity. The two targets of change addressed in this chapter are kinetic energy (the overall kinetic energy of the molecules within the reactor population) and the reactor population itself.

Our experiment results were inconclusive. Although the artificial chemistry produced exact multipliers in approximately one-third of all runs, a relationship between environmental variability and the numbers of multipliers was unproven. This is primarily a consequence of the low number of multipliers observed. So what then might be the cause of the low numbers? First, it is important to remember that this is not unexpected, as the lack of successful previous work attests. Even in biology, where we benefit from having an existence proof, we lack a complete and compelling description for the equivalent series of steps from molecular chemistry to replication, although there are several competing narratives. It is clear though that in artificial domains, in the absence of a shortcut mechanism (such as described in section 1.2.4), the evolution of multipliers, based on autocatalytic sets, is quite unlikely. As observed by Faulconbridge (2011, p173), when discussing the choice of metric for evaluating fitness in RBN-World, "As with hypercycles, however, autocatalytic sets are too rare to be useful...". Although the context is as a fitness function for a genetic algorithm, the point remains valid that autocatalytic cycles, and hence multipliers, are unusual. This is likely exacerbated by our adoption of a sampling algorithm for detecting multipliers; as noted in section 5.3 this undercounts multipliers, but is necessary for analysis times that are measured in days rather than weeks per experimental run.

Similarly, although the artificial chemistry generated a number of variable replicator candidates—limited sets of repeated reaction cycles—none of the candidates were other than endemic to a single run. As the condition for a full variable replicator is that each state should be equivalent under selection, the lack of alternative forms in other runs prevents us from properly testing this condition. However, we would have expected a variable replicator to oscillate between a set of states while our candidates never repeated the same state. On balance we feel that the results cannot be other than inconclusive.

Biology gives us many examples of functional replicators, but provides little guidance on alternative approaches; we don't wish to bias our simulation model in a direction which may very well be wrong. Unfortunately, although we have no prior knowledge to estimate probabilities, it is clear that even without the constraint of attempting to replicate the process by which replication evolved on Earth, replication is very unlikely.

The critical step appears to be the shift from a limited heredity replicator to an unlimited one. At that point, the necessary leap in information that must be maintained and managed by the entity's genome seems to require a quite different mechanism for infor-

mation storage—a shift from an analog genome to a digital (for example, templated) one, associated with a process to map between genome and phenome. The outcome of this change is a capacity for increased expressiveness, and an increased ability to respond to complexity in the environment.

Table 6.4: Cycle sequences for each multiplier and variable replicator identified in the reaction graph. All cycles of same species (see text) replaced by the same single character or digit (see table A.3.)

Dataset	Experiment	Environment	Replicate	Cycle Sequence
1489554358	0	0	0	A A A B B B
1489554358	0	2	0	C C C
1489554358	0	2	1	D D D E E E F F F
1489554358	0	3	0	G G G
1489554358	0	3	0	H H H
1489554358	0	3	0	I I I
1489554358	0	4	0	J J J
1489554358	0	4	1	K K K L L L
1489554358	1	0	0	M M M
1489554358	1	1	0	N N N O O O
1489554358	1	2	0	P P P P P Q Q Q Q Q
1489554358	1	2	1	P P P
1489554358	1	3	0	R R R S S S
1489554358	1	3	0	T T T U U U
1489554358	1	3	0	V V V W W W
1489554358	1	3	1	X X X
1489554358	1	4	1	Y Y Y
1489554358	2	0	0	M M M
1489554358	2	1	1	Z Z Z a a a
1489554358	2	2	1	M M M
1489554358	2	4	1	G G G b b b
1489565574	0	0	0	c c c
1489565574	0	0	1	d d d e e e
1489565574	0	3	1	f f f
1489565574	0	3	1	g g g
1489565574	0	3	1	h h h
1489565574	1	0	0	i i i
1489565574	1	0	1	j j j
1489565574	1	1	0	k k k
1489565574	1	2	1	l l l m m m
1489565574	1	4	1	n n n o o o p p p q q q r r r s s s t t t u u u v v v w w w x x x y y y z z z
1489565574	2	4	0	0 0 0 1 1 1 2 2 2
1489565574	2	4	1	3 3 3
1489565574	2	4	1	4 4 4
1489565574	3	0	0	5 5 5
1489565574	3	0	1	6 6 6
1489565574	3	2	0	i i i
1489565574	3	3	0	G G G
1489565574	3	3	1	Y Y Y
1489565574	3	4	0	7 7 7
1489565574	3	4	1	8 8 8

7 Conclusions

This thesis has explored two related topics in the emergence of evolutionary replicators from artificial chemistries: first, how environmental variability affects heritability in a simple evolutionary model, and second, how the reactions in an artificial chemistry can result in simple replicating entities. The combination of these two topics suggests a pathway towards the formation of full informational replicators in artificial chemistries purely through evolutionary bootstrapping.

We shall now return to our research questions from section 1.3, and for each question summarize the major results and findings. Then, in order to qualify the scope of the results, we critique some aspects of the work in section 7.3, which naturally leads in to our final section, a consideration of future work in section 7.4.

7.1 RQ1: In what way does selective pressure drive changes in heredity in a population of evolving replicators?

This research question served as the primary motivation behind the work documented in chapter 2, and in particular for the model relating heritability and fitness described in that chapter. By making the usually implicit measures of heritability and fitness explicit for each entity, the model allows us to measure them directly as the population of an otherwise standard evolutionary system evolves. Our first experiment tested the evolution of heritability in a stable environment (without external selective pressure), before we introduced an explicit model of environmental change for the second and final experiment in the chapter.

The changing environmental model incorporates a novel application of an AR-timeseries, combining a regular component and an unpredictable component. At each generation of the evolutionary model, the corresponding value from the timeseries is added to the fitness value of each entity in the population, so efficiently simulating the effects of external environmental change.

7.1.1 Main findings

In summary, the main findings related to RQ1 are:

1. Heritability is proportional to the predictability of the environment, and is at a minimum in conditions of maximum unpredictability, and at a maximum in stable conditions.
2. The variation in heritability, $\sigma_{heritability}$, is proportional to the degree of environmental variability. As a corollary, the $\sigma_{heritability}$ under changing conditions is greater than that under stable conditions.

These results align with an intuitive understanding of the effects of environmental change: in varying environments, a reservoir of variation provides flexibility, whereas perfect inheritance restricts the variation in a population, and limits the population's ability to respond to change.

7.2 RQ2: Can variable replicators emerge from a molecular artificial chemistry?

Our approach to this research question is based upon the exploration of a hierarchy of replication proposed by Zachar and Szathmáry (2010). We began in chapter 4 by describing a modularized artificial chemistry for experimentation, ToyWorld, and conjectured that a count of reaction cycles might serve as a proxy for the number of higher level replicators. We showed that the strategies used to select reactants and products in the artificial chemistry affected the higher level entities (such as reaction cycles) observed in the resulting reaction networks.

The next two chapters used the ToyWorld artificial chemistry to search for first multipliers and then variable replicators in the reaction graphs that result from a selection of reactant and product selection strategies. We also extended the environmental model from chapter 2 from one to two dimensions, to include not only the shape or form of the timeseries, but the element of the evolutionary model to receive the change (the target).

To the single shape of an AR-timeseries from chapter 2, we also added a second shape in the form of a bistate series that switches between two alternate values at regular intervals.

For the target of change, we defined two alternatives, one being kinetic energy (where each molecule's kinetic energy is increased or decreased by the change value) and the other directly modifying the population by adding or removing the number of molecules specified by the change value.

Finally, to detect replicators within the reaction networks, we defined three novel algorithms:

1. Reaction cycle detection from a set of seed molecules selected by sampling from the molecules in the reaction graph (alg. 8.)

2. Identification of exact multipliers from reaction cycles, based on a specific definition for exact multipliers as two or more copies of the same reaction cycle species, where the reaction cycle species has stoichiometry greater than one, and where each cycle in the multiplier is connected to at least one other multiplier cycle by a molecule that is a product of one cycle and a reactant in the other (alg. 7.)

3. Identification of variable replicator candidates (that is, without consideration of selection) from reaction cycles, where we define variable replicators as multipliers that can occupy any of a limited set of states without losing their underlying identity (alg. 11.)

7.2.1 Main findings

The main findings for RQ2 are summarized below.

1. Exact multipliers do arise in the ToyWorld artificial chemistry (in approximately one-third of all runs), but not in any large number, and when they do, they do not persist for long.

2. Multipliers occur as the result of a non-neutral combination of Product and Reactant selection strategies, and not purely by chance.

3. The hypothesised relationship between environmental variability and the numbers of multipliers remains unproven.

4. Although ToyWorld produces a number of variable replicator candidates (sequences of repeated reaction cycles), none were observed that meet all of our criteria for variable replication. Specifically, although we observe exact multipliers, and variable replicator candidates, we do not observe candidates that only occupy a restricted or limited set of states, or that exist in multiple runs (see discussion in section 7.3)

5. The choice of S_{Reactant} has a significant effect on the emergence of reaction cycles in ToyWorld; S_{Product} is of lesser effect.

6. A Kinetic Reactant selection strategy is more effective for cycle emergence than a Uniform one. The number of cycles, and length of longest cycle, are maximized under the combination of a Kinetic Reactant selection strategy and a LeastEnergy Product selection strategy.

7. Other parameters to the ToyWorld model have lesser effects on cycle formation; only the relationship of E_{bonds} to the number of cycles produced has any significance.

7.3 Limitations

Identifying cycles in a reaction network by searching from a sampled set of seed molecules dramatically improves the performance of the cycle detection algorithm, and as shown in chapters 5 and 6, produces sufficient density of cycles to establish the presence of multipliers and variable replicators. However, a significant limitation is that sampling influences the likelihood of detecting all cycles in a multiplier or variable replicator and so the length of a replicator is likely to be under-reported. If a cycle in the middle of sequence of cycles is not detected through sampling, the algorithms will identify two shorter replicators instead of one longer one.

The next topics are closely linked. First, it seems clear from both the experiments in this work, and from our knowledge of early life, that the probability of complex replicators arising within 100,000 generations (around the current limit of practicality for analysis) under the conditions described is extremely low. Some combination of a significant increase in generations and a change of conditions will be required to increase the likelihood of observing a significant step such as the emergence of an informational replicator. At present, the only approach is to conduct many extremely long-duration trials and so leverage probabilities. This is clearly unsatisfactory.

Second, the current work does not provide any guidance as to how that step from a variable to an template-based informational replicator might arise, and yet it is essential if complex replicators are to form.

Third, our current algorithms rely upon network analysis at the reaction cycle level, and cannot inherently detect any higher-level structure. The distinction between genotype and phenotype in an informational replicator occurs at a different conceptual level to the component reactions, and our level of interest and investigation needs to change accordingly. We need an approach that adapts to different levels of emergence, from cycles to elements built from cycles, to yet more complex elements, and so on.

Finally, the current approach of graph analysis of complete reaction networks, even dynamic analysis, cannot scale to the network sizes needed. It seems clear that either the performance of the current algorithms must be dramatically improved, perhaps by rewriting in a lower-level higher-performance computer language, or more profitably, the approach to analysing the generated data must change if larger networks are mandated. The graph structure we generate at present forms a single connected component in which every molecule and reaction are contained (with the exception of those molecules that never take part in a reaction). The lack of obvious substructures within this single graph makes it difficult to naturally subset the graph to improve the speed of the analysis.

7.4 Future work

The previous section has identified issues primarily with the scope of the present work. In this section, we concentrate on those improvements that could be made within the current scope.

Although the AR-timeseries generator described in section 2.7.1 produces a time series for environmental change with the property of stationarity, the δ term makes the evolutionary model of fitness non-stationary. However, any change still remains steady and gradual. An extension would be to co-opt the idea of concept drift from time series analysis to induce an abrupt change with probability p at each generation. Each change would therefore form a new 'concept'. Instead of the environment changing in a predictable and describable way from one generation to another, the change could not be predictable from the earlier history.

There are some obvious extensions of the model from section 2.5 that have been left for future work. First, the model currently assumes only single-parent inheritance, whereas many biological species have two parents. Extending to two parents would be a useful enhancement to increase the model's scope. Second, the model does not include any influence from development (the production of the phenome from the genome). However, it is unclear at this stage what effect development would have on the model as its effects are bundled into the overall *fitness* parameter. Finally, although outside of the overall scope of this work in evolutionary systems, the effect of acquired characteristics would be interesting to explore. Others (*e.g.* Gaucherel and Jensen (2012), Paenke et al. (2007), and Sasaki and Tokoro (2000)) have studied the differences between general models based on acquired and non-acquired characteristics, finding a difference between models in changing environments. This could be another area of exploration for the future.

A limitation of the experiment design in chapter 4 is that the values chosen for the high and low values of E_{Bonds} make it impossible to determine the cause of the difference observed in section 4.10.3. There are two alternative explanations: first, the energy required to make or break bonds is simply different between the two factor levels; second, in the low factor level, based on real-world values, the bond make and break energies for even a single bond vary depending on the atoms involved, while in the high factor level these values are consistent for all bonds of the same degree. To distinguish between the two explanations, the average levels at each degree should be the same for each factor; this is a suggestion for a future experiment.

As mentioned in section 7.3, the sampling algorithm for cycles means that the sizes of replicators may be underestimated. The sampling proportion p could certainly be increased, but this is currently impractical for large reaction graphs. Alternatively, repeating the cycle detection with a new set of seed molecules could eliminate any sampling gaps.

After identifying an initial set of replicators, the seed molecules for the repeated cycle detection should include the product and reactant molecules from the replicator cycles that are not consumed or produced, respectively, by the replicator.

In chapter 6 we have assumed that each environmental change affected all entities equally; however, this isn't necessarily the only option. We can identify three levels of scope, or the proportion of entities in the population to receive a particular set of changes, from most homogeneous to least:

1. The group of all entities. All entities receive the same set of changes.

2. A group for each set of "related" entities, where the most natural and obvious relation is that between parent and child; this is unambiguous and straightforward in our model where each entity has only one parent. We refer to a group of entities related by inheritance as a *lineage*. A separate set of changes is provided for each lineage.

3. A single-member group for each entity. Each entity receives a unique set of changes.

The first level is the simplest application of environmental change, and the one adopted throughout this work, while the second represents the common scenario where we expect similar entities to react in similar ways to change, and where similarity is a result of descent: entities that share a common ancestor are more similar to each other than they are to other lineages.

The third scope level implies that each entity has an independent response to environmental changes. This seems problematic; environmental response is a function of phenotypes, and we would expect related entities to have related phenotypes[1]. Thus instead of single-member groups we would expect lineage-related groups.

7.5 Personal reflection

The work in this thesis has been driven by a belief that the most effective approach to replication was likely to be one in which replicators could emerge from a simpler artificial chemical system without explicit external design or direction. The experiments described in this work have therefore followed a "big data" approach, but given limited resources, they have been only to an exploratory scale. Given the probabilities involved, it is clear that ideally we would greatly increase the number of trials, or reactions simulated.

This could be done by rearchitecting the current implementation to use cloud-based computing resources in parallel. Running multiple experiments in parallel, with cycle, multiplier and variable replicator detection happening in real-time from streamed results

[1] In general, although in biology there can be significant phenotypic differences between related entities.

would effectively remove the current limits on the number of reactions that can be modelled and analysed.

A Experiment datasets and procedures

Table A.1: Summary of parameters for each set of replicates in the experiments. The columns dataset, experiment and environment define each experiment, with parameters given by columns reactant strategy, product strategy, and environmental change target and shape. The two final columns, DFA and sample entropy, provide two alternative measures for the variability of the environmental model for the experiment: DFA is the calculated Hurst parameter (H) using detrended fluctuation analysis, while sample entropy is assessed by Richman and Moorman (2000). Sample Entropy scaled by a factor of 1000 (see section 6.3).

Dataset	Experiment	Environment	Reactant strategy	Product strategy	Target	Shape	DFA	Sample Entropy
1489554358	0	0	KineticReactantSelection	LeastEnergy	KE	BISTATE	0.00	0.00
1489554358	0	1	KineticReactantSelection	LeastEnergy	KE	BISTATE	0.00	0.00
1489554358	0	2	KineticReactantSelection	LeastEnergy	KE	BISTATE	0.00	0.00
1489554358	0	3	KineticReactantSelection	LeastEnergy	KE	BISTATE	0.00	0.00
1489554358	0	4	KineticReactantSelection	LeastEnergy	KE	BISTATE	0.00	0.00
1489554358	1	0	KineticReactantSelection	LeastEnergy	KE	AR		0.00
1489554358	1	1	KineticReactantSelection	LeastEnergy	KE	AR	1.92	1.69
1489554358	1	2	KineticReactantSelection	LeastEnergy	KE	AR	1.90	1.67
1489554358	1	3	KineticReactantSelection	LeastEnergy	KE	AR	1.91	1.60
1489554358	1	4	KineticReactantSelection	LeastEnergy	KE	AR	1.92	1.70
1489554358	2	0	KineticReactantSelection	LeastEnergy	POPULATION	BISTATE	0.00	0.00
1489554358	2	1	KineticReactantSelection	LeastEnergy	POPULATION	BISTATE	0.00	0.00
1489554358	2	2	KineticReactantSelection	LeastEnergy	POPULATION	BISTATE	0.00	0.00
1489554358	2	3	KineticReactantSelection	LeastEnergy	POPULATION	BISTATE	0.00	0.00
1489554358	2	4	KineticReactantSelection	LeastEnergy	POPULATION	BISTATE	0.00	0.00
1489554358	3	0	KineticReactantSelection	LeastEnergy	POPULATION	AR		0.00
1489554358	3	1	KineticReactantSelection	LeastEnergy	POPULATION	AR	1.92	1.66
1489554358	3	2	KineticReactantSelection	LeastEnergy	POPULATION	AR	1.90	1.85
1489554358	3	3	KineticReactantSelection	LeastEnergy	POPULATION	AR	1.90	1.79

Continued on next page

Dataset	Experiment	Environment	Reactant strategy	Product strategy	Target	Shape	DFA	Sample Entropy
1489554358	3	4	KineticReactantSelection	LeastEnergy	POPULATION	AR	1.94	1.44
1489565574	0	0	KineticReactantSelection	Uniform	KE	BISTATE	0.00	0.00
1489565574	0	1	KineticReactantSelection	Uniform	KE	BISTATE	0.00	0.00
1489565574	0	2	KineticReactantSelection	Uniform	KE	BISTATE	0.00	0.00
1489565574	0	3	KineticReactantSelection	Uniform	KE	BISTATE	0.00	0.00
1489565574	0	4	KineticReactantSelection	Uniform	KE	BISTATE	0.00	0.00
1489565574	1	0	KineticReactantSelection	Uniform	KE	AR		0.00
1489565574	1	1	KineticReactantSelection	Uniform	KE	AR	1.94	1.20
1489565574	1	2	KineticReactantSelection	Uniform	KE	AR	1.90	1.72
1489565574	1	3	KineticReactantSelection	Uniform	KE	AR	1.93	1.71
1489565574	1	4	KineticReactantSelection	Uniform	KE	AR	1.93	1.67
1489565574	2	0	KineticReactantSelection	Uniform	POPULATION	BISTATE	0.00	0.00
1489565574	2	1	KineticReactantSelection	Uniform	POPULATION	BISTATE	0.00	0.00
1489565574	2	2	KineticReactantSelection	Uniform	POPULATION	BISTATE	0.00	0.00
1489565574	2	3	KineticReactantSelection	Uniform	POPULATION	BISTATE	0.00	0.00
1489565574	2	4	KineticReactantSelection	Uniform	POPULATION	BISTATE	0.00	0.00
1489565574	3	0	KineticReactantSelection	Uniform	POPULATION	AR		0.00
1489565574	3	1	KineticReactantSelection	Uniform	POPULATION	AR	1.94	1.28
1489565574	3	2	KineticReactantSelection	Uniform	POPULATION	AR	1.93	1.12
1489565574	3	3	KineticReactantSelection	Uniform	POPULATION	AR	1.93	1.45
1489565574	3	4	KineticReactantSelection	Uniform	POPULATION	AR	1.91	1.65
1489951262	0	0	UniformReactantSelection	Uniform	POPULATION	BISTATE	0.00	0.00

Table A.2: Pipelines used in this work to transform input data of one form, such as raw reaction data or cycle data into another form for analysis.

Input	Algorithm	Output
Reaction network	alg. 8	Cycles by molecule
Cycles by molecule	alg. 7	All multipliers
Multipliers	alg. 11	All variable replicators

Table A.3: All reaction cycles in all observed multipliers from chapters 5 and 6 labelled by an arbitrarily-assigned single character or digit species identifier. Each cycle is represented as a series of reactions, numbered in sequence.

Species	Steps in Cycle
0	1. [H+]+[H][O]->[H+]+[H+]+[O-] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[H+]+[H+]+[O-]
1	1. [H+]+[O-]->[H][O] 2. [H+]+[H][O]->[H+]+[H+]+[O-] 3. [H+]+[O-]->[H][O] 4. [H][O]+[H+]->[H+]+[H+]+[O-]
2	1. [H+]+[H][O]->[O-]+[H+]+[H+] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[H+]+[O-] 4. [H+]+[O-]->[H][O] 5. [H][O]+[H+]->[H+]+[H+]+[O-]
3	1. [H][O]+[H+]->[H+]+[O-]+[H+] 2. [O-]+[H+]->[H][O] 3. [H][O]+[H+]->[O-]+[H+]+[H+]
4	1. [H][O]+[H+]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H+]+[H][O]->[H+]+[H+]+[O-]
5	1. [H+]+[H][O]->[H+]+[H+]+[O-] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[H+]+[O-]+[H+]
6	1. [H+]+[O-]->[H][O] 2. [H][O]+[H+]->[H+]+[H+]+[O-] 3. [H+]+[O-]->[H][O] 4. [H+]+[H][O]->[H+]+[H+]+[O-]
7	1. [H+]+[H][O]->[H+]+[O-]+[H+] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[H+]+[O-]
8	1. [H][O]+[H+]->[O-]+[H+]+[H+] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[H+]+[H+]+[O-]
A	1. [O-]+[H+]->[H][O] 2. [H][O]+[H+]->[H+]+[O-]+[H+] 3. [H+]+[O-]->[H][O] 4. [H+]+[H][O]->[H+]+[H+]+[O-]
B	1. [H][O]+[H+]->[H+]+[O-]+[H+] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[H+]+[O-]
C	1. [H+]+[O-]->[H][O] 2. [H+]+[H][O]->[H+]+[H+]+[O-] 3. [H+]+[O-]->[H][O] 4. [H+]+[H][O]->[H+]+[H+]+[O-]
D	1. [H+]+[O-]->[H][O] 2. [H+]+[H][O]->[H+]+[O-]+[H+] 3. [O-]+[H+]->[H][O] 4. [H][O]+[H+]->[O-]+[H+]+[H+] 5. [O-]+[H+]->[H][O] 6. [H+]+[H][O]->[H+]+[H+]+[O-]
E	1. [H+]+[O][O-]->[H+]+[O]+[O-] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[O-]+[H+] 4. [O-]+[H+]->[H][O] 5. [H][O]+[H+]->[O-]+[H+]+[H+] 6. [O-]+[H+]->[H][O] 7. [H+]+[H][O]->[H+]+[H+]+[O-]
F	1. [H+]+[H][O]->[H+]+[O-]+[H+] 2. [O-]+[H+]->[H][O] 3. [H][O]+[H+]->[O-]+[H+]+[H+] 4. [O-]+[H+]->[H][O] 5. [H+]+[H][O]->[H+]+[H+]+[O-]
G	1. [H+]+[H][O]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H][O]+[H+]->[H+]+[O-]+[H+]
H	1. [H][O]+[H+]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H][O]+[H+]->[H+]+[H+]+[O-]
I	1. [O-]+[H+]->[H][O] 2. [H+]+[H][O]->[H+]+[H+]+[O-] 3. [H+]+[O-]->[H][O] 4. [H][O]+[H+]->[O-]+[H+]+[H+]
J	1. [H][O]+[H+]->[H+]+[O-]+[H+] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[O-]+[H+]+[H+]
K	1. [H+]+[O-]->[H][O] 2. [H][O]+[H+]->[H+]+[H+]+[O-] 3. [O-]+[H+]->[H][O] 4. [H+]+[H][O]->[O-]+[H+]+[H+]
L	1. [H][O]+[H+]->[H+]+[H+]+[O-] 2. [O-]+[H+]->[H][O] 3. [H+]+[H][O]->[O-]+[H+]+[H+]
M	1. [H][O]+[H+]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H+]+[H][O]->[H+]+[O-]+[H+]
N	1. [H+]+[O-]->[H][O] 2. [H][O]+[H+]->[H+]+[H+]+[O-] 3. [H+]+[O-]->[H][O] 4. [H][O]+[H+]->[O-]+[H+]+[H+]
O	1. [H][O]+[H+]->[H+]+[H+]+[O-] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[O-]+[H+]+[H+]
P	1. [H+]+[H][O]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H][O]+[H+]->[H+]+[H+]+[O-]
Q	1. [H+]+[O-]->[H][O] 2. [H+]+[H][O]->[O-]+[H+]+[H+] 3. [O-]+[H+]->[H][O] 4. [H][O]+[H+]->[H+]+[H+]+[O-]
R	1. [H][O]+[H+]->[H+]+[H+]+[O-] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[O-]+[H+]
S	1. [O-]+[H+]->[H][O] 2. [H][O]+[H+]->[H+]+[H+]+[O-] 3. [H+]+[O-]->[H][O] 4. [H+]+[H][O]->[H+]+[O-]+[H+]
T	1. [O-]+[H+]->[H][O] 2. [H][O]+[H+]->[H+]+[H+]+[O-] 3. [H+]+[O-]->[H][O] 4. [H][O]+[H+]->[H+]+[H+]+[O-]
U	1. [H][O]+[H+]->[H+]+[H+]+[O-] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[H+]+[H+]+[O-]
V	1. [H][O]+[H+]->[H+]+[O-]+[H+] 2. [O-]+[H+]->[H][O] 3. [H+]+[H][O]->[H+]+[H+]+[O-]
W	1. [H+]+[O-]->[H][O] 2. [H][O]+[H+]->[H+]+[O-]+[H+] 3. [O-]+[H+]->[H][O] 4. [H+]+[H][O]->[H+]+[H+]+[O-]
X	1. [H+]+[O-]->[H][O] 2. [H][O]+[H+]->[H+]+[O-]+[H+] 3. [H+]+[O-]->[H][O] 4. [H+]+[H][O]->[H+]+[O-]+[H+]
Y	1. [H][O]+[H+]->[H+]+[O-]+[H+] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[O-]+[H+]
Z	1. [H][O]+[H+]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H][O]+[H+]->[O-]+[H+]+[H+]
a	1. [H+]+[O-]->[H][O] 2. [H][O]+[H+]->[O-]+[H+]+[H+] 3. [O-]+[H+]->[H][O] 4. [H][O]+[H+]->[O-]+[H+]+[H+]
b	1. [H+]+[O-]->[H][O] 2. [H+]+[H][O]->[O-]+[H+]+[H+] 3. [O-]+[H+]->[H][O] 4. [H][O]+[H+]->[H+]+[O-]+[H+]
c	1. [H][O]+[H+]->[O-]+[H+]+[H+] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[H+]+[O-]+[H+]
d	1. [H+]+[H][O]->[H+]+[O-]+[H+] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[H+]+[H+]+[O-]
e	1. [H+]+[O-]->[H][O] 2. [H+]+[H][O]->[H+]+[O-]+[H+] 3. [H+]+[O-]->[H][O] 4. [H][O]+[H+]->[H+]+[H+]+[O-]
f	1. [H+]+[H][O]->[H+]+[H+]+[O-] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[O-]+[H+]
g	1. [H+]+[H][O]->[H+]+[H+]+[O-] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[H+]+[O-]
h	1. [H+]+[H][O]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H+]+[H][O]->[H+]+[O-]+[H+]
i	1. [H][O]+[H+]->[H+]+[H+]+[O-] 2. [H+]+[O-]->[H][O] 3. [H+]+[H][O]->[H+]+[H+]+[O-]
j	1. [H+]+[H][O]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H+]+[H][O]->[O-]+[H+]+[H+]
k	1. [H][O]+[H+]->[O-]+[H+]+[H+] 2. [O-]+[H+]->[H][O] 3. [H+]+[H][O]->[O-]+[H+]+[H+]
l	1. [H+]+[O-]->[H][O] 2. [H][O]+[H+]->[H+]+[H+]+[O-] 3. [O-]+[H+]->[H][O] 4. [H+]+[H][O]->[H+]+[O-]+[H+]
m	1. [H][O]+[H+]->[H+]+[H+]+[O-] 2. [O-]+[H+]->[H][O] 3. [H+]+[H][O]->[H+]+[O-]+[H+]
n	1. [O-]+[H+]->[H][O] 2. [H][O]+[H+]->[H+]+[O-]+[H+] 3. [H+]+[O-]->[H][O] 4. [H][O]+[H+]->[H+]+[O-]+[H+]
o	1. [H+]+[C]=O->[H+]+[C][O] 2. [H+]+[C][O]->[H+]+[C]=O 3. [C]=O+[H+]->[O]+[C]+[H+] 4. [O-][O-]+[O]->[O-]+[O-]+[O] 5. [O-]+[H+]->[H][O] 6. [H][O]+[H+]->[H+]+[O-]+[H+] 7. [H+]+[O-]->[H][O] 8. [H][O]+[H+]->[H+]+[O-]+[H+]
p	1. O=O+[H+]->[H+]+[O]+[O] 2. [H][O]+[O]->[H]O[O] 3. [H]O[O]+[H+]->[H+]+[O][O-]+[H+] 4. [H+]+[O][O-]->[O-]+[H+]+[O] 5. [O-]+[O-]->[O-][O-] 6. [O-][O-]+[O]->[O-]+[O-]+[O] 7. [O-]+[H+]->[H][O] 8. [H][O]+[H+]->[H+]+[O-]+[H+] 9. [H+]+[O-]->[H][O] 10. [H][O]+[H+]->[H+]+[O-]+[H+]
q	1. [H][O]+[H+]->[O-]+[H+]+[H+] 2. [O-]+[O-]->[O-][O-] 3. [O-][O-]+[O]->[O-]+[O-]+[O] 4. [O-]+[H+]->[H][O] 5. [H][O]+[H+]->[H+]+[O-]+[H+] 6. [H+]+[O-]->[H][O] 7. [H][O]+[H+]->[H+]+[O-]+[H+]
r	1. [C]=O+[H+]->[O]+[C]+[H+] 2. [O-][O-]+[O]->[O-]+[O-]+[O] 3. [O-]+[H+]->[H][O] 4. [H][O]+[H+]->[H+]+[O-]+[H+] 5. [H+]+[O-]->[H][O] 6. [H][O]+[H+]->[H+]+[O-]+[H+]
s	1. [H+]+[C][O]->[H+]+[C]=O 2. [C]=O+[H+]->[O]+[C]+[H+] 3. [O-][O-]+[O]->[O-]+[O-]+[O] 4. [O-]+[H+]->[H][O] 5. [H][O]+[H+]->[H+]+[O-]+[H+] 6. [H+]+[O-]->[H][O] 7. [H][O]+[H+]->[H+]+[O-]+[H+]
t	1. [H]O[O]+[H+]->[H+]+[O][O-]+[H+] 2. [H+]+[O][O-]->[O-]+[H+]+[O] 3. [O-]+[O-]->[O-][O-] 4. [O-][O-]+[O]->[O-]+[O-]+[O] 5. [O-]+[H+]->[H][O] 6. [H][O]+[H+]->[H+]+[O-]+[H+] 7. [H+]+[O-]->[H][O] 8. [H][O]+[H+]->[H+]+[O-]+[H+]
u	1. [H][O]+[H+]->[H+]+[O-]+[H+] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[H+]+[O-]+[H+]
v	1. [H+]+[O][O-]->[O-]+[H+]+[O] 2. [O-]+[O-]->[O-][O-] 3. [O-][O-]+[O]->[O-]+[O-]+[O] 4. [O-]+[H+]->[H][O] 5. [H][O]+[H+]->[H+]+[O-]+[H+] 6. [H+]+[O-]->[H][O] 7. [H][O]+[H+]->[H+]+[O-]+[H+]
w	1. [H+]+[O-]->[H][O] 2. [H][O]+[H+]->[O-]+[H+]+[H+] 3. [O-]+[O-]->[O-][O-] 4. [O-][O-]+[O]->[O-]+[O-]+[O] 5. [O-]+[H+]->[H][O] 6. [H][O]+[H+]->[H+]+[O-]+[H+] 7. [H+]+[O-]->[H][O] 8. [H][O]+[H+]->[H+]+[O-]+[H+]
x	1. [H+]+[O-]->[H][O] 2. [H][O]+[O]->[H]O[O] 3. [H]O[O]+[H+]->[H+]+[O][O-]+[H+] 4. [H+]+[O][O-]->[O-]+[H+]+[O] 5. [O-]+[O-]->[O-][O-] 6. [O-][O-]+[O]->[O-]+[O-]+[O] 7. [O-]+[H+]->[H][O] 8. [H][O]+[H+]->[H+]+[O-]+[H+] 9. [H+]+[O-]->[H][O] 10. [H][O]+[H+]->[H+]+[O-]+[H+]
y	1. [O-]+[H+]->[H][O] 2. [H+]+[H][O]->[H+]+[O-]+[H+] 3. [H+]+[O-]->[H][O] 4. [H][O]+[H+]->[O-]+[H+]+[H+] 5. [O-]+[O-]->[O-][O-] 6. [O-][O-]+[O]->[O-]+[O-]+[O] 7. [O-]+[H+]->[H][O] 8. [H][O]+[H+]->[H+]+[O-]+[H+] 9. [H+]+[O-]->[H][O] 10. [H][O]+[H+]->[H+]+[O-]+[H+]
z	1. [H+]+[H][O]->[H+]+[O-]+[H+] 2. [H+]+[O-]->[H][O] 3. [H][O]+[H+]->[O-]+[H+]+[H+] 4. [O-]+[O-]->[O-][O-] 5. [O-][O-]+[O]->[O-]+[O-]+[O] 6. [O-]+[H+]->[H][O] 7. [H][O]+[H+]->[H+]+[O-]+[H+] 8. [H+]+[O-]->[H][O] 9. [H][O]+[H+]->[H+]+[O-]+[H+]

Glossary

Alife Artificial Life.

ANOVA Analysis of Variance.

Artificial Chemistry The computational abstraction of biomolecular processes.

Constructive An artificial chemistry is constructive if it is closed under all reactions. That is, if the products of a reaction may become reactants in turn..

DFA Detrended Fluctuation Analysis.

EA Evolutionary Algorithm.

ENS Evolution by Natural Selection.

Heredity The passing on of traits from parent to offspring.

Heritability The proportion of the total variation between individuals in a given population due to genetic variation, or the degree to which environment or genetics influences the phenotype. (see `https://www.ncbi.nlm.nih.gov/books/NBK22001/`).

Horizontal Gene Transfer "The transfer of genetic material between reproductively isolated species" (Pace et al. 2008).

Hypercycle An organisation of self-replicating molecules connected in an autocatalytic cycle (Eigen 1971)..

IDA Initial Darwinian Ancestor.

LUCA Last Universal Common Ancestor.

ODE Ordinary Differential Equation.

Open-Ended Evolution "the ability to continuously produce novelty and/or complexity" (Banzhaf et al. 2016).

RAF Reflexively autocatalytic and F-generated.

RBN Random Boolean Network.

Replicate "…an independent repeat of each factor combination." (Montgomery 2009).

Run "…when an apparatus has been set up and allowed for function under a specific set of experimental conditions." (Box, Hunter, and Hunter 2005) Essentially a synonym for replicate, but without the sense of membership in a series under the same conditions.

References

Adami, C (2002). "What is complexity?" In: *Bioessays* 24.0265-9247 (Linking), pp. 1085–1094.

Adami, Christoph et al., eds. (2012). *Proceedings of the thirteenth international conference on the simulation and synthesis of living systems (artificial life 13)*.

Alife xii, odense, Denmark, August 2010 (2010). MIT Press.

Altman, Sidney (2013). "The rna-protein world". In: *RNA* 19.5, pp. 589–590. ISSN: 1469-9001. URL: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3677274/.

Arthur, W. Brian (2009). *The nature of technology: what it is and how it evolves*. Free Press, Simon and Schuster, Inc.

Bagley, R. J. and J. Doyne Farmer (1991). "Spontaneous emergence of a metabolism". In: *Artificial Life II, Santa Fe Institute Studies in the Sciences of Complexity, vol.X*. Addi, pp. 93–140.

Banzhaf, Wolfgang (1994). "Self-organisation in a system of binary strings". In: *Artificial Life IV*. MIT Press, pp. 109–118.

Banzhaf, Wolfgang et al. (2016). "Defining and simulating open-ended novelty: requirements, guidelines, and challenges". In: *Theory in Biosciences*, p. 1. ISSN: 1611-7530. DOI: 10.1007/s12064-016-0229-7. URL: http://dx.doi.org/10.1007/s12064-016-0229-7.

Benkö, Gil, Christoph Flamm, and Peter F. Stadler (2003). "A graph-based toy model of chemistry". In: *Journal of Chemical Information and Computer Sciences* 43.4. PMID: 12870897, pp. 1085–1093. DOI: 10.1021/ci0200570. eprint: http://pubs.acs.org/doi/pdf/10.1021/ci0200570.

— (2005). "The toychem package: a computational toolkit implementing a realistic artificial chemistry model". URL: http://www.tbi.univie.ac.at/~xtof/ToyChem/.

Benkö, Gil et al. (2009). "A topological approach to chemical organizations". In: *Artificial Life* 15.1, pp. 71–88. DOI: 10.1162/artl.2009.15.1.15105.

Bourrat, Pierrick (2015). "Reconceptualising evolution by natural selection". PhD thesis. University of Sydney.

Box, George E.P., J. Stuart Hunter, and William G. Hunter (2005). *Statistics for experimenters: design, innovation, and discovery*. 2nd. Wiley-Interscience.

Braakman, Rogier and Eric Smith (2012). "The emergence and early evolution of biological carbon-fixation". In: *PLoS Comput Biol* 8.4, e1002455. DOI: 10.1371/journal.pcbi.1002455. URL: http://dx.doi.org/10.1371%2Fjournal.pcbi.1002455.

Brockwell, Peter J. and Richard A. Davis (2002). *Introduction to time series and forecasting*. 2nd ed. Springer Berlin / Heidelberg.

Cheng, Leslie K.L. and Peter J. Unrau (2010). "Closing the circle: replicating rna with rna". In: *Cold Spring Harbor Perspectives in Biology* 2.10. DOI: 10.1101/cshperspect.a002204. eprint: http://cshperspectives.cshlp.org/content/2/10/a002204.full.pdf+html. URL: http://cshperspectives.cshlp.org/content/2/10/a002204.abstract.

Darwin, Charles (1859). *On the origin of species by means of natural selection, or, the preservation of favoured races in the struggle for life*. London: J. Murray.

Dawkins, Richard (1976). *The selfish gene*. Oxford University Press.

Daylight Chemical Information Systems, Inc. (2011). *Daylight theory manual*. URL: http://www.daylight.com/dayhtml/doc/theory/index.html.

De Jong, Kenneth A. (2006). *Evolutionary computation : a unified approach*. Cambridge, Mass.: MIT Press, pp. ix, 256.

Dittrich, Peter, Jens Ziegler, and Wolfgang Banzhaf (2001). "Artificial chemistries-a review". In: *Artificial Life* 7.3, pp. 225–275. DOI: 10.1162/106454601753238636. eprint: http://www.mitpressjournals.org/doi/pdf/10.1162/106454601753238636.

Doolittle, W. Ford and Eric Bapteste (2007). "Pattern pluralism and the Tree of Life hypothesis". In: *Proceedings of the National Academy of Sciences* 104.7, pp. 2043–2049. DOI: 10.1073/pnas.0610699104. eprint: http://www.pnas.org/content/104/7/2043.full.pdf. URL: http://www.pnas.org/content/104/7/2043.abstract.

Dorin, A. and K. B. Korb (2006). "Building virtual ecosystems from artificial chemistry". In: *9th European Conference on Artificial Life, ECAL 2006*. Ed. by F. A. E. Costa et al. Springer-Verlag Berlin, pp. 103–112.

Ducharme, Vincent, Richard Egli, and Claude Y. Legault (2012). "Energy-based artificial chemistry simulator". In: *Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems (Artificial Life 13)*. Ed. by Christoph Adami et al., pp. 449–456.

Eigen, Manfred (1971). "Selforganization of matter and the evolution of biological macromolecules". In: *Naturwissenschaften* 58 (10), pp. 465–523. ISSN: 0028-1042. DOI: 10.1007/BF00623322.

Farmer, J.Doyne, Stuart A Kauffman, and Norman H Packard (1986). "Autocatalytic replication of polymers". In: *Physica D: Nonlinear Phenomena* 22.1,Äì3. Proceedings of the Fifth Annual International Conference, pp. 50–67. ISSN: 0167-2789. DOI: 10.1016/0167-

2789(86)90233-2. URL: http://www.sciencedirect.com/science/article/pii/0167278986902332.

Faulconbridge, Adam (2011). "Rbn-world: sub-symbolic artificial chemistry for artificial life". PhD thesis. University of York.

Faulconbridge, Adam et al. (2010). "RBN-world: the hunt for a rich AChem". In: *ALife XII, Odense, Denmark, August 2010*. MIT Press, pp. 261–268.

Faulon, Jean-Loup and Allen G. Sault (2001). "Stochastic generator of chemical structure. 3. reaction network generation". In: *Journal of Chemical Information and Computer Sciences* 41.4. PMID: 11500106, pp. 894–908. DOI: 10.1021/ci000029m. eprint: http://pubs.acs.org/doi/pdf/10.1021/ci000029m.

Fernando, Chrisantha and J. Rowe (2007). "Natural selection in chemical evolution". In: *Journal of Theoretical Biology* 247.1, pp. 152–167.

Fernando, Chrisantha and Jonathan Rowe (2008). "The origin of autonomous agents by natural selection". In: *Biosystems* 91.2. 0303-2647 doi: DOI: 10.1016/j.biosystems.2007.05.012, pp. 355–373.

Flamm, Christoph et al. (2010). "Evolution of metabolic networks: a computational framework". In: *Journal of Systems Chemistry* 1.1, p. 4. ISSN: 1759-2208. DOI: 10.1186/1759-2208-1-4. URL: http://www.jsystchem.com/content/1/1/4.

Fontana, W (1992). "Algorithmic chemistry". In: *Artificial Life II*. Ed. by C.G. Langton et al. Vol. 10. Santa Fe Institute Studies In The Sciences Of Complexity - Proceedings Volumes. Reading: Addison-Wesley, pp. 159–209. ISBN: 0-201-52571-2.

Fontana, W. and L. W. Buss (1994). "What would be conserved if "the tape were played twice"?" In: *Proceedings of the National Academy of Sciences of the United States of America* 91.2, pp. 757–761. ISSN: 1091-6490. URL: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC43028/.

Fontana, W., G. P. Wagner, and L. W. Buss (1994). "Beyond digital naturalism". In: *Artificial Life* 1.2, pp. 211–227.

Funes, Pablo (2001). "Evolution of complexity in real-world domains". PhD thesis. Department of Computer Science, Brandeis University.

Gánti, Tibor (2003). *The principles of life*. Oxford ; New York: Oxford University Press, p. 201.

Gardiner, Ken, James Harland, and Margaret Hamilton (2007). "A framework for the co-evolution of genes, proteins and a genetic code within an artificial chemistry reaction set". In: *Progress in Artificial Life, Proceedings*. Ed. by M Randall, HA Abbass, and J Wiles. Vol. 4828. Lecture Notes in Artificial Intelligence. 3rd Australian Conference on Artificial Life, Gold Coast, AUSTRALIA, DEC 04-06, 2007, pp. 281–291. ISBN: 978-3-540-76930-9.

Gaucherel, Cédric and Henrik Jeldtoft Jensen (2012). "Origins of evolution: non-acquired characters dominates over acquired characters in changing environment". In: *Journal of*

Theoretical Biology 304, pp. 111–120. ISSN: 0022-5193. DOI: 10.1016/j.jtbi.2012.02.028. URL: http://www.sciencedirect.com/science/article/pii/S0022519312001130.

Gillespie, Daniel T (1976). "A general method for numerically simulating the stochastic time evolution of coupled chemical reactions". In: *Journal of Computational Physics* 22.4, pp. 403–434. ISSN: 0021-9991. DOI: 10.1016/0021-9991(76)90041-3. URL: http://www.sciencedirect.com/science/article/pii/0021999176900413.

Godfrey-Smith, Peter (2007). "Conditions for evolution by natural selection". In: *The Journal of Philosophy* 104.10, pp. 489–516. ISSN: 0022362X. URL: http://www.jstor.org/stable/20620050.

Griesemer, J. (2001). "The units of evolutionary transition". In: *Selection* 1.1-3, pp. 67–80. DOI: 10.1556/Select.1.2000.1-3.7. eprint: http://dx.doi.org/10.1556/Select.1.2000.1-3.7.

— (2005). "The informational gene and the substantial body: on the generalization of evolutionary theory by abstraction". In: *Idealization XII: Correcting the Model, Idealization and Abstraction in the Sciences*. Ed. by Martin R. Jones and Nancy Cartwright. Vol. 86. Poznan Studies in the Philosophy of the Sciences and the Humanities. Amsterdam/New York, NY: Rodopi, pp. 59–115.

Hartenfeller, Markus et al. (2011). "A collection of robust organic synthesis reactions for in silico molecule design". In: *Journal of Chemical Information and Modeling* 51.12, pp. 3093–3098. DOI: 10.1021/ci200379p. eprint: http://pubs.acs.org/doi/pdf/10.1021/ci200379p.

— (2012). "Probing the bioactivity-relevant chemical space of robust reactions and common molecular building blocks". In: *Journal of Chemical Information and Modeling* 52.5, pp. 1167–1178. DOI: 10.1021/ci200618n. eprint: http://pubs.acs.org/doi/pdf/10.1021/ci200618n.

Hickinbotham, Simon et al. (2011). "Molecular microprograms". In: *Advances in Artificial Life. Darwin Meets von Neumann*. Ed. by George Kampis, István Karsai, and Eörs Szathmáry. Vol. 5777. Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 297–304. ISBN: 978-3-642-21282-6. DOI: 10.1007/978-3-642-21283-3_37.

— (2012). *Specification of the stringmol chemical programming language version 0.2*. Tech. rep. YCS-2010-458. Univ. of York.

Högerl, Daniel (2010). "Simulation of prebiotic chemistries". MA thesis. Institute for Theoretical Chemistry, University of Vienna.

Hordijk, Wim, Stuart A. Kauffman, and Mike Steel (2011). "Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems". In: *International Journal of Molecular Sciences* 12.5, pp. 3085–3101. ISSN: 1422-0067. DOI: 10.3390/ijms12053085. URL: http://www.mdpi.com/1422-0067/12/5/3085.

Hordijk, Wim, Joshua Smith, and Mike Steel (2015). "Algorithms for detecting and analysing autocatalytic sets". In: *Algorithms for Molecular Biology* 10.1, p. 15. ISSN: 1748-7188. DOI: 10.1186/s13015-015-0042-8. URL: http://www.almob.org/content/10/1/15.

Hutton, Tim J. (2002). "Evolvable self-replicating molecules in an artificial chemistry". In: *Artificial Life* 8.4, pp. 341–356. URL: http://www.sq3.org.uk/Evolution/Squirm3/EvSelfReps/.

— (2007). "Evolvable self-reproducing cells in a two-dimensional artificial chemistry". In: *Artificial Life* 13.1, pp. 11–30. DOI: 10.1162/artl.2007.13.1.11.

— (2009). "The organic builder: a public experiment in artificial chemistries and self-replication". In: *Artificial Life* 15.1, pp. 21–28. DOI: 10.1162/artl.2009.15.1.15102.

Jablonka, Eva et al. (1995). "The Adaptive Advantage of Phenotypic Memory in Changing Environments". In: *Philosophical Transactions: Biological Sciences* 350.1332, pp. 133–141. ISSN: 09628436. URL: http://www.jstor.org/stable/56329.

Jeschke, Matthias et al. (2008). "Parallel and distributed spatial simulation of chemical reactions". In: *Proceedings of the 22nd Workshop on Principles of Advanced and Distributed Simulation*. PADS '08. Washington, DC, USA: IEEE Computer Society, pp. 51–59. ISBN: 978-0-7695-3159-5. DOI: 10.1109/PADS.2008.20.

Johnson, Donald B. (1975). "Finding All the Elementary Circuits of a Directed Graph". English. In: *SIAM Journal on Computing* 4.1. Copyright - Copyright] © 1975 Society for Industrial and Applied Mathematics; Last updated - 2012-01-30, pp. 77–8. URL: http://search.proquest.com.ezproxy.canterbury.ac.nz/docview/918546620?accountid=14499.

Kauffman, Stuart A. (1969). "Metabolic stability and epigenesis in randomly constructed genetic nets". In: *Journal of Theoretical Biology* 22.3, pp. 437–467.

— (1986). "Autocatalytic sets of proteins". In: *Journal of Theoretical Biology* 119.1, pp. 1–24. ISSN: 0022-5193. DOI: 10.1016/S0022-5193(86)80047-9. URL: http://www.sciencedirect.com/science/article/pii/S0022519386800479.

Kimura, M. (1968). "Evolutionary rate at the molecular level". In: *Nature* 217.5129, pp. 624–626.

King, G. A. M. (1978). "Autocatalysis". In: *Chem. Soc. Rev.* 7 (2), pp. 297–316. DOI: 10.1039/CS9780700297. URL: http://dx.doi.org/10.1039/CS9780700297.

Kotz, John C., Paul M. Treichel, and Gabriela C. Weaver (2006). *Chemistry & chemical reactivity*. 6th edition. Thomson Brooks/Cole.

Landrum, Greg (2013). *Rdkit: open-source cheminformatics*. URL: http://www.rdkit.org.

Lenaerts, Tom and Hugues Bersini (2009). "A synthon approach to artificial chemistry". In: *Artificial Life* 15.1, pp. 89–103. DOI: 10.1162/artl.2009.15.1.15106.

Lewontin, R. C. (1970). *The units of selection*.

Lifson, Shneior (1997). "On the crucial stages in the origin of animate matter". In: *Journal of Molecular Evolution* 44.1, pp. 1–8. ISSN: 0022-2844. DOI: 10.1007/PL00006115.

Lincoln, Tracey A. and Gerald F. Joyce (2009). "Self-sustained replication of an rna enzyme". In: *Science* 323.5918, pp. 1229–1232. DOI: 10.1126/science.1167856. eprint: http://www.sciencemag.org/content/323/5918/1229.full.pdf. URL: http://www.sciencemag.org/content/323/5918/1229.abstract.

Lovelock, James E. and Lynn Margulis (2011). "Atmospheric homeostasis by and for the biosphere: the gaia hypothesis". English. In: *Tellus A* 26.1-2.

Lucht, Michael W. (2012). "Size selection and adaptive evolution in an artificial chemistry". In: *Artificial Life* 18.2, pp. 143–163. ISSN: 1064-5462.

Madina, Duraid, Naoaki Ono, and Takashi Ikegami (2003). "Cellular evolution in a 3D lattice artificial chemistry". In: *Advances in Artificial Life*. Ed. by Wolfgang Banzhaf et al. Vol. 2801. Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 59–68. ISBN: 978-3-540-20057-4.

Markovitch, Omer et al. (2012). "Is there an optimal level of open-endedness in prebiotic evolution?" In: *Origins of Life and Evolution of Biospheres* 42.5, pp. 469–474. ISSN: 0169-6149. DOI: 10.1007/s11084-012-9309-y.

McShea, Daniel W. (1998). "Possible largest-scale trends in organismal evolution: eight "live hypotheses"". In: *Annual Review of Ecology and Systematics* 29, pp. 293–318. ISSN: 00664162. URL: http://www.jstor.org/stable/221710.

Montgomery, Douglas C. (2009). *Design and analysis of experiments*. 7th Edition. John Wiley & Sons, Inc.

Muller, H. J. (1964). "The Relation of Recombination to Mutational Advance". In: *Mutation Research* 106, pp. 2–9.

Nellis, Adam (2012). "Towards meta-evolution via embodiment in artificial chemistries". PhD thesis. University of York. URL: http://etheses.whiterose.ac.uk/id/eprint/3701.

Nellis, Adam and Susan Stepney (2014). "Computational novelty: phenomena, mechanisms, worlds". In: DOI: 10.7551/978-0-262-32621-6-ch081. URL: https://mitpress.mit.edu/sites/default/files/titles/content/alife14/978-0-262-32621-6-ch081.pdf.

Ochman, Howard, Jeffrey G. Lawrence, and Eduardo A. Groisman (2000). "Lateral gene transfer and the nature of bacterial innovation". In: *Nature* 405.6784, pp. 299–304. DOI: 10.1038/35012500.

Ofria, Charles and Claus O. Wilke (2004). "Avida: a software platform for research in computational evolutionary biology". In: *Artificial Life* 10.2, pp. 191–229. ISSN: 1064-5462. DOI: 10.1162/106454604773563612.

Ono, Naoaki and Takashi Ikegami (2000). "Self-maintenance and self-reproduction in an abstract cell model". In: *Journal of Theoretical Biology* 206.2, pp. 243–253. ISSN: 0022-5193. DOI: 10.1006/jtbi.2000.2121. URL: http://www.sciencedirect.com/science/article/pii/S0022519300921210.

Ostwald, Wilhelm (1890). "Über autokatalyse". In: *Berichte über die Verhandlungen der Königlich-Sächsischen Gesellschaft der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse* 42, pp. 189–191.

Pace, John K. et al. (2008). "Repeated horizontal transfer of a DNA transposon in mammals and other tetrapods". In: *Proceedings of the National Academy of Sciences of the United States of America* 105.44, pp. 17023–17028.

Paenke, Ingo et al. (2007). "On the adaptive disadvantage of lamarckianism in rapidly changing environments". In: *Advances in Artificial Life*, pp. 355–364. DOI: 10.1007/978-3-540-74913-4_36.

Paixão, Tiago et al. (2015). "Toward a unifying framework for evolutionary processes". In: *Journal of Theoretical Biology* 383, pp. 28–43. ISSN: 0022-5193. DOI: 10.1016/j.jtbi.2015.07.011. URL: http://www.sciencedirect.com/science/article/pii/S0022519315003409.

Pargellis, A. N. (2001). "Digital Life Behavior in the Amoeba World". In: *Artificial Life* 7.1, pp. 63–75. ISSN: 1064-5462. DOI: 10.1162/106454601300328025. URL: http://dx.doi.org/10.1162/106454601300328025.

Pigliucci, M. (2008). "Opinion - is evolvability evolvable?" In: *Nature Reviews Genetics* 9.1, pp. 75–82.

Plasson, Raphael et al. (2010). "Autocatalyses". In: *ALife XII, Odense, Denmark, August 2010*. MIT Press.

Powner, Matthew W., Beatrice Gerland, and John D. Sutherland (2009). "Synthesis of activated pyrimidine ribonucleotides in prebiotically plausible conditions". In: *NATURE* 459.7244, pp. 239–242. ISSN: 0028-0836. DOI: 10.1038/nature08013.

Prokopenko, Mikhail, Fabio Boschetti, and Alex J. Ryan (2009). "An information-theoretic primer on complexity, self-organization, and emergence". In: *Complexity* 15.1, pp. 11–28. ISSN: 1099-0526. DOI: 10.1002/cplx.20249.

Pross, Addy and Robert Pascal (2013). "The origin of life: what we know, what we can know and what we will never know". In: *Open Biology* 3.3. DOI: 10.1098/rsob.120190.

Ragan, Mark A., James O. McInerney, and James A. Lake (2009). "The network of life: genome beginnings and evolution". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 364.1527, pp. 2169–2175. DOI: 10.1098/rstb.2009.0046. eprint: http://rstb.royalsocietypublishing.org/content/364/1527/2169.full.pdf+html. URL: http://rstb.royalsocietypublishing.org/content/364/1527/2169.short.

Rasmussen, Steen et al. (1990). "The coreworld: emergence and evolution of cooperative structures in a computational chemistry". In: *Physica D: Nonlinear Phenomena* 42.1–3, pp. 111–134. ISSN: 0167-2789. DOI: 10.1016/0167-2789(90)90070-6. URL: http://www.sciencedirect.com/science/article/pii/0167278990900706.

Ray, T.S. (1991). "An approach to the synthesis of life". In: *2nd International Conference on Artificial Life, Alife II*. Ed. by C Langton. Addison-Wesley, pp. 371–408.

Richman, Joshua S. and J. R. Moorman (2000). "Physiological time-series analysis using approximate entropy and sample entropy". English. In: *American Journal of Physiology - Heart and Circulatory Physiology* 278.6, pp. 2039–2049.

Sasaki, Takahiro and Mario Tokoro (2000). "Comparison between lamarckian and darwinian evolution on a model using neural networks and genetic algorithms". In: *Knowledge and Information Systems* 2 (2), pp. 201–222. ISSN: 0219-1377. DOI: 10.1007/s101150050011.

Saunders, Peter T. (1994). "Evolution without natural selection: further implications of the daisyworld parable". In: *Journal of Theoretical Biology* 166.4, pp. 365–373.

Schuster, Peter (2011). "Mathematical modeling of evolution. solved and open problems". In: *Theory in Biosciences* 130 (1), pp. 71–89. ISSN: 1431-7613. DOI: 10.1007/s12064-010-0110-z.

Segré, Daniel et al. (1998). "Graded autocatalysis replication domain (gard): kinetic analysis of self-replication in mutually catalytic sets". In: *Origins of life and evolution of the biosphere* 28.4-6, pp. 501–514. ISSN: 0169-6149. DOI: 10.1023/A:1006583712886.

Shenhav, Barak, Aia Oz, and Doron Lancet (2007). "Coevolution of compositional protocells and their environment". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 362.1486, pp. 1813–1819. ISSN: 1471-2970. URL: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2442395/.

Shenhav, Barak et al. (2005). "Polymer Gard: Computer Simulation of Covalent Bond Formation in Reproducing Molecular Assemblies". In: *Origins of Life and Evolution of Biospheres* 35.2, pp. 111–133. ISSN: 1573-0875. DOI: 10.1007/s11084-005-5578-z. URL: http://dx.doi.org/10.1007/s11084-005-5578-z.

Steel, Mike, Wim Hordijk, and Joshua Smith (2013). "Minimal autocatalytic networks". In: *Journal of Theoretical Biology* 332.0, pp. 96–107. ISSN: 0022-5193. DOI: 10.1016/j.jtbi.2013.04.032. URL: http://www.sciencedirect.com/science/article/pii/S0022519313002002.

Sugiura, Komei et al. (2003). "Evolution of Rewriting Rule Sets Using String-Based Tierra". In: *Advances in Artificial Life*. Springer. ISBN: 978-3-540-20057-4. DOI: 10.1007/978-3-540-39432-7_8. URL: http://dx.doi.org/10.1007/978-3-540-39432-7_8.

Suzuki, H (2006). "A molecular network rewiring rule that represents spatial constraint". In: *International Conference on Complex Systems (ICCS2006)*. URL: http://necsi.edu/events/iccs6/viewabstract.php?id=313.

Suzuki, Hideaki and Peter Dittrich (2008). "Artificial chemistry". In: *Artificial Life* 15.1, pp. 1–3. DOI: 10.1162/artl.2009.15.1.15100.

Suzuki, Keisuke and Takashi Ikegami (2008). "Shapes and self-movement in protocell systems". In: *Artificial Life* 15.1, pp. 59–70. DOI: 10.1162/artl.2009.15.1.15104.

Szathmáry, E. (1999). "Chemes, genes, memes: a revised classification of replicators". In: *Lect. Math. Life Sci.* 26, pp. 1–10. URL: http://www.researchgate.net/publication/266076188_Chemes_genes_memes_A_revised_classification_of_replicators.

Szathmáry, Eörs (2006). "The origin of replicators and reproducers". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 361.1474, pp. 1761–1776.

Taylor, T (1999a). "On self-reproduction and evolvability". In: *ADVANCES IN ARTIFICIAL LIFE, PROCEEDINGS*. Ed. by D Floreano, JD Nicoud, and F Mondada. Vol. 1674. LECTURE NOTES IN ARTIFICIAL INTELLIGENCE. 5th European Conference on Artificial Life (ECAL), LAUSANNE, SWITZERLAND, SEP 13-17, 1999, pp. 94–103. ISBN: 3-540-66452-1.

Taylor, Timothy (1999b). "From artificial evolution to artificial life". PhD thesis.

— (2001). "Creative evolutionary systems". In: ed. by D.W. Bentley P.J. Corne. Morgan Kaufman. Chap. Creativity in evolution: Individuals, interactions and environments, pp. 74–108.

Tominaga, Kazuto (2004). *A formal model based on affinity among elements for describing behavior of complex systems*. Tech. rep. Technical report UIUCDCS-R-2004-2413. Department of Computer Science, University of Illinois at Urbana-Champaign. URL: http://hdl.handle.net/2142/10785.

Tominaga, Kazuto et al. (2007). "Modeling molecular computing systems by an artificial chemistry–its expressive power and application". In: *Artificial Life* 13.3, pp. 223–247. DOI: 10.1162/artl.2007.13.3.223.

Ullrich, Alexander et al. (2010). "In silico evolution of early metabolism". In: *Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems (Artificial Life XII)*. Ed. by Harold Fellermann et al. The MIT Press, Cambridge, MA.

Varela, F. G., H. R. Maturana, and R. Uribe (1974). "Autopoiesis: the organization of living systems, its characterization and a model". In: *Biosystems* 5.4, pp. 187–196.

Vasas, Vera and Chrisantha Fernando (2012). "Evolution before life". In: *Models of the Ecological Hierarchy – From Molecules to the Ecosphere*. Ed. by Ferenc Jordán and Sven Erik Jørgensen. Vol. 25. Developments in Environmental Modelling. Elsevier, pp. 15–33. DOI: 10.1016/B978-0-444-59396-2.00002-X.

Vasas, Vera, Eörs Szathmáry, and Mauro Santos (2010). "Lack of evolvability in self-sustaining autocatalytic networks constraints metabolism-first scenarios for the origin of life". In: *Proceedings of the National Academy of Sciences* 107.4, pp. 1470–1475. DOI: 10.1073/pnas.

0912628107. eprint: http://www.pnas.org/content/107/4/1470.full.pdf+html. URL: http://www.pnas.org/content/107/4/1470.abstract.

Vasas, Vera et al. (2012). "Evolution before genes". In: *Biology Direct* 7.1, p. 1. ISSN: 1745-6150. DOI: 10.1186/1745-6150-7-1. URL: http://www.biology-direct.com/content/7/1/1.

Vasas, Vera et al. (2015). "Primordial evolvability: impasses and challenges". In: *Journal of Theoretical Biology* 381. Tibor Gánti (1933- 2009): Towards the Principles of Life and Systems Chemistry, pp. 29–38. ISSN: 0022-5193. DOI: 10.1016/j.jtbi.2015.06.047. URL: http://www.sciencedirect.com/science/article/pii/S0022519315003264.

Virgo, Nathaniel and Takashi Ikegami (2013). "Autocatalysis before enzymes: the emergence of prebiotic chain reactions". In: ed. by Pietro Liò et al., pp. 240–247. DOI: 10.7551/978-0-262-31709-2-ch036. URL: http://mitpress.mit.edu/sites/default/files/titles/content/ecal13/ch036.html.

Von Neumann, John (1966). *Theory of self-reproducing automata*. Ed. by Arthur W. Burks. Urbana: University of Illinois Press.

Vose, Michael D. (1999). *The simple genetic algorithm : foundations and theory*. Complex adaptive systems. Cambridge, Mass.: MIT Press.

Wagner, Andreas (2008). "Robustness and evolvability: a paradox resolved". In: *Proceedings of the Royal Society B-Biological Sciences* 275.1630, pp. 91–100.

Wagner, G. P. and L. Altenberg (1996). "Perspective: complex adaptations and the evolution of evolvability". In: *Evolution* 50.3, pp. 967–976.

Wagner, G. P. and P. Krall (1993). "What is the difference between models of error thresholds and Muller's ratchet?" In: *Journal of Mathematical Biology* 32.1, pp. 33–44. ISSN: 1432-1416. DOI: 10.1007/BF00160372. URL: http://dx.doi.org/10.1007/BF00160372.

Watson, Richard A. (2012). "Is evolution by natural selection the algorithm of biological evolution?" In: *Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems (Artificial Life 13)*. Ed. by Charles Ofria Christoph Adami David M. Bryson and Robert T. Pennock, pp. 121–128.

Watson, Richard A. et al. (2010). "Adaptation without natural selection". In: *ALife XII, Odense, Denmark, August 2010*. MIT Press.

Wilson, Timothy J. and David M. J. Lilley (2009). "The evolution of ribozyme chemistry". In: *Science* 323.5920, pp. 1436–1438. DOI: 10.1126/science.1169231. eprint: http://www.sciencemag.org/content/323/5920/1436.full.pdf. URL: http://www.sciencemag.org/content/323/5920/1436.short.

Yarus, Michael (2011). "Getting past the rna world: the initial darwinian ancestor". In: *Cold Spring Harbor Perspectives in Biology* 3.4. DOI: 10.1101/cshperspect.a003590. eprint: http://cshperspectives.cshlp.org/content/3/4/a003590.full.pdf+html. URL: http://cshperspectives.cshlp.org/content/3/4/a003590.abstract.

Young, Thomas and Kourosh Neshatian (2013). "A constructive artificial chemistry to explore open-ended evolution". In: *AI 2013: Advances in Artificial Intelligence*. Ed. by Stephen Cranefield and Abhaya Nayak. Vol. 8272. LNCS.

Young, Thomas J. (2013a). *Toyworld: a parameterized chemistry for artificial life experimentation*. URL: https://github.com/th0mmeke/toyworld.

— (2013b). *Toyworld2*. URL: https://github.com/th0mmeke/toyworld2.

Young, Thomas J. and Kourosh Neshatian (2015). "The effect of reactant and product selection strategies on cycle evolution in an artificial chemistry". In: *Artificial Life and Computational Intelligence*. Ed. by Stephan K. Chalup, Alan D. Blair, and Marcus Randall. Vol. 8955. Lecture Notes in Computer Science. Springer International Publishing, pp. 310–322. ISBN: 978-3-319-14802-1. DOI: 10.1007/978-3-319-14803-8_24.

Zachar, Istvan and Eors Szathmáry (2010). "A new replicator: a theoretical framework for analysing replication". In: *BMC Biology* 8.1, p. 21. ISSN: 1741-7007. DOI: 10.1186/1741-7007-8-21. URL: http://www.biomedcentral.com/1741-7007/8/21.