Summer 7-14-2015

The Role of Environmental Dynamics in the Emergence of Autocatalytic Networks

Joe Fusion
Portland State University

Let us know how access to this document benefits you.

Follow this and additional works at: http://pdxscholar.library.pdx.edu/open_access_etds

Part of the Catalysis and Reaction Engineering Commons

The Role of Environmental Dynamics in the Emergence of Autocatalytic Networks

by

Joe Fusion

A dissertation submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy
in
Systems Science

Dissertation Committee:
Martin Zwick, Chair
Mark Bedau
Wayne Wakeland
Niles Lehman

Portland State University
2015

ABSTRACT

For life to arise from non-life, a metabolism must emerge and maintain itself, distinct from its environment. One line of research seeking to understand this emergence has focused on models of autocatalytic reaction networks (ARNs) and the conditions that allow them to approximate metabolic behavior. These models have identified reaction parameters from which a proto-metabolism might emerge given an adequate matter-energy flow through the system. This dissertation extends that research by answering the question: can dynamically structured interactions with the environment promote the emergence of ARNs? This question was inspired by theories that place the origin of life in contexts such as diurnal or tidal cycles. To answer it, an artificial chemistry system with ARN potential was implemented in the dissipative particle dynamics (DPD) modeling paradigm. Unlike differential equation (DE) models favored in prior ARN research, the DPD model is able to simulate environmental dynamics interacting with discrete particles, spatial heterogeneity, and rare events. This dissertation first presents a comparison of the DPD model to published DE results, showing qualitative similarity with some interesting differences. Multiple examples are then provided of dynamically changing flows from the environment that promote emergent ARNs more than constant flows. These include specific cycles of energy and mass flux that consistently increase metrics for ARN concentration and mass focusing. The results also demonstrate interesting nonlinear interactions between the system and cycle amplitude and period. These findings demonstrate the relevance that environmental dynamics has to ARN research and the potential for broader application as well.

ACKNOWLEDGMENTS

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS AND SYMBOLS

ABM – agent-based modeling

AChem – artificial chemistry

ALife – artificial life

ARN – autocatalytic reaction network

B-F – Bagley & Farmer (specifically, the 1991 publication)

DE – differential equation

DPD – dissipative particle dynamics

GARD – graded autocatalytic replication domain

OOL – origin of life

R – mass ratio

SD – system dynamics

Λ – slope of concentration profile

1. INTRODUCTION

In the beginning of modern research into life's origins, the great challenge was to determine what natural process could have produced the molecular building blocks of life. When Miller (1953) brought a spark of light to this darkness, it may have seemed like the answers would soon be found. Instead, their experiment revealed new challenges, and new gaps between the primordial soup and modern life. In the years since, some of those challenges have been resolved, in whole or in part, and new ones have arisen. Today, there are many known natural sources of organic molecules, from extra-planetary amino acids to geothermal nucleotides, but critical questions remain in the story of life's emergence (Branciamore et al. 2009; Koonin & Martin 2005).

One unanswered question is how, from the combinatorial sea of possible molecules, a living system could emerge that is composed of relatively few types, but is able to maintain itself amidst the variety. Life as we know it solves the problem with three interlinked internal subsystems: a metabolism to ingest and convert food (matter-energy); information storage to serve as template to identity and allow replication, short-term adaptation, and long-term (populational) evolution; and a membrane to create and control a boundary with the environment (Gánti 2003). These three systems appear to be inextricably linked, so none could exist without the others, and there is not yet an accepted model for how they might arise by natural process.

Many of the candidate models for the Origin of Life (OOL) theorize the initial emergence of one subsystem before the others, as in the membrane-first 'Lipid World' (Segré et al. 2001), or of a single structure that fulfills multiple functions, as in the combined metabolism and template of the 'RNA World' (Higgs & Lehman 2015). Some

OOL models have been constructed around specific environments or conditions where life could have emerged, such as tide pools or deep-sea thermal vents. The environmental aspects of these models include factors that may have been critical to abiogenesis (life from non-life), such as mineral surfaces with specific catalytic properties, or a consistent source of energy and food (Corliss 1990; Russell & Hall 1997). Among these environments, those that contribute a structured, dynamic interaction to the chemistry of the model system are of particular interest, and are central to this dissertation. The dynamic structure of these features could derive from natural periodicity, such as the day/night cycle or the ebb and flow of tides, and could interact with the chemistry by providing free energy, by evaporation or condensation of a solution, and so on. To an emerging proto-life system, this type of dynamic interaction may act as a predictable source of variety or as a selective force, which is thought to be conducive to the emergence of complex adaptive systems such as life (Hazen 2005). For example, a periodic inflow of new molecules might drive a select set of chemical reactions, so that both reactants and catalysts remain in the system at levels far from equilibrium. A constant high flow would favor the new molecules, washing away the reaction products, while a constant low inflow would fail to drive the system, with catalyzed reactions undistinguished from steady-state level. Even a constant moderate flow, cumulatively equal to the periodic flow, could have a deficient outcome if it never provides the surge of metabolites that helps the set of reactions to become established, perhaps failing to focus mass into robust loops or larger molecular products.

While some origin-of-life theories have explicit connections to an environment, many others do not. In the field of Artificial Life (ALife), models of life and of

emergence are typically simple and abstract, with little consideration for environmental interactions (Langton 1984). This can be by design, as some research seeks to discover the minimal system necessary for a behavior, or it may be pragmatic, limited by available computational power or laboratory tools. Within these limits, ALife models have been successful in a number of ways. One area of success has been the study of autocatalytic reaction networks (ARNs), which are sets of mutually catalytic molecules and reactions, able to form reinforcing feedback loops and thereby maintain higher concentrations than other molecule species, as long as the reactions continue to be driven (Bagley & Farmer 1991; Bagley et al. 1991). The model systems in this research are sometimes analogized to amino acids or nucleotides (Kauffman 1986), but more often they use highly abstracted chemistries to investigate the behaviors of differential equation kinetics and graph-theoretic relationships (Fraser & Reidys 1997; Hordijk & Steel 2004). Even so, research into ARNs has delivered promising results, and continues to push in new directions (Hordijk et al. 2014; Vasas et al. 2010).

The challenge presented by these abstract findings is how to relate them to other research in the Origin of Life or in Artificial Life that uses different modeling paradigms, with different assumptions and different levels of physical validity (realism). The present research addresses this challenge, as it applies to the question of whether dynamically structured interactions with the environment promote the emergence of ARNs, by extending ARN research to another modeling paradigm. While the present research did not aim for strict physical or chemical validity, it was structured so that its qualitative findings could be generally applicable, and so the choice of modeling paradigm might enable increased physical realism in future work. Additional details are given below.

Research Objectives

The primary objective of this dissertation project, as stated above, was to determine the extent to which dynamic environmental interactions can promote the emergence of autocatalytic reaction networks. The expected results for this objective were to identify dynamic interactions that increase both the *probability* of ARN emergence relative to base conditions, and also the *variety* of system parameters in which emergence occurs. (It was also expected that some dynamic interactions would decrease emergence or have little effect.) The emergence of an ARN within a system of artificial chemical reactions is marked by the ARN's ability to maintain substantially higher concentrations of its molecules, relative to other molecules in the same system or a steady-state system without catalysis (Bagley & Farmer 1991). The distance of this separation from background or steady state levels also measures the strength of the ARN, and can be calculated in terms of *concentration profile* (mean counts of same-length molecules, simplified to a single scalar value) or *mass focusing* (how much of the chemical system's mass is focused into the ARN, expressed as a percent or ratio) (Bagley & Farmer 1991). For either metric, measured values can be used to quickly compare instances of the model that have different parameters.

In order to maximize the likelihood of success on the primary objective, it was determined that a preliminary step and objective would be necessary first. The preliminary step was the implementation of an ARN-capable system in a different modeling paradigm than what has been used in the existing research. The matching objective was to establish a basic level of comparability between the outputs of a model from the new paradigm and previously published results. The determination to switch

modeling paradigms was informed by general review of modeling techniques (Borshchev & Filippov 2004; Epstein 1999; Minar et al. 1996; Scholl 2001), by the specifics of the prior research, and by the research question at hand. Previous research into ARNs has primarily used models from the differential equation (DE) paradigm (Bagley & Farmer 1991; Farmer et al. 1991), which have certain advantages and disadvantages. For example, DE models work well for large, well-mixed populations, and for system attributes measured at the aggregate level, as with concentrations or reaction rates (Borshchev & Filippov 2004). The disadvantages include difficulties in tracking individual particles, in modeling heterogeneity of location or attribute, and in local interactions between particles. These disadvantages limit the usefulness of the DE modeling paradigm for this project's primary objective of identifying increased variety and for its focus on system-environment interactions. Research into available modeling paradigms instead suggested the use of a multi-agent (or agent-based) modeling approach. The multi-agent modeling approach facilitates qualitatively different system features than the DE approach, including spatial heterogeneity, local interactions, and greater sensitivity to small effects and randomness (Scholl 2001). Within the broader category of multi-agent models, the dissipative particle dynamics (DPD) model paradigm was selected to use in this dissertation, for its appropriateness (specialized in middle-scale molecules) and generality (not overly constrained by physical chemistry) (Flekkøy & Coveney 1999), and for its current application to related topics in OOL research (the formation of protocells [Fellermann et al. 2007, Gazzola et al. 2007]). A particular DPD system was found (Gazzola et al. 2007), and experiments from foundational ARN literature were selected for replicating in the DPD model (Bagley & Farmer 1991). The

preliminary objective of replicating DE results in DPD models was undertaken to add credibility to the primary findings about the importance of environmental dynamics, and to give these findings greater relevance in the field of ARN research. The models are explained further in the LITERATURE REVIEW and RESEARCH METHODS chapters.

Summary of Findings

The research was successful in both of its objectives. The preliminary objective was met, establishing qualitatively similar behaviors between the project's new DPD-based model and the DE-based models that dominate the literature. Using one of the foundational papers of ARN research (Bagley & Farmer 1991), the new model was able to demonstrate several of the DE model's aggregate behaviors with little tuning or optimization. Differences in the model were not surprising, and served to elucidate differences between the systems, as well as weaknesses in the measurements used for quantifying deviations from the reference non-autocatalytic steady state.

The main objective of the dissertation was successful in demonstrating sizeable differences in the emergence of ARNs when the DPD model's artificial chemistry system was subjected to dynamically structured variations in energy and mass flux. The findings include a consistent increase in ARN strength for cyclic mass flux in a wide range of parameter values, and for energy cycles with small amplitude. The gathered results also suggest interaction effects between the dynamic parameters of period and amplitude (for both energy and mass) and other parameters of the system, perhaps as interaction between the various time-based aspects of the chemistry, such as catalyst binding time and reaction rates. The results also suggest possible limitations in the metrics of

concentration profile and mass ratio as reliable indicators of ARN strength or system behavior.

These findings constitute a useful contribution by this dissertation to the research of ARNs and emergent metabolisms, as well as the encompassing areas of study, including Origin of Life and Artificial Life. The dissertation also suggests a rich future of potential research directions that are likely to produce many publication-worthy results and other useful contributions.

2. LITERATURE REVIEW AND PROBLEM CONTEXT

This project draws upon research in the fields of Artificial Life, Origin of Life, and Systems Science. As all of these are interdisciplinary fields, this project builds upon a wide variety of work. The primary phenomenon being researched here—the emergence of an autocatalytic reaction network—lives primarily in the field of ALife (Bagley & Farmer 1991). The dissipative particle dynamics model that was adapted for this research is most commonly used in chemistry and biology (Flekkøy & Coveney 1999), but the particular software arrived via an ALife project in protocell membrane formation (Gazzola et al. 2007). Models of the "hard science" DPD paradigm resemble the multi-agent models of many "soft science" fields, from sociology to economics. The inspiration for the main research question, as well as much of the anticipated future work, comes from the variety of existing OOL theories (Hazen 2005), as well as from general principles in systems science and complexity theory (Ashby 1956; Emery & Trist 1965). Systems science is critical to the core ideas of the project, including the synthetic interdisciplinary approach and the emphasis on an explicit role for the environment in a model system, and its influence runs so deep that it cannot be properly represented in references.

Origin of Life

In his book on the origin of life, Robert Hazen (2005) describes four factors that contribute to the emergence of complexity: the concentration of agents, the interconnectivity of agents, an energy flow through the system, and a periodic cycling of that energy flow. He suggests that these four factors may combine through some

undetermined function to provide an upper bound on the complexity of a system. Of

these four factors, he suggests that the role of cycling is "a frontier area of study," and

that the natural cycles of day and night, tides, and wetness may have contributed to the

emergence of life. The first three of these factors contributing to the emergence of

complexity are present in the previous research on ARNs (Bagley & Farmer 1991), but

the fourth is noticeably absent.
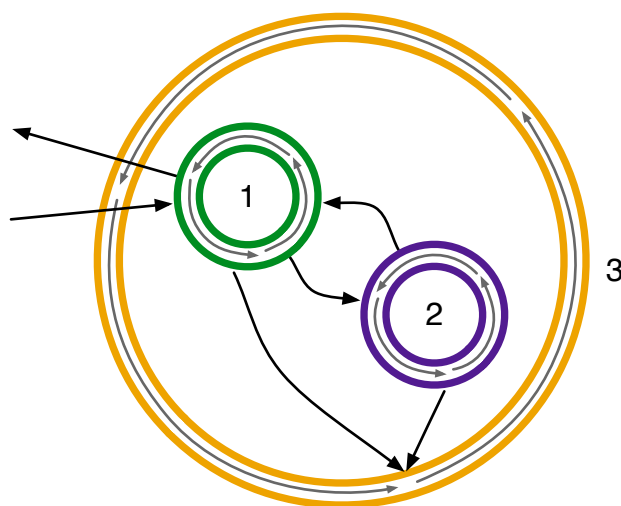
Figure 1. A simplified representation of the chemoton model.
The three autocatalytic cycles present are (1) metabolism, (2) template replication, and
(3) membrane. Food particles pass through the membrane into the metabolic cycle, and
waste particles are allowed out. The metabolism and template cycles interact to maintain
the membrane. The three cycles are shown as separate here, but may not be distinct. For
instance, in a metabolism-first model, a single cycle would fill the roles of (1) and (2).
(Adapted from Bersini [2010].)

One of the most well known abstract models for a living system is the chemoton

(Gánti 2003). According to this model, a living system must have three subsystems:

metabolism, template replication, and enclosure by a membrane (Figure 1). Each of the

subsystems in the chemoton is an autocatalytic cycle, and the three work together to

maintain the whole. Matter and energy enter the chemoton from the environment,

selectively filtered by the membrane. The metabolism uses the material from the environment as food to maintain itself, and also to produce the raw materials that the template system and membrane need to maintain themselves. The template system stores the condensed information of the chemoton, used for maintaining identity, replicating, and directing the transformation of raw materials into specific forms, such as structural and catalytic elements. The membrane maintains the concentrations of the other particles, and selectively allows unneeded particles out, just as it lets food in.
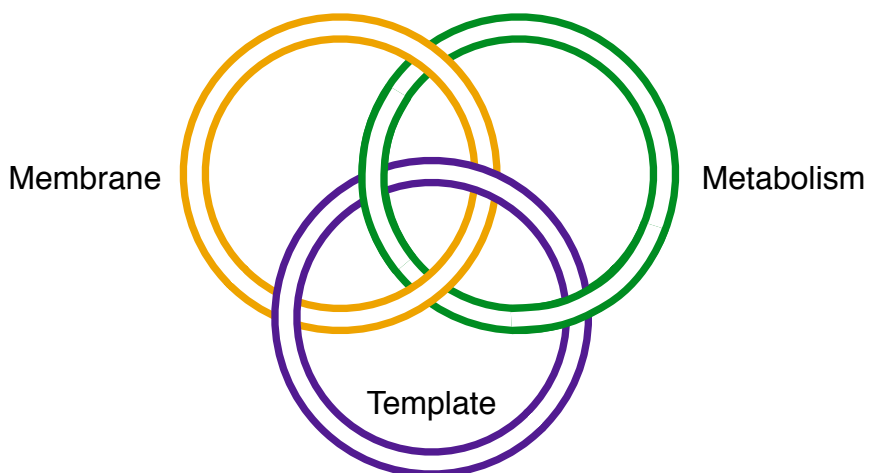
Figure 2. The three autocatalytic cycles of life, as Borromean rings.
Removing any one of the rings unlinks the other two, and the triad disbands.

Perhaps the biggest unanswered question in Origin of Life (OOL) research is: in what order did the three subsystems emerge? It appears that each needs the others, and could not exist without them. The three cycles act like a set of Borromean rings (Figure 2), where each ring holds the other two together (Cromwell et al. 1998). Remove any one ring, and the other two are separated, or, in the present context, may not be able to persist: without metabolism, the other cycles cannot maintain themselves; without

templates, the other cycles lose informational specificity; without a membrane, the other cycles cannot differentiate themselves from the environment.

Many theories have a preference for which cycle appeared first: replication-first models, such as the RNA world, favor a template molecule with catalytic functions (Higgs & Lehman 2015); membrane-first models, such the Lipid world, suggest the formation of lipid-based vesicles first (Segré et al. 2001); and metabolism-first models, such as the iron-sulfur world, begin with the emergence of a self-maintaining set of reactions (Wächtershäuser 2006). Theories for OOL can also be distinguished by the location or conditions in which they are believed to have occurred. These locations include deep-sea thermal vents (Corliss 1990; Russell & Hall 1997), tide pools, mineral surfaces (Koonin 2007; Orgel 1998), subterranean minerals (Koonin & Martin 2005; Wächtershäuser 2006), the ocean surface (Dobson et al. 2000), and more. The models of both this dissertation and the prior research it draws upon most heavily (Bagley & Farmer 1991) are perhaps most similar to metabolism-first theories. However, the ARN results may be just as applicable to metabolism emergence, regardless of order, and may also apply to autocatalytic networks of nucleotide polymers in the RNA World.

It has been suggested that augmenting Origin of Life studies, and ARN models in particular, to include spatiality may be the key to uniting the phenomena involving metabolism and membrane (McMullin 1999). Regardless of which of these came first, or which gives rise to the other, the concept of a membrane is difficult to even consider without spatiality. If one holds the metabolism-first view, then spatiality must be added to the model for metabolism to give rise to a membrane. If one holds a membrane-first view, then spatiality precedes metabolism, and is an inherent part of the system. It has

even been argued that the ALife concept of *autopoiesis* can be characterized as "collective autocatalysis plus spatial individuation" (McMullin 1999). Autopoiesis (literally, "self-creating") is a term coined to describe a general system capable of both self-making (creating the self as a bounded, individual entity) and self-maintenance (replenishing its parts by taking in matter and energy) (Varela et al. 1974). The term is meant to include Earth's living systems, with enough abstraction to include non-biological life. Reducing autopoiesis to autocatalysis and spatiality makes the point that the difference between a self-*maintaining* autocatalytic system and a self-*making* autopoietic system is intrinsically tied to spatiality. Though the current project does not attempt the unification of self-maintenance and self-making, it is a compelling direction for future work, and directly made feasible by the translation of the ARN system to the DPD model. One next step toward autopoiesis would be to integrate protocell-style, lipid-like membrane components into an ARN system model.

Artificial Life

The field of Artificial Life (ALife) has much in common with OOL research, but with more concern for the general idea of "life-as-it-might-be," rather than "life-as-we-know-it" (Langton 1992). The overlap with OOL includes the subfield of "wet" ALife, which aims to create biochemical life from scratch, in the lab. There is also overlap between OOL and software-based "soft" ALife, with its computational models aimed at simulating particular aspects of living systems, or at breaking life down to its simplest components. While these models are all abstracted to a greater or lesser extent, from the ALife point of view this is a virtue. A major theme of ALife is to discover the generalities

and underlying rules of all living systems. This variety of purpose and method in ALife research has created a broad spectrum of models, from relatively realistic simulations of molecules forming protocells, to highly abstract self-reproducing cellular automata.

Somewhere in the middle of this spectrum is the study of *autocatalytic reaction networks* (ARNs) (Bagley and Farmer 1991; Bagley et al. 1991; Farmer et al. 1986; Kauffman 1993). As described above, this work focuses on the emergence of a metabolism from a soup of prebiotic molecules, in the form of a network of autocatalytic reactions. A foundational work in this area is Bagley & Farmer (1991), which features a differential equation model with parameters of reaction rates and concentrations, with its varied configurations solved analytically. The Bagley & Farmer model (or *B-F* model, within this dissertation) attempts some amount of biochemical justification in the parameters of its artificial chemistry, and like a set of chemistry equations, it treats the "world" of its system as a formless space holding a large, well-mixed population of molecules. The B-F model has demonstrated emergence of a self-sustaining autocatalytic reaction network, given adequate levels of available energy and food concentration, as well as an artificial chemistry with catalysis, condensation-favoring reaction parameters, and so on. The B-F model does not address either the creation of a membrane, or the use or need for templates. The particle population is selected to be approximately the size of a bacterium, but there is nothing "outside" the non-spatial system to interact with or exclude. There is some discussion of the robustness of an ARN against the dramatic reduction of a member type, or against changes in the food set, but any information that facilitates this exists only in the population or the reaction graph. There is a brief mention

that the findings could also apply to RNA-like molecules, but there is no template function in the model.

Further work with this model (Bagley et al. 1991) began to explore its potential for evolution. Small changes were made by stochastically selecting polymers from "nearby" autocatalytic subgraphs, and adding them to the reaction network. These altered networks were allowed to return to a steady state, and if the added subgraph was maintained at high concentrations, it was kept. By repeating this process, the research showed how a network of reacting chemicals could "evolve," collecting variations that were favorable to system's ability to maintain its distance from equilibrium. (This is not evolution in the Darwinian sense, but is perhaps more appropriately called adaptation.) As with other studies of emergent metabolism, the researchers suggest that the "genetic information" of these networks is distributed among the polymers that make them up. (This is similar to the "composome" idea in Lipid World research [Wu & Higgs 2008], though the Bagley et al. [1991] model does not use that term. This term will be revisited later.) The discussion accompanying Bagley et al. (1991) suggests the possibility of modeling multiple reaction networks, each in its own container with occasional interactions between them. However, it does not appear that this work has ever been published.

For the most part, the study of autocatalytic reaction networks has been connected with the metabolism-first theories for the OOL. This class of models has been criticized for its lack of biochemical realism, because the entities in these models have limited potential for heredity (Anet 2004; Orgel 2008; Szathmáry 2006; Vasas 2010). However, findings in autocatalytic research may also be applicable elsewhere, despite their current

association with metabolism models. As the chemoton model in Figure 1 demonstrated, all three loops could be characterized as autocatalytic. Bagley and Farmer (1991), though they labeled their networks as metabolisms, were aware that the model's components could have been nucleic acids such as RNA rather than amino acids and proteins.

Even if the metabolism-first concept falls out of favor within OOL research, it may yet have other applications. Even if metabolisms did not emerge first, they did emerge from systems that did not possess metabolisms. Indeed, the replication-first OOL models could perhaps be faulted for their silence on how metabolism emerges. Because the transition to metabolism happened and was necessary for life as we know it, the study of the emergence of metabolism will continue to be relevant to OOL studies.

And, within the field of ALife, where Earth's actual history and biology carry less weight, the emergence of metabolism-first life continues to be a possibility that is interesting to explore. For these reasons and more, the study of autocatalytic reaction networks remains active and relevant (Benkö et al. 2009; Gordon-Smith 2009; Hordijk et al. 2014; Stassinopoulos et al. 1998; Wu & Higgs 2012).

Autocatalytic Reaction Networks

Briefly defined, an *autocatalytic reaction network* (ARN) is a set of molecules where every member is a participant (reactant or product) in a catalyzed reaction that involves only other members of the set as participants and catalysts. This interdependence can create a tightly bound group of reinforcing feedback loops whose reactions are able to dominate non-catalyzed reactions, in terms of activity or member populations. This is most relevant in a system that is being driven away from equilibrium:

since catalysis affects reaction rates rather than equilibrium levels, an ARN could not differentiate itself in a system at chemical equilibrium.

In ARN research, the default is for reactions to be reversible, so that molecules on either side of an equation's directional arrow are treated equivalently in the ARN definition. This leads to the interesting phenomenon that an ARN may possess a basic facility for self-maintenance: if a single molecule type is removed from a network, the stock of other molecules and their supporting reactions will replenish some of the lost type. Because of these behaviors—the ability to "out-compete" other reactions for food, and the simple self-maintenance—the emergence of an ARN has been described as the emergence of a simple metabolism (Bagley & Farmer 1991). Arguably, this simple "proto-metabolism" is only fulfilling the *constructive* aspect of a metabolism, as it builds structure and maintains itself from component parts, but does not demonstrate the *destructive* metabolism of organic life, where long molecules are broken down for energy and waste is produced.

The reactions and molecule types of an ARN can be specified as a list of chemical equations or as a network graph. This specification separates the potential for reaction from the "actual" reactions that might happen in a simulation with energy, reaction rates, and molecule concentrations. In the literature, this aspect of an ARN is sometimes called the *autocatalytic set*, and some research focuses on the characteristics of these sets that make an ARN more or less likely to emerge (Hordijk & Steel 2004; Hordijk et al. 2014). Figure 3 depicts the graph of an autocatalytic set that is used extensively in Bagley & Farmer (1991). Since some of the results from Bagley & Farmer (1991) are targets of attempted replication in this dissertation, the graph is also the basis for most of the ARNs

and results presented here. This graph contains 13 reactions, each catalyzed by one member of the set. (Each reaction must have at least one catalyst to be in the ARN, but there is no restriction against multiple catalysts.) There are 2 food set molecules (the monomers *A* and *B*) and 13 polymer products of various lengths. Some of the reactions are catalyzed by their own products and so qualify as single-reaction ARNs, nested within the larger ARN. There may be other subset ARNs hidden in the graph as well.

Figure 3. The autocatalytic network of Bagley and Farmer (1991).
This network is used in many of the experiments in this dissertation, for purposes of replicating results in the literature, and for new research. Bullets are reactions, and arrows connect molecules to them, showing favored direction. Dashed lines show catalytic effects, and the circled particles are the food set. There are 15 molecule types and 13 catalyzed reactions in this ARN. (Reprinted from Figure 2 of Bagley and Farmer [1991].)

A simpler example of an autocatalytic set can be found in the reaction network graph of Figure 4. The autocatalytic set consists of the molecular species *A, B, BA, BAA,*

and *BAAA* (marked by ovals in the figure). By the definition above, an autocatalytic set

must be all of the reacting molecules for a particular set of catalyzed reactions, for which

all of the catalysts are also members of the set. "Reacting molecules" refers to all of the

types on both sides of an equation, regardless of favored direction, but not the catalyst.

The five types in the set meet these criteria: each is a reactant in a catalyzed reaction

where all components are also in the set. For instance, *BAA* is a product of the reaction

BAA+A⇨BAAA that is catalyzed by *BA*, and the other three participants (*A, BA, BAAA*)

are also members of the set. On the other hand, *BB* and *AA* are not members of the set, as

they don't meet the criteria: *BB* is a catalyst, but is not a reactant in a catalyzed reaction;

AA is reactant in a catalyzed reaction, but the catalyst (*BB*) is not in the set.
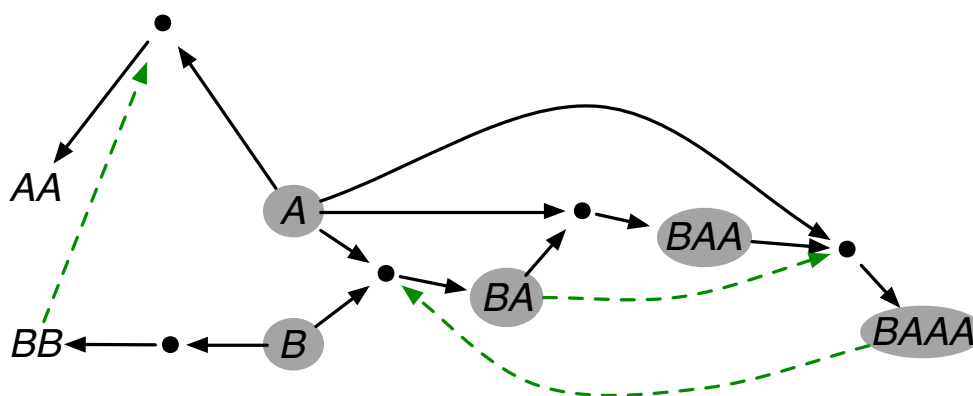
Figure 4. A simple reaction network graph containing a small ARN.
The graph shows only a subset of all possible combinations of A and B. Dots represent
reactions, and solid lines connect them to reactants, with arrows showing favored
direction. Dashed lines are catalytic effects. Molecule species marked with ovals
comprise an ARN, or autocatalytic set.

Thankfully, there are methods to algorithmically locate or verify autocatalytic sets

within a reaction network. For instance, to decide on the presence of an autocatalytic set

in a RN, a list may be created of all molecules from catalyzed reactions. A minimum

concentration or population may be used as a threshold for inclusion. Any particles that were not catalyzed by members of the list are removed. This process is repeated until all molecules on the list are catalyzed by other molecules on the list. If the resulting list is non-empty, it contains at least one autocatalytic set. More efficient algorithms exist as well (Hordijk & Steel 2004).

The presence of an autocatalytic set is necessary but not sufficient for the emergence of an ARN. As Bagley & Farmer (1991, p.106) states, "We must emphasize, however, that the presence of an autocatalytic set in a reaction network, in and of itself, does not imply that there will be any interesting departures from equilibrium behavior. To achieve catalytic focusing it is critical that the kinetic parameters are favorable." In this context, *focusing* refers to the effect where the mass of the system exists mostly in the relatively few molecule types that make up the autocatalytic set. Because the system is driven away from equilibrium, catalysis can cause this focusing effect, and maintain a higher fraction of the system's mass in catalyzed products than would otherwise be possible. Bagley & Farmer (1991) suggests three "favorable" kinetic parameters that might be critical: strong catalytic effects; reaction energies that favor polymerization (increasing molecule length); and a driving force (such as a constant inflow of new particles). The claim for the necessity of these three parameters was not made strongly, though it was consistent with the results (Bagley & Farmer 1991). In the present research, a system configuration was encountered where it appears that the driving force is not necessary. This could not be proven within the scope of this dissertation, however. (See Result 1.3, especially Figure 11, Table 5, and the intervening discussion.)

While studying a particular ARN, it can be helpful to describe and measure certain other sets of molecules in the system. In the literature, the *background set* is the name for the set of all molecule types that are not members of the ARN. The chemistry of a typical model (such as the B-F model) allows for all possible reactions to occur, though most are not catalyzed. A chemical model with 2 monomer types has 2^n polymer types of length n, and $(2^{n+1} - 2)$ total types of length $\leq n$. For the B-F model's network (Figure 3), there are 15 types in the ARN's set, and 32,751 ($2^{15} - 2 - 15$) in the background set. (That count is only up to length 14 – background reactions aren't limited in length.) Without a process that selects for some reactions to occur more than others, specific types become exponentially more rare with length.

It is also useful to split the ARN itself into two sets: the *food set* and the ARN without the food set. (The latter has no convenient short name.) The food set contains all of the molecule types that are allowed by the model's design to enter the system as *mass flux*. This mass flux is the driving force that pushes the system away from equilibrium, and that driving is what allows the ARN to focus mass. (A formal description of the relation between food set and catalytic reaction network is given in Hordijk & Steel [2004], along with an algorithm for detecting autocatalytic sets.) Separating the food set from the rest of the ARN makes it easier to study the effects of changes to the mass flux. For instance, in Figure 12 and Figure 13, the distribution of mass as a function of mass flux reveals relationships between food set, ARN without food, and background set, which could be obscured if the ARN were not separated in this way.

There are no definitive measures or boundaries to decide whether the potential in an autocatalytic set has been fulfilled in a model and the ARN has "emerged." The

literature states that an emerged ARN is "an autocatalytic set whose concentrations make significant departures from the values they would have if none of the reactions were catalyzed," and at the same time admits that this definition is "vague" and "subjective" (Bagley & Farmer 1991, p.107). One attempt to measure emergence is the "concentration profile," which condenses the populations of a set of molecules (ARN or background) down to a handful of points, with molecule types grouped by polymer length. An example is shown in Figure 5, using data from a simulation of the DPD model.
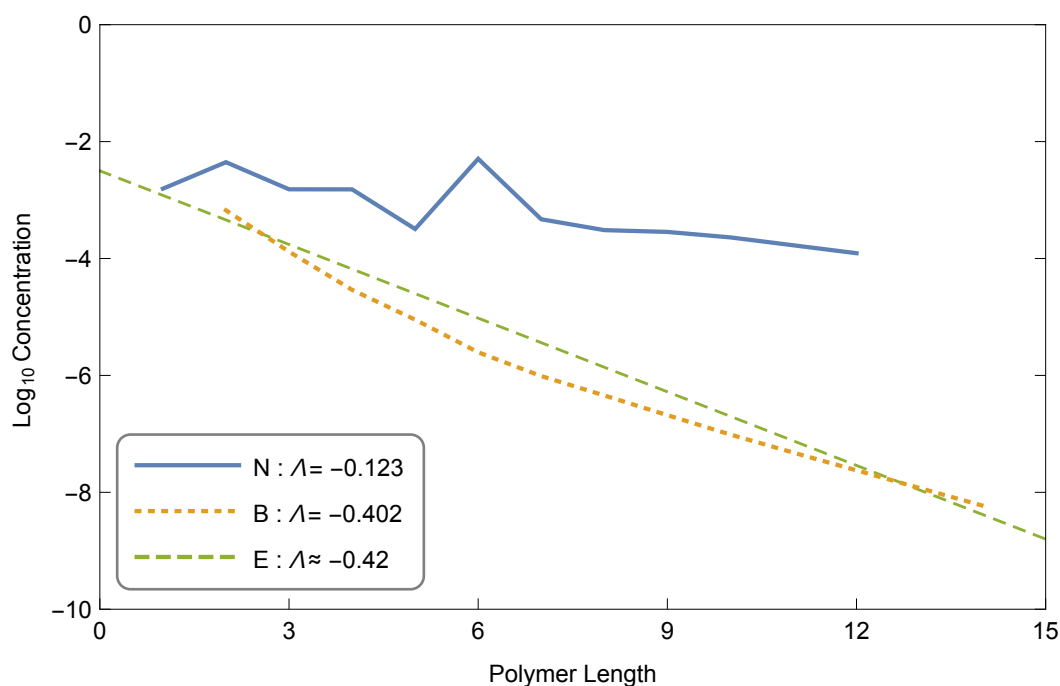
Figure 5. Concentration profile of a single simulation run with the DPD model.
The lines N and B show mean concentrations for ARN (Network) and Background sets, respectively, grouped by polymer length. The line E shows approximate concentrations in a non-catalyzed network at steady state. The legend includes Λ values for each line, which are the slopes of best-fit lines for each set. (This figure is identical to Figure 9, and is explained further there, along with a similar B-F figure.)

In Figure 5, the concentration profile for the ARN (N, for Network) is plotted on the same axes as the concentration profile for the background (B). There is also a third

line that represents a similar system without catalysis (marked E for Equilibrium in the literature, though "non-catalyzed steady state" may be a more accurate label). When an ARN has emerged, the difference in concentrations will be substantial, and the N-points will be far from the B-points. This may be especially apparent for longer polymers, since concentrations should drop off exponentially with polymer length. Note that the vertical axis uses a logarithmic scale, and the concentrations of the longest polymers in the ARN are a few orders of magnitude more common than the background in this figure.

The slope of a concentration profile line can also be used, further condensing the populations of a set of molecules down to a single value. For the ARN, this acts as an approximation for the strength of the network, summarizing its ability to maintain its members' concentrations at levels higher than the background. This measure is labeled "Λ" (capital Greek lambda) in the reference literature (Bagley & Farmer 1991; Bagley et al. 1991). To computer the metric Λ, a line is fit to a concentration profile (ARN or Background) by a least-mean-squares method, and the slope of that line is taken to be Λ. The values for Λ are always negative, since concentrations are measured as fractions of the total molecule count and then plotted in the profile on a logarithmic scale. Λ values can be compared between simulations with different parameters, with larger (closer to zero) values suggesting higher concentrations. For ARNs, higher values imply greater separation from the background and from the steady-state value, which represents the concentrations of a non-catalyzed system.

As a simple, single-value indicator for an entire population, the concentration slope Λ is convenient for characterizing an ARN-capable system's sensitivity to

variations in other parameters, such as mass influx rate, reaction rate, and probability of catalysis. This method of experimentation is prominent in the reference literature (Bagley & Farmer 1991), and because of that, it also figures prominently in this dissertation's results. See the section COMPARING THE MODELS in the RESULTS chapter for specific examples of the measure Λ being used for sensitivity analysis.

Modeling ARNs

The differential equation (DE) modeling paradigm is the typical structure for ARN simulation, including the B-F model, as mentioned previously. A DE model describes systems through a series of differential equations, and the behavior is characterized by rates of change and flows between populations. A model in the DE paradigm typically assumes that there are no spatial or distance effects; that members of populations are identical; and that fractional population quantities are acceptable. When used for molecular simulation, a DE model might be thought of as simulating a large population (on the order of 10^6 or more particles) that exists in a well-mixed container, such that every particle has equal opportunity to interact with every other particle. These models are typically deterministic, though random factors can be included (see OTHER RELATED MODELS, below). The DE paradigm is essentially identical to the system dynamics (SD) paradigm (Scholl 2001), though one term may be more common than the other in a particular domain.)

The DPD modeling paradigm that is used in this dissertation can be considered a sub-type of the general category of Agent-Based Modeling (ABM). In the ABM paradigm, systems are modeled as a population of discrete, individually distinguishable

agents with well-defined states (Epstein 1999). These agents typically inhabit a simulated

physical space in 2D or 3D, moving about and interacting in discrete time. The rules for

interaction may be the same for all agents, or they may vary for different types of agents.

Models in the ABM paradigm prefer smaller populations than can be modeled using DE,

and typically require more computational resources relative to population. ABM models

emphasize local effects and interactions, spatial variations, and sensitivity to random

initial conditions. Because of random effects, experiments with ABM models often

require multiple executions, the results of which are collected are described via statistical

distributions.

Models in the DPD paradigm specialize in simulating systems of chemical

particles with more realism than a general ABM model. DPD modeling is most suitable

for modeling *mesoscopic* systems (between microscopic and macroscopic) comprised of

larger molecules or clusters of molecules, and the time- and space-scales over which such

particles interact (Flekkøy & Coveney 1999). The related paradigm of molecular

dynamics (MD) is designed for modeling much smaller systems and effects, close to the

atomic scale, and is therefore not as useful for modeling at the scale of an ARN (Flekkøy

& Coveney 1999). DPD models are also well suited to the study of hydrodynamic

systems (like an ARN) and have been used to model the behaviors of lipids and lipid-like

particles, including the formation of vesicles and protocells from these particles

(Fellermann et al. 2007; Gazzola et al. 2007; Solé 2009).

The specific DPD modeling software used in this dissertation is derived from a

package that has been used in other ALife research (Gazzola et al. 2007). The source

code for the software was generously made available, so that it could be modified as

necessary for the ARN simulations. It had previously been used for several different projects, most recently to study catalytic activity between lipid-like amphiphiles (molecules with both hydrophobic and hydrophilic parts) and smaller molecules, with some realistic aspects of chemistry included (Gazzola et al. 2007).

Other Related Models

Many other models used in ALife or OOL research have some relation to those used in ARN research, and so may find some relevance in the results of the dissertation. For instance, some work on the 'RNA World' concept has been influenced by the idea of a 'mineral honeycomb' whose inorganic compartments provide protection and spatial separation for replicators (Branciamore et al. 2009). That research used a cellular automaton model and differed greatly from the models of this project, and yet it is similar in the aspects of spatiality and an externally provided compartment. The use of the mineral honeycomb was influenced by OOL ideas involving deep-sea thermal vents and mineral surfaces (Koonin & Martin 2005; Wächtershäuser 2003), and this dissertation counts those same ideas among its influences.

OOL research has also made use of spatial molecular dynamics models, especially in the study of protocells (Boerlijst & Hogeweg 1991; McCaskill et al. 2007; Morowitz et al. 1988; Segré & Lancet 2000). This research seeks to produce a minimal cell, either in simulation or *in vitro*. Much of the work is concerned with membrane formation, which is inherently spatial in its goal of creating a boundary that separates the interior of a cell from the rest of the world. Some protocell research uses the DPD modeling paradigm (Fellermann et al. 2007; Solé 2009), and the DPD models of this dissertation were

adapted from source code used in protocell research (Gazzola et al. 2007). This presents

the possibility that the DPD paradigm could help to bring together protocell boundary

research and the proto-metabolism of ARN research. Another potential bridge between

protocell research and ARNs may exist in a greater role for peptides in protocell

formation, on par with lipids (Fishkis 2007). Since peptides are more closely related to

the amino-acid-like chemistry typical in an ARN, a protocell theory that includes

peptides might be even more conducive to studying the linked emergence of membranes

and metabolism.

Another ALife model that is related to this project is the GARD model (Segré et

al. 1998, 2001; Segré & Lancet 1999). GARD (for "graded autocatalytic replication

domain") is similar to the B-F model in several ways. It features an artificial chemistry

with autocatalytic features, with reactions determined by differential equations or

stochastic methods. However, the GARD model is more concerned with amphiphiles

(molecules that both attract and repel water) and origin theories of The Lipid World.

Much of the research with the GARD model relies upon the possibility of heredity that

occurs via a composite genome, or *composome*. The idea of the composome is that

templates in early life may have been characteristics of populations of molecules, rather

than a single molecule like RNA or DNA (Wu & Higgs 2008). Though the GARD model

is focused on membranes, the composome aspect is relevant to this research for its

treatment of replication without a template. Some research using the GARD model has

even considered the role of the environment, though in a different form than in this

project (Shenhav et al. 2007). In that research, composomes were measured in models

featuring environments of either infinite or finite size. In the finite model, variety in

composomes was found to correlate over time with aspects of the environment, while composome variety in the infinite model did not.

Another common molecular modeling method is to use stochastic methods to simulate interactions. These models are referred to as *stochastic molecular* models or *stochastic molecular collision* models (Bagley et al. 1991; Farmer et al. 1986; Fernando et al. 2007; Segbroeck 2009). This paradigm is similar to the DE paradigm in that it models systems with no spatial aspect. However, it typically assumes smaller, whole number populations like an ABM paradigm model, which allows individual particles to be tracked. Interactions between particles are modeled by selecting candidates at random from the population, completing a reaction if conditions are right, then returning the results to the population. By repeating these steps many times for a large enough population, the results approximate those of a DE model with the potential for more realistic treatment of individual particles. A drawback of this discrete treatment of interactions is that parallel events become more difficult to model. Another drawback is that a stochastic molecular model has higher computational resource requirements than a DE model, and so has lower practical limits on system size. In this sense, it likely falls somewhere between the DE and ABM modeling paradigms, trading speed for the additional modeling capabilities.

The subject of the stochastic molecular model in Fernando et al. (2007) is notable here for its other similarities to the present research. That work attempts to answer questions about the competing processes of elongation and replication in RNA, related to the RNA World concept. One issue in this area is to determine under what conditions replication may be favored over elongation, since this is required for RNA to serve the

template function in a prebiotic system. Fernando et al. (2007) specifically considers whether the effects of tidal cycles favor replication, as has been suggested elsewhere. By applying the stochastic molecular model (in some ways a compromise between the DE and ABM paradigms, as described above) to a problem featuring polymer formation and a cyclic environment, Fernando et al. (2007) is among the closest works to the present research. Though the differences are still sizeable, the similarities may prove useful if the two areas of research can be bridged, or can "cross-pollinate" each other. In particular, the dissertation attempts substantially less chemical realism than Fernando et al. (2007), and so might be able to learn from that line of research should such validity be sought in the future.

The contrast between the ABM and DE paradigms has been made repeatedly, in areas of research from molecular dynamics to sociology and economics (Borshchev & Filippov 2004; Epstein 1999; Schieritz & Milling 2003; Scholl 2001). These comparisons often highlight how the two model types complement each other, and how their different strengths can be applied to the same problem for different research questions. The strengths mentioned typically are spatiality, individualism and emergence in ABM models, and feedback, causality and aggregation in DE models.

Systems Science and Complexity

The influence of Systems Science ideas exists throughout this research. The most obvious is the focus on the environment as a rich part of the model, rather than an implicit background for the system of interest, present only as simple parameters. In the models of this dissertation, the environment is considered as a source of *causal texture*.

This means that structured variation in the environment can cause change and create structure in the modeled system. In the work of Emery and Trist (1965), causal texture is classified into four types:

1. unchanging, random, and inhabited by one system
2. unchanging, clustered, and inhabited by one system
3. inhabited by competing dynamic systems
4. a dynamically changing environment, perhaps containing other systems, and perhaps structured at a higher level

The previous models for ARNs would be classified as the first two types. The models in this dissertation were designed to have the fourth type of causal texture. This level of dynamism, where the world that the system inhabits undergoes structured change, can be a source of variety and order for the system. It is also a more reasonable approximation of the environments in which life is observed and is expected to exist. This idea of causal texture has been applied to the evolution of adaptive systems (Fletcher et al. 1997, 1998), where it was found that the structure of the environment had a strong effect on the ability of mobile agents to adapt. Though the structure of the environment was held constant for models in that work, agents were forced to move around, so that they experienced a dynamic structure. The information content of the environment, relative to agents' abilities, was predictive of their adaptive success.

This project was also inspired by Ashby's Law of Requisite Variety (Ashby 1956). This law says, roughly, that for a system to be able to regulate itself optimally, it must be able to respond (deterministically) to external disturbances with a variety of action that meets or exceeds the variety of disturbances. This suggests that a dynamically structured environment can be viewed as a source of selective pressure for adaptive systems. Any systems that lack the requisite variety to respond will have much less

chance of regulating themselves, and will therefore be less fit. This also suggests a drawback of modeling with a simple environment: the absence of selective pressure leads to less-fit systems. For instance, in a simple environment that provides constant abundant resources, any system may survive regardless of its fitness. In an environment that is constrained but still simple, the selective pressure is also simple. This may lead to a monoculture, with a single system that has the best response to the simple constraint. Such a system is unlikely to be adaptive, and probably could not self-regulate outside of the single, simple environment in which it emerged. Modeling this sort of system-environment interaction might be providing insights that are only applicable in simple worlds, while the worlds that we inhabit are decidedly complex.

3. RESEARCH METHODS

This chapter contains three sections. The first, Model Adaptation and Analysis, describes the process of adapting DPD modeling software to the purposes of this project, as well as the support and analysis tools that were created and used for the dissertation. The second section, Replication and Comparison Experiments, describes the methods used to complete the four experiments aimed at replicating results from the reference work in Bagley & Farmer (1991). The third section, Dynamic Environment Experiments, explains the methods used to answer the main research question through a series of four experiments. These experiments applied dynamic effects to the reaction network system in the form of cyclic mass flux or energy, varying in period and amplitude.

Model Adaptation and Analysis

The DPD models used in this dissertation were executed in software that was customized and adapted specifically for this project, beginning with a body of code that had been used for other DPD models in ALife (Gazzola 2007). The core of the existing code that handles particle movements and interactions was usable in its received state. Modifications to the software were necessary to improve the mechanism of catalysis, to implement the dynamic environmental interactions, and to enable data logging needed for the analysis. Other modifications were made to improve automation and performance.

The software is written in the C programming language, and was compiled by the open-source compiler that is available for free with the Mac OS X operating system. (The software should also compile and run on a Linux-based operating system without substantial modification.) To run a single simulation, the compiled software is executed

from the command line with a configuration file that contains all desired parameters, both for the workings of the DPD paradigm and for model system. Output from the simulation is written as text to the console and saved into text files, and also displayed graphically in a separate window. The graphical output is designed to show a 2-dimensional view of the system's particles and bonds, updated at a configurable frequency, though this may be limited by system resources. Among the modifications added for the dissertation is an option to disable the graphics within a simulation, which improves performance and simplifies execution during parallelization.

The DPD software, before modifications, allowed for a form of catalysis (Equation 1) where one particle type (*X*) would be transformed to second type (*Y*) by reacting with a catalyst (*E*), with the catalyst treated as both reactant and product. However, this method of modeling catalysis was not feasible for the polymerization reactions of the B-F model or other ARN research. In these models, there may be exponentially many possible chains built from a few monomer types, which would be difficult to represent in the older catalyst system. That system would also make it difficult to track and conserve the monomers throughout a simulation. An example of the desired form of catalyzed reaction is presented in Equation 2, in which two chains of monomers (*AABA* and *BA*) form a bond, using the catalytic path provided by a third chain (*BBB*).

Equation 1. Transformative catalysis in the original DPD software. *X* and *Y* are molecules, and *E* is the catalyst.

$$X + E \rightarrow Y + E$$
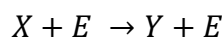
Equation 2. Catalyzed reaction between polymers. The polymer chains at left bond into the long chain at right by energetic preference. *BBB* is the catalyst.

$$AABA + BA \xrightarrow{BBB} AABABA$$
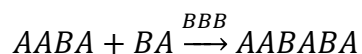
To implement this in the DPD software, additional logic was added to the functions that determine whether two particles might bond (if separate) or cleave (if already bonded). The existing algorithm for this determination was based on factors such as distance, forces between the particles, and configurable reaction probabilities. The new logic adds a check for whether the reaction has potential catalysts, and if so, whether any are nearby and available for reaction. If all conditions are met, the catalytic efficiency of that catalyst is factored into the reaction probability. The probability is tested against a generated random number, and if the reaction occurs, the catalyst is marked as 'bound' along with the reactants. This lasts for a small, parameterized amount of time, during which the participants are bound together and unable to engage in other reactions. Besides the logic added to include catalysis in reactions, options were added to the input file format for specifying catalyst properties, and data structures were built into the software to efficiently maintain the list of possible catalyzed reactions.

The other major modification to the software added the functionality of dynamic values for mass flux and free energy. For both of these dynamic interactions, the configuration options allow for choosing a cycle type and specifying its parameters. The only cycle type that is currently implemented is a sine wave with parameters for period, amplitude, and phase offset. (The code is structured so that other functions can be added easily.) Both sorts of dynamics build upon existing model parameters, which provide static values that are combined with results of the cycle calculation. The cycle parameters are used to calculate a time-dependent amount to be combined with the static value.

For mass flux, the previously existing parameter is a base rate that specifies the inflow to the system in mass per time, in the form of molecules drawn randomly from the

food set. By design, any inflow of mass is balanced by an equal outflow, taken at random from the system. In the DPD model, the location of inflow and outflow could be relevant to results, but this was controlled for in the current experiments by randomizing inflow location and outflow selection. The conservation method for the DPD model is analogous to the system "spilling over." Polymers are removed while total mass exceeds the desired set point. If the outflow of a long chain overshoots the set point, no additional outflow occurs until inflow "refills" the system. The matched inflow and outflow effectively conserve mass in the system, which is equivalent in these models to maintaining monomer population, since all monomers have equal mass. These built-in assumptions serve to simplify the experimental design by controlling for those properties.

The addition of the cyclic inflow to mass flux was relatively straightforward to implement in the code, since the two values have the same units and are combined additively, maintaining the inflow/outflow balancing process and all other code for handling particle population. The implementation of cyclic energy was not as straightforward, in conception or in implementation.

In standard chemistry, the most direct way to manipulate a system's energy is arguably through temperature, and this is the form of dynamic energy that was selected for the model. In hydrodynamic systems such as these, an increase in temperature generally causes an increase in reaction rate, and a decrease in temperature causes a decrease in rate. The quantification of the rate change can be approximated with the Arrhenius equation (Equation 3) as a function of the reaction's rate constant and activation energy, and the system's temperature (Housecroft & Constable 2010). The

Arrhenius equation provides a guideline that reaction rate doubles for each 10 K rise in

temperature, approximately accurate within a limited temperature range.

Equation 3. The Arrhenius equation. k=reaction constant, A=frequency factor, E_a=activation energy, R=molar gas constant, T=temperature.

$$k = Ae^{-E_a/RT}$$

Equation 4. Cumulative distribution function for the exponential distribution. The parameter in the equation, λ, is a rate parameter.

$$F(a) = 1 - e^{-\lambda a}, \text{ for } a \geq 0$$

Equation 5. "Pseudo-code" equation for probability of reaction in the DPD model. Configuration parameters for *reaction constant* and *catalytic efficiency* (see **Table 3**) have been converted to factors. The equation, similar to Equations 3 and 4, combines the factors with a *dynamic energy* factor to compute an adjusted reaction probability.

$$\text{Prob(reaction)} = 1 - e^{-(\text{reaction rate})(\text{energy factor})(\text{catalysis factor})}$$

In the DPD model, reactions are parameterized with probabilities rather than rate

constants, and the doubling relationship is not directly applicable. However, this gap can

be bridged using the relationship between an event's probability and the chance of it

occurring in a given time. The constant-rate aspect of the event can be modeled as a

Poisson process, and the probability of such an event is then described by an exponential

distribution function (Equation 4) (Dekking et al. 2005). This function is notably similar

to the Arrhenius equation, both featuring the negative event rate in the exponent. The

guideline described above—a doubling of rate for a linear increase in temperature—is in

fact a property of the exponential component in these equations. These relationships

suggested the method for implementing dynamic energy as temperature change in the

DPD model. The implementation is designed so that a parameterized energy cycle

function acts as a model for variation in the temperature of the system. The value from

the cycle function is dynamically converted into a factor that can be combined with

factors for the base reaction rate and for any catalytic effect (Equation 5). This equation is used to calculate the probability of a given reaction occurring within an integration step

The software was also modified to collect and output additional data during a simulation, to monitor the new additions described above and also to serve other requirements of the dissertation's experiments. The output of reaction counts was modified to include catalysts when present. The census and reaction outputs were both modified to make it easier to read and process their data, over many simulations and experiments. To work with the new options of catalysis and dynamics, the logged control file and the onscreen run-time output were also modified.

Throughout the process of adapting the DPD modeling software to the needs of the dissertation, there were challenges in understanding the history and purpose of the code. The code had been created and modified for multiple purposes by different authors, and it was at times necessary to trace and disentangle that history to understand the existing code before making changes. On the other side of this issue, it remains to be determined whether the modifications for this dissertation should be integrated back into the code that other authors may be using still. Nor is it clear if the source code can be shared publicly. That decision may depend on the licensing and ownership claims of the various contributors and their institutions. The preferred solution would allow for the sharing of the source code along with the data and configuration files for the simulations, so that the findings have the potential to be confirmed and extended by the community.

As designed, this dissertation called for a large number of runs with the DPD software, due to the many separate sets of parameters and also to repeated runs of each parameter set. Many of the different parameter sets are detailed in the sections below, and

many others were run while testing code modifications or building understanding of the models. The repeated runs of each parameter set are needed to determine a distribution of behavior, due to the random effects in the model. As described in the LITERATURE REVIEW, models from the DPD paradigm (and more general ABM paradigm) are typically not deterministic, so that data must be gathered from many runs to gain confidence in the statistical findings. For the majority of simulation configurations in the dissertation, 10-20 runs were completed and analyzed together. This provided enough statistical strength (for example, in the form of standard errors, as in Figure 22) to meet the objectives of the research by identifying trends and demonstrating proofs-of-existence for behaviors of interest. If the goal had been to focus on a specific quantification of the results, additional runs may have been required to achiever greater understanding and confidence.

The large number of model runs created a need for automated systems to monitor those runs, organize their outputs, and then gather and analyze the data. This was accomplished using several different tools. Short 'script' programs were created with the python programming language and with the command-line shell, both included with the Mac OS X operating system. Some of the scripts simplified file-handling tasks such as synchronization, archiving, and duplicating. Other scripts were created and used to run the simulations in parallel on a workstation, by managing queues and maintaining a desired (high) level of processor utilization. After sets of simulations were completed and their data files gathered together, most of the analysis was completed in Mathematica (Wolfram Research 2015). Mathematica code was created to parse the various log files, extracting data as single events or as time series, for single simulations or for repeated

sets, and across multiple variations of model parameters. For instance, when performing a

sensitivity analysis, a basic control file would be cloned and modified into multiple

variations for each parameter value (also in Mathematica). The control files were then

entered into the queue on the parallel workstation, executed by scripts as processor time

became available, and then transferred to an analysis directory upon completion. The

analysis code then could extract select time series of data from each repetition of each

parameter variation. This data might consist of simulation outputs from 200 time points

for each of 20 repetitions and 10 parameter values. The outputs at a time point include a

count of all polymers in the system, a summary of all of the reactions that occurred since

the previous output, and a few summary values. Approximately 20,000 simulations were

run in total, producing more than 4GB of data. The majority of the simulations took place

in the testing, optimization, and exploration phases of the research. These were often

shorter runs, taking less than 30 minutes of processor time each. The results presented in

the dissertation represent around 2,000 simulations, with most taking 45-60 minutes of

processor time on a workstation. Though these only represent around 10% of the

simulations, they account for about 25% of the data, due to longer runs and more frequent

data recording.

Replication and Comparison Experiments

The preliminary dissertation objective was to establish grounds for comparison

between the DPD model that would be employed for the main objective and the DE

models more common in the literature. One of the foundational publications in the study

of ARNs (Bagley & Farmer 1991) was selected as the target for this comparison, and it

was then necessary to select a few results for replication from the many results presented

in that work. Four results were eventually selected, which are summarized in Table 1. All

four results are presented graphically and were judged to be amenable to qualitative

description and comparison. This was necessary, since a quantitative calibration and

validation between the DPD and DE models would be beyond the scope of this

dissertation. Three of the four results use the 13-reaction network described in Figure 3,

which was the only reaction network fully specified for the B-F models, and one of the

few points of complete translation between the models. The reactions that make up this

autocatalytic set are also listed completely in Table 2.

Table 1. Summary of replicated experiments.

Experiment	Result title	B-F Result from DE model	Dissertation Result from DPD model
1.1	Polymer concentrations over time	Figure 6	Figure 7
1.2	Concentration profile of a single simulation run	Figure 8	Figure 9
1.3	Sensitivity of concentration slope Λ to varying mass flux	Figure 10	Figure 11
1.4	Sensitivity of mass distribution to varying mass flux	Figure 12	Figure 13

The first of the four selected results demonstrates how the concentrations of

particle in a system with an ARN approach a steady state over time. This result was

selected because it clearly presents behavior typical of a DE model, but it is also the type

of behavior that should be relatively easy to replicate in the DPD model, at least

qualitatively. Differences in these results would likely be instructive about how the two

model paradigms differ, and how they compare favorably. This was the only result

selected that did not use the 13-reaction network of Figure 3. The content of the B-F

figure was not completely described, nor was the reaction network that was used, so this replication could not aim for a highly detailed comparison.

The second result selected for replication is a plot of the concentration profile for a single simulation. This result was selected because it explains the derivation of the frequently used ARN metric Λ, or concentration slope, the details of which can be difficult to understand. This result is also heavily dependent on the reaction network graph that is used, which was perhaps the most translatable feature from the DE model to the DPD model. This result specified its parameters clearly, and although they can't be translated directly to the DPD model, their relation to other parameters given in Bagley & Farmer (1991) provided confidence that the figure represents typical results, which the DPD model would likely be able to approximate. (See Table 3 below for discussion of the difficulty in translating system parameters.)

The remaining two results show sensitivity analyses for the concentration slope Λ and for the distribution of mass, each as a function of mass flux (the flow of particles through the system: food set enters, arbitrary molecules leave). These analyses were selected for their concise and informative results, showing qualitatively distinct responses to mass flux. The range of mass flux values common to both figures includes rates too low to drive the system and too high for any reaction network to persist. The transitions from those extremes to the central 'sweet spot' where the ARN can differentiate itself are deeply informative, both about ARN behavior and about the different ways to measure it.

Table 2. The 13-reaction autocatalytic network used in all dissertation experiments.
The network is adapted from Bagley & Farmer (1991), as in **Figure 3** and the simulations
resulting in **Figure 8**, **Figure 10**, and **Figure 12**. Monomers *A* and *B* are this network's
food set types, which 'feed' the system via mass flux. The network's *background set*
contains all other polymer permutations, since all polymerization reactions are possible,
though the vast majority are not catalyzed. All reactions are reversible.

Cleavage Product 1	Cleavage Product 2	Bonding Product	Catalyst
BAABA	*BAABAAB*	*BAABABAABAAB*	*ABBAAB*
BAAB	*AAB*	*BAABAAB*	*ABBAAB*
BAAB	*A*	*BAABA*	*AAB*
BAABAAB	*B*	*BAABAABB*	*BAABAAB*
B	*AAB*	*BAAB*	*AB*
AAB	*AB*	*AABAB*	*BAABAAB*
A	*AB*	*AAB*	*AAB*
BAAB	*B*	*BAABB*	*BAABA*
A	*B*	*AB*	*B*
A	*AABABBAAB*	*AAABABBAAB*	*AB*
AB	*BAAB*	*ABBAAB*	*ABBAAB*
AAB	*ABBAAB*	*AABABBAAB*	*AABABBAAB*
BAABB	*AABABBAAB*	*BAABBAABABBAAB*	*BAABBAABABBAAB*

The two metrics for ARN emergence/strength that are used throughout this

dissertation, and taken from the literature, are the concentration slope Λ and the mass

ratio R. The slope Λ, as described in the LITERATURE REVIEW with Figure 5 as example,

is the slope of a line fitted to the mean concentrations of the particles in a set, gathered by

length. There are multiple intricacies to the process of calculating this value. First, the

points to which the line is fit, and which make up the jagged curves in the figure, each

represent a mean of the concentrations of a group of particles. That is, not the

concentration itself, which would be the count of all particles of that length, divided by

the total particles in the system. Instead, each point is the mean of the concentrations of

all types of the given length from the set of interest (either the ARN or its complement,

the background set). However, the calculation only includes the types that have positive

counts in the system, so that zero-count types do not affect the mean concentration. The

line is fit to the points using a standard least-squared-error method, as is available in most statistical software. (The calculations for DPD models were carried out in Mathematica, along with the rest of the analysis.) The line-fitting step can also be problematic, such as when only a few particle types exist. In these cases, Λ can be highly variable over small changes in concentration, and this may not be apparent when aggregating results. For example, this behavior occurs when the system only contains polymers of length 1 and 2, such as when the mass flux is very high. Further, Λ is not defined when only one type is present. These subtleties make it difficult to describe a clear definition for Λ, and occasionally lead to values that are hard to understand. This frequently arises near extremes of system behavior, such as very high or very low mass flux.

The metric for mass ratio R is easier to understand and compute. Put simply, R is a ratio of the entire reaction network split into two groups. Because all molecules in the network are included, calculating the mass of the two parts is equivalent to finding the mass of one part and knowing the total mass. The separation rule to create the two groups is the core of the metric, and the only point of likely confusion. R is defined as the mass of the ARN set *minus* the food set, divided by the remainder of the network's mass. That remainder is, by necessity, the combined mass of the food set and the background set. The two groups are defined in Equation 6, and the ratio R is described using various combinations of their calculated masses. The food set is counted separately from the ARN to clarify certain system states. For instance, a system with high mass flux will be dominated by the food set, and the total mass of the ARN would appear very high with the food set included, though no other ARN species are maintained.

Equation 6. Calculation of the mass ratio R for a reaction network.

$$A = \{p \mid p \in \text{ARN} - \text{Food}\}$$

$$B = \{p \mid p \in \text{Background} \cup \text{Food}\}$$

$$m_A = \sum_{a \in A} \text{Mass}(a)$$

$$m_B = \sum_{b \in B} \text{Mass}(b)$$

$$\text{Total Mass} = m_T \equiv m_A + m_B$$

$$R = \frac{m_A}{m_B} = \frac{m_A}{m_T - m_A} = \frac{m_T - m_B}{m_B}$$

The main set of parameters that are used to configure both the DE model and the

DPD model are listed in Table 3. Because of differences in the two paradigms, several of

the parameters take different forms in the two models. Reaction rate constants from the

DE model are not applicable to the discrete events of the DPD model, and are replaced by

reaction probabilities. The unbinding constant is also a rate in the DE model, but is

implemented as a constant time in the DPD model, as a simplification. Mass flux is most

easily translated of these, as a simple rate of particle inflow per time in both systems.

Catalytic efficiency is almost comparable between the two models, since it functions as a

factor in reaction equations in both cases. However, the equations differ in their terms—

rates vs. probabilities—as well as their calibrations, so the factors have different effects

and different scales. The three related system properties of monomer population, density,

and volume are interdependent, with only two degrees of freedom. Both models hold

monomer population constant as a convenience. (Since all monomers are modeled with

equal mass, this assumption simplifies multiple calculations and experimental design

issues. When mass flux would increase the population, random particles are selected to

flow out of the system until balance is restored.) The DE model chooses density as the

remaining configurable parameter, while in the DPD model it made more sense to

manipulate volume.

Table 3. Analogous parameters for the DE and DPD models that approximately maximize Λ. DE values are from Table 2(a) of Bagley & Farmer (1991). DPD values are the result of an optimizing search, and are at best a local maximum.

Parameter	DE value	DPD value
Forward reaction parameter	6.490×10^2 (rate constant)	1.0×10^{-5} (probability)
Reverse reaction parameter	2.50 (rate constant)	9.5×10^{-8} (probability)
Catalytic efficiency	8.97×10^5 (factor change to reaction rate)	1.35×10^4 (factor change to reaction rate)
Unbinding parameter	5.00×10^4 (rate constant)	2.0 (time)
Mass flux (rate, both in and out)	1.79×10^1 (monomer masses / time)	2.0 (monomer masses / time)
Density or Size (only one needs to be specified)	2.0 (density)	$6 \times 6 \times \sim 0$ (size; essentially 2D)

The parameter values given in Table 3 for the DE model were selected from a

series of multivariate sensitivity analyses, as approximate values that maximized Λ in

those results. (This is the described in the caption of Table 2(a) from Bagley & Farmer

[1991].) The analogous values for the DPD model were found through an optimizing

search, in a process that varied the parameters by small amounts, keeping those variations

that resulted in an improvement of the metric Λ. The parameters were selected in random

order, and changed by a small random amount. Each parameter set was simulated 10

times and the mean of the results was taken. The results had a large random component,

and the search was certainly not exhaustive, and at best may have reached a local

maximum. Even so, the process was able to improve Λ substantially, going from an

initial value of –0.19 to a final value of –0.092 in one search process. That result was found in the 63rd iteration of that particular search, which was ended after 100 steps when no better values had been found. No claim is made that these parameters are closely analogous to the DE model parameters. Instead, the goal of the search process was to find a set of parameters for which some attempt at optimization had been made, so that the experiments of the dissertation could be initialized with values that had a better than random chance of producing emergent ARNs.

Dynamic Environment Experiments

The research method for the experiments in dynamic environmental interactions is relatively simple compared to the series of steps required to get to this point. Once the capability for dynamic interactions was successfully added to the DPD software (as described above), what remained was to run many simulations with dynamic free energy or mass flux while varying the cyclic parameters, and then to analyze the results. The objective for these experiments is to demonstrate an increase in ARN deviation from the background in terms of the two metrics R (mass ratio) and Λ (concentration slope). An increase in these deviations will be taken as a potential step toward ARN emergence, putting aside considerations of step size and determination of emergence for now.

The experimental plan that was carried out called for four sets of simulations, each of which varied one of two parameters (amplitude or period) for one of the two dynamic effects. All four sets of simulations used the same control file and non-dynamic parameters as a starting point for their variations, so that the results could be compared to that system as a baseline, as well as to each other. The baseline parameter values were

selected from the earlier simulations, including those from the model replication phase, to determine a system where ARN emergence was likely, but not where the metrics were maximized. That is, the values previously found to result in a (local) maximum for R or Λ were avoided, so that the experiments would not have a bias toward lower metrics from starting at a maximum value. Parameters that were associated with very poor metrics were likewise avoided, so the experiments wouldn't have a bias toward success either.

The four sets of simulations are summarized in Table 4, along with the baseline experiment. For each experiment, an initial range of dynamic parameter values was selected and simulated and the results analyzed. The ranges were then extended or narrowed as necessary, so that the following simulations would span a range of parameter values and a variety of results. In this exploration of the parameter ranges, only a few values were dropped, as the initial ranges seemed to provide satisfactory variety, whether by luck or by generality of that variety. As more repetitions were completed and analyzed, potentially interesting results started to appear, such as a sharp change between two sequential parameter values. In some of these cases, extra simulations were added with intermediate parameter values, to confirm and refine the results. An example of this can be seen in Figure 22, in which the R value at the mass flux amplitude of 80 showed a sharp drop from the two adjacent amplitude parameter values of 60 and 100. Simulations with the intermediate values of 70, 75, 85, and 90 were completed, and the results confirmed the interesting sharp drop near amplitude = 80.

Table 4. Summary of experiments with varied environmental dynamics.

Experiment	Cycle	Varied parameter	Values for varied parameter * effective baseline value	Cycle's associated static parameter
2.0	None	None	None	None
2.1	Energy	Amplitude	1*, 1.15, 1.3, 1.5, 1.7, 2, 3, 4, 6, 8	Period = 1000
2.2	Energy	Period	0*, 50, 100, 150, 200, 250, 350, 500, 800, 950, 1100, 1250, 1400, 1700, 1850	Amplitude = 10
2.3	Mass Flux	Amplitude	0*, 10, 20, 40, 60, 70, 75, 80, 85, 90, 100	Period = 1000
2.4	Mass Flux	Period	0*, 100, 200, 350, 500, 800, 950, 1100, 1250, 1400, 1550, 1700, 1850	Amplitude = 50

All of the varied-parameter simulations are listed in Table 4, and these same

experiments provide all data for the EFFECTS OF ENVIRONMENTAL DYNAMICS section in

RESULTS. For every varied parameter value in Table 4, 20 repetitions of the simulation

were run with random initial conditions, each extending for 10,000 time steps. For the

majority of the cases, data would only be used from the last 4,000 steps, so that the

system had a chance to reach an approximation of a dynamic steady state. In experiments

where the varied parameter was a cycle period, a time span of fewer than 4,000 steps was

used when necessary to include a whole number of cycles in the data range. This was

done to ensure that no bias was introduced by data from a fractional number of cycles.

For each set of simulations and its associated time range, the data that had been

recorded by the model were extracted and grouped together. Of the outputs that the DPD

software can produce, the main source of data for these analyses was the 'census,' which

provides a list of all existing particle types with quantities. This list, along with contents

of the corresponding network graph (Table 2 or Figure 3), provides all of the information

that is required to calculate both metrics of interest: the mass ratio R and the

concentration profile Λ. The metrics can be computed for each point in time when the census was recorded, which was typically specified to be every 20 times steps for these experiments. In the span of 4,000 time steps, there would be 200 samples of the census data, each resulting in counts of all polymers and values for R and Λ. The metric values could also be computed statistically over a time series with basic tools such as mean, standard error, variance, or a moving average. Alternatively, the metrics could be aggregated across the 20 repetitions to find how behavior varies statistically over the random initial conditions. The time series for either metric could be plotted for an entire set of repetitions to get a visual understanding of the behavior and variation, along with the mean time series created by aggregating the repetitions at every time point. The mean time series might also be plotted along with the analogous curves for other parameter values, as well as the values from the non-dynamic baseline system. Figure 14 is an example of data condensed across time and repetitions to a few values, allowing for the comparison of behavior from different parameter values. Figure 17 provides a less condensed view, showing more detail in the form of time series, though fewer parameter values can be viewed at once. The combinations and variations of these analytical steps provided all of the results that appear in the EFFECTS OF ENVIRONMENTAL DYNAMICS section of the RESULTS chapter.

4. RESULTS

The results of this dissertation are presented here in two sections. The first section provides the results of experiments designed to show that a dissipative particle dynamics (DPD) model of an autocatalytic network is comparable to other models in the literature, specifically the differential equation model of Bagley and Farmer (1991). These results show substantial similarities between the models in aggregate, qualitative behavior, despite the fundamental differences in the modeling paradigms. The second section presents the results of the experiments designed for the main objective: discovering whether dynamic environmental interactions can promote ARN emergence by increasing both the probability of ARN emergence and the variety of conditions in which ARNs emerge. This objective has been met multiple times by the results from DPD model simulations with dynamic environment flows of both energy and matter, in terms of the aggregate measures established in the ARN literature.

The two principal metrics used here for quantifying ARNs are *concentration slope*, denoted Λ (the slope of a linear fit to a molecule set's mean concentrations, as a function of polymer length) and *mass ratio*, denoted R (the ratio of mass in the ARN without the food set over the combined mass of the background and food sets). Concentration slope is by definition negative, with larger (less negative) values signifying higher concentrations than a set with a smaller, more negative Λ. For an ARN to demonstrate emergence, or separation from the state of the rest of the system, Λ for the ARN can be compared to Λ for the background set in the same simulation, or to a Λ value for the dynamic steady state of a reaction network without catalysis. For these

experiments, both the background and the steady state values for Λ are near –0.4, while

an emerged ARN may have Λ closer to –0.2. The background value varies in simulation,

in response to the ARN or to system parameters, though Λ for the ARN is likely to see

more variation. Because of this interactive variability, concentration slope is typically

compared between ARN and background, rather than the steady state, since the ARN and

background are likely to be responding to the same system parameters or to each other.

Unlike Λ, the mass ratio R is a single-value measure for an entire reaction network,

rather than measuring a specific set of molecules. Mass ratio ranges from 0 to 1, with

higher values marking a healthier ARN in control of more of the system's mass. A typical

R measurement of a "strong" ARN might be 0.5 or higher. These two metrics will be

explained further in context, below. They are also detailed in the LITERATURE REVIEW

(p.15) and RESEARCH METHODS (p.38), and their advantages and disadvantages as

metrics are addressed in the DISCUSSION chapter.

Comparing the Models

As described in RESEARCH METHODS, four of the graphical results from Bagley &

Farmer (1991) were selected as focal points for replication, to establish a comparison

between the DE and DPD models. These four results were selected for being the most

informative, for being the most detailed in the source paper, and for having the highest

chance of success. There are still several sources of uncertainty for the replication goals

from the B-F methods. These include certain details of DE model execution, of metric

calculation, of parameter implementation, and particularly of the optimization process

used for finding optimal parameter sets. The replication targets were selected to minimize the effects of these uncertainties, but they could not be avoided entirely.

The four results cover several aspects of the behavior of the DE model for reaction networks containing ARNs. Each of the original four figures is included below (with permission) along with an analogous plot of data from experiments with the DPD model. The text accompanying each pair of figures provides detailed explanation and comparison of the results, though the majority of the interpretation is concentrated in the DISCUSSION chapter. Additional details of the replication process that are not relevant to the results were described in the RESEARCH METHODS chapter.

The first two plots, Figure 6 and Figure 7, show the concentrations of molecule types over time within a single simulation run. Figure 8 and Figure 9 show concentration levels by polymer length, from which the concentration slope Λ is computed. In Figure 10 and Figure 11, the sensitivity of the concentration slope Λ is plotted against changes in mass flux, with data collected from separate simulations for a series of flux rates. Using the same set of simulations, Figure 12 and Figure 13 plot the impact of mass flux on where the mass resides in the system. In this single case, the B-F result deconstructs the mass ratio R into its components of Network, Food Set, and Background to see how the proportion of each varies with the changes to mass flux.

Result 1.1. Polymer Concentrations

The B-F plot in Figure 6 demonstrates the behavior of a system transitioning from an early state with few polymer types present to a dynamic steady state in which an ARN has emerged. In this model, the system was initialized with only food set types (of length

1), and it reaches a steady state with 22 molecule types present above the *threshold level*.

The threshold level in the DE model is the concentration for which at least one molecule

is consistently present in the system. For instance, in a model that represents 10,000

molecules, only types with a concentration of (1/10,000) are considered present. This rule

is necessary because the DE paradigm uses continuous values for population counts,

which can cause quirks such as fractional molecules.
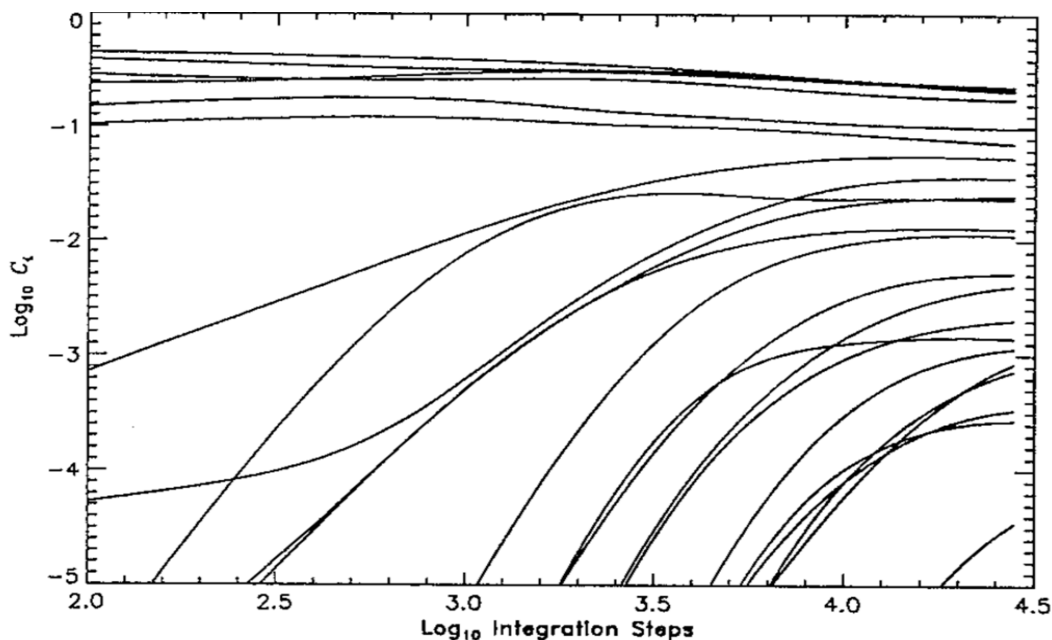
Figure 6. Polymer concentrations over time (DE model).
Each curve plots the concentration of a molecule type, through a single run of the DE
model. The x-axis is Log_{10} Integration Steps and the y-axis is Log_{10} Concentration,
measured as the molecule type's fraction of the total reacting particle count. (Reprinted
from Figure 3 of Bagley and Farmer [1991].)

The figure begins 100 time steps into the simulation, with curves for the

concentrations of each the 8 molecule types already present. Most likely, the uppermost

curves are for the concentrations of short food-set monomers, which have concentrations

between 0.1 and 0.4, and which together make up the majority of the reaction network.

The remaining curves represent types that derive from reactions within the food set. These are likely to be short molecules from the ARN that are closest to the food types in the reaction graph. (The ARN reaction graph for this simulation is not the same as in the models below and is not specified in the literature, thus the ambiguity in the molecule types.) The lowest line at this time marks a concentration below 1/10,000, with the next lowest around 1/1000. When existing molecules achieve high enough concentration to produce another new type above threshold level, that new type enters the plot from the bottom. The typical arc for each of these is a rapid rise in concentration until it too begins engaging in condensation reactions, at which point the curve levels off, perhaps with a slight decline. The right-hand edge of the plot shows 22 types in existence after $10^{4.5}$ (approximately 31,000) time steps, and the system is described as reaching its steady state with that set of types (Bagley & Farmer 1991).

The results from the DPD model (Figure 7) show behavior that is qualitatively similar to the results from the DE model, but with conspicuous differences, such as the variability—or lack of smoothness—in each curve. This figure displays data from a single run of a DPD model that was started with only food-set particles, as with the DE model above. In both figures, the results show high concentrations in the food set at the start, accompanied by a few other particle types that are nearby on the reaction graph. As time passes (moving from left to right), new types appear and rise in concentration. The types that dominated at the start decrease in concentration as their molecules react to produce other, longer types, just as in the previous results.

Unlike the DE model described in Figure 6, the concentration curves from the DPD model do not appear to be smoothly approaching a steady state. Each curve remains

variable throughout the timespan of the simulation and experiences many inflection points, as opposed to 0-2 in the DE model. Another difference between the results is seen in the curves in Figure 7 that rise from the bottom of the plot and later return to the bottom. These represent molecule types that exist in the system temporarily but whose concentration later returns to zero, after all molecules of that type are consumed in reactions. (They could also be removed from the system by the mass flux, typically with a much lower probability.) The DE model does not exhibit this behavior, and perhaps cannot, due to the paradigm's larger populations and deterministic reactions, as well as the threshold rule.
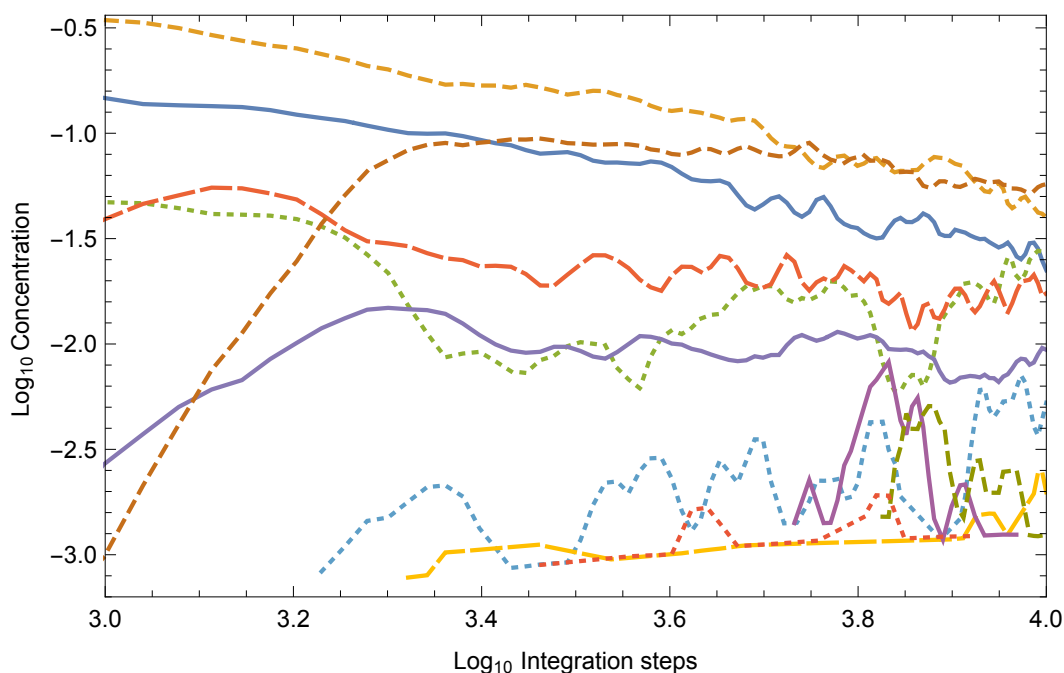
Figure 7. Polymer concentrations over time (DPD model).
Each curve describes the concentration of a molecule type over a single run of the DPD model. Early in the run, food types are most common (at top), and only a few types are present. As time continues, food concentration drops and more, longer types exist. The rising "floor" to the right shows how the minimum concentration rises, as there are fewer molecules in the system. Mass is conserved, so average molecule length must increase.

Some of the difference between the two models could perhaps be corrected with a looser definition of "steady state," or by smoothing the DPD data over time or across multiple simulations. Such repetition is standard in the DPD paradigm, where each model configuration is run many times with different (random) initial conditions, and the set of results are combined to give a statistical description of system behavior. That method is used throughout this dissertation, for most of the results that follow. However, in this case it is useful to look at the results of a single simulation run, to demonstrate and understand some of the key differences. The variability in the DPD model's concentrations is intrinsic to the paradigm, and it enables interesting behavior that is not possible in the DE model, such as the types that exist briefly and randomly. Results that showed only the mean concentrations of many runs would seem to dampen the variability of the DPD model, and might therefore look more similar to Figure 6, but such results would be incomplete and potentially misleading.

Result 1.2. Concentration Profiles

Another visualization of the steady state of a DE model run is presented in Figure 8, as the second result to be replicated from Bagley & Farmer (1991). This plot describes the "concentration profile" of a single simulation through two lines, one for the ARN (N) and one for the background (B). Each line plots the mean concentrations of member polymers grouped by length, which is on the horizontal x-axis.. The graph also includes a line (E) that marks what the concentrations would be for a non-catalyzed version of the system, as reference. For both the model lines and the non-catalyzed line, the concentration values are taken from a system that has reached a dynamic steady state.

Each line can be reduced to a single value by calculating its slope, or the slope of its least

squares linearization, for the jagged lines. This single value is the measure Λ, or

concentration slope, which can be used to concisely describe an ARN or other set.
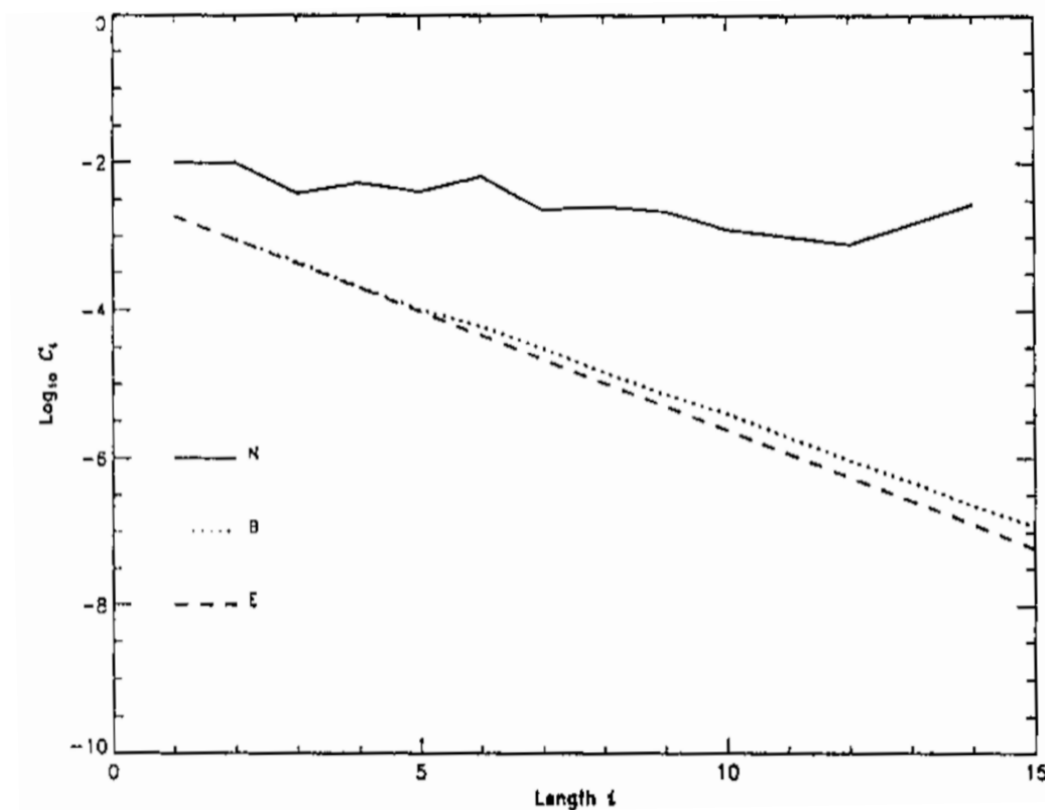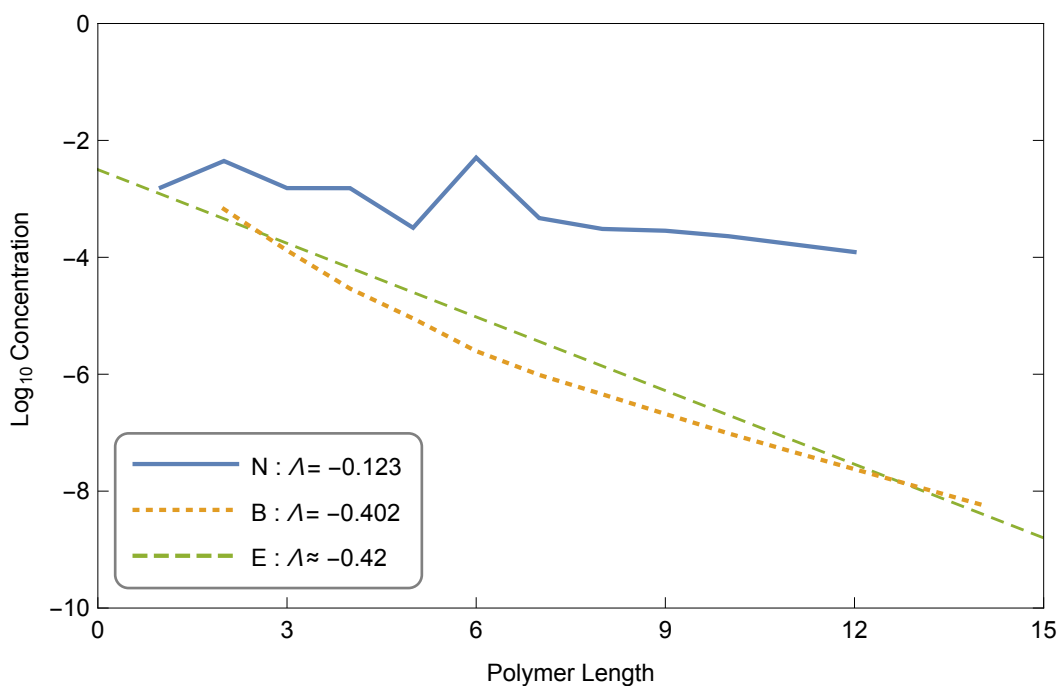

Figure 8. Concentration profile of a single simulation run (DE model).
This plot shows mean concentrations for an ARN (N) and its Background (B), and also a
non-catalyzed system for reference (E). Precise Λ values were not provided with this
figure, but manual measurement suggests Λ for N is near –0.1, and near –0.3 for both B
and E. The x-axis is Polymer Length (from 0 to 15), and the y-axis is Log_{10} of Mean
Concentration (from -10 to 0). (Reprinted from Figure 6 of Bagley & Farmer [1991].)

The analogous plot for a DPD model appears in Figure 9. The ARN reactions in

both models are defined by the same reaction graph (Figure 3), but they differ

substantially in their kinetic parameters. This is due to model paradigm differences and

the inability to match unknown aspects of the DE model, such as the method for selecting

the 'optimal' parameter set from a large state space. Despite the differences, the
concentration profiles in the two plots are qualitatively similar, especially in the way the
segments of the ARN profile alternately climb and fall through their first seven points in
both figures. The ARN profile for the DPD model is more jagged and angular, but it is
still quite similar to the DE model results.

Figure 9. Concentration profile of a single simulation run (DPD model).
This plot shows mean concentrations for molecules in the ARN (N) and its Background
(B). Also plotted is an approximate value for the concentrations of a non-catalyzed
system (E). This model used the same parameter configuration as the models for Figure
11 and Figure 13, with a Mass Flux of 2.0.

The similarity between the Network and Background slopes within each figure is
also comparable between the two, even though the parameters of each are derived from
separate optimization processes. The Λ quantities appear to be different between the two
systems, but this is not surprising. Some similarities could be predicted, like the zigzag of

the ARN profile that derives largely from the shared ARN reaction graph. In other cases, quantitative similarity was not expected, as with the slopes of the background and non-catalyzed lines. These lines are determined by the kinetic parameters of the artificial chemistry, such as the ratio between forward and backward reaction rates, and these were known to be different between the models.

The two figures do have some differences. The Network line in the DPD model (Figure 9) is shorter, only reaching to polymer length 12, while the DE model reached to length 14. This is likely to be due to the smaller population of the DPD model. Given that the longest polymer types are absent and that the population is between 1 and 2 orders of magnitude smaller in the DPD model, it is interesting that the results are as similar as they are. The other significant difference is that the Background line for the DPD model is not straight, though it appears to be in the DE model's figure. This could be attributed to the increased variability in the DPD model's chemical reactions, or again, to the differences in population size.

Result 1.3. Sensitivity of Concentration Slope to Mass Flux

The third result selected for replication is a sensitivity analysis of the concentration slope Λ to variation in *mass flux*. Mass flux (denoted δ in B-F figures) is a key parameter in both models, as described in RESEARCH METHODS. It is the rate at which new mass enters the system, and that influx is one way in which the system is driven away from chemical equilibrium, which is necessary for ARN emergence. New mass enters the system in the form of particles drawn randomly from a "food set" of simple types, which are usually monomers or short chains. Because mass (and monomer count,

in these models) is designed to be constant, the rate of incoming mass is matched by the rate of exiting mass, which works by selecting particles randomly for removal. As a consequence of this, a higher mass flux translates to a higher probability of displacing existing molecules. If the incoming mass consists of monomers, as is typical, the displaced molecules are more likely to be longer (since they can't be shorter), and so mass flux can have a diluting effect. As mentioned above, some minimum level of mass flux is required to drive the system away from equilibrium, and that driven separation is what allow catalyzed reactions to process (or focus) more than their share of the mass in the system.
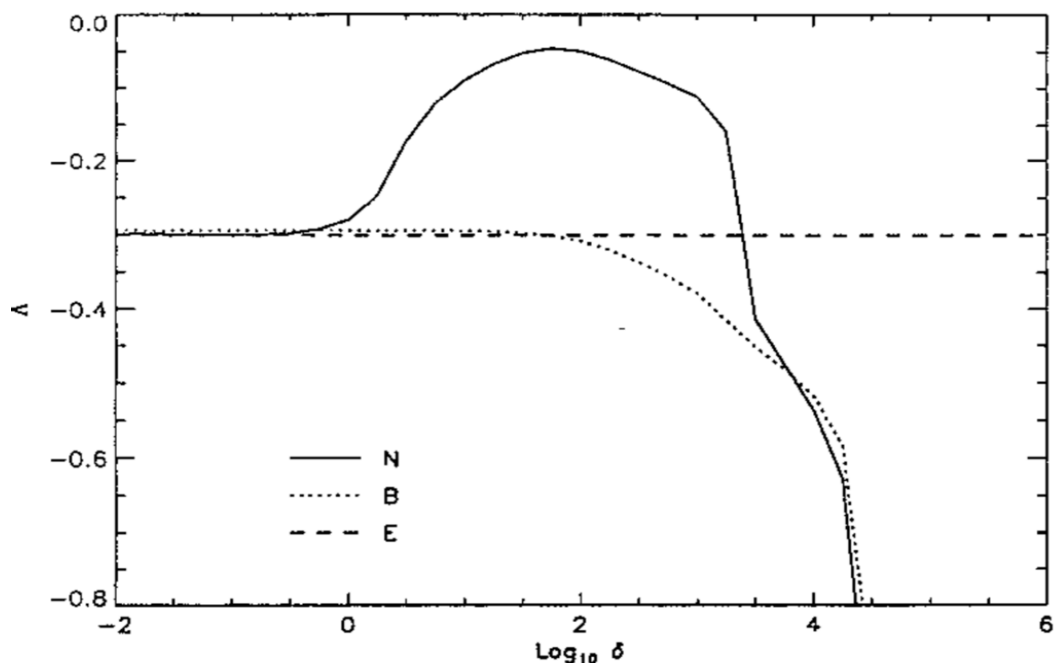
Figure 10. Sensitivity of concentration slope Λ to varying mass flux (DE model).
Lines mark Λ for the ARN Network (N, solid), the Background (B, dotted), and a non-catalyzed Steady State (E, dashed), from a DE simulation for each mass flux value. The x-axis is Log_{10} of Mass Flux, and the y-axis is the concentration slope Λ. (Reprinted from Figure 7 of Bagley and Farmer [1991].)

These issues at extreme values of mass flux suggest a "sweet spot" in the middle where ARN emergence is most likely. This sweet spot is apparent in the results of the sensitivity analysis of Λ to mass flux, shown for the DE model in Figure 10 (Bagley & Farmer 1991). The upper values for mass flux seem to overwhelm the system, so that Λ for both Network and Background drop off the chart. (Such low slope values are the result of the system being dominated by monomer food types. The "linear fit" to get slope becomes less meaningful when there are very few points to fit.) At the low end of the mass flux range, Network and Background again have equal Λ values. At such a low flux, there is no "drive" to the system to give the catalyzed reactions their advantage.
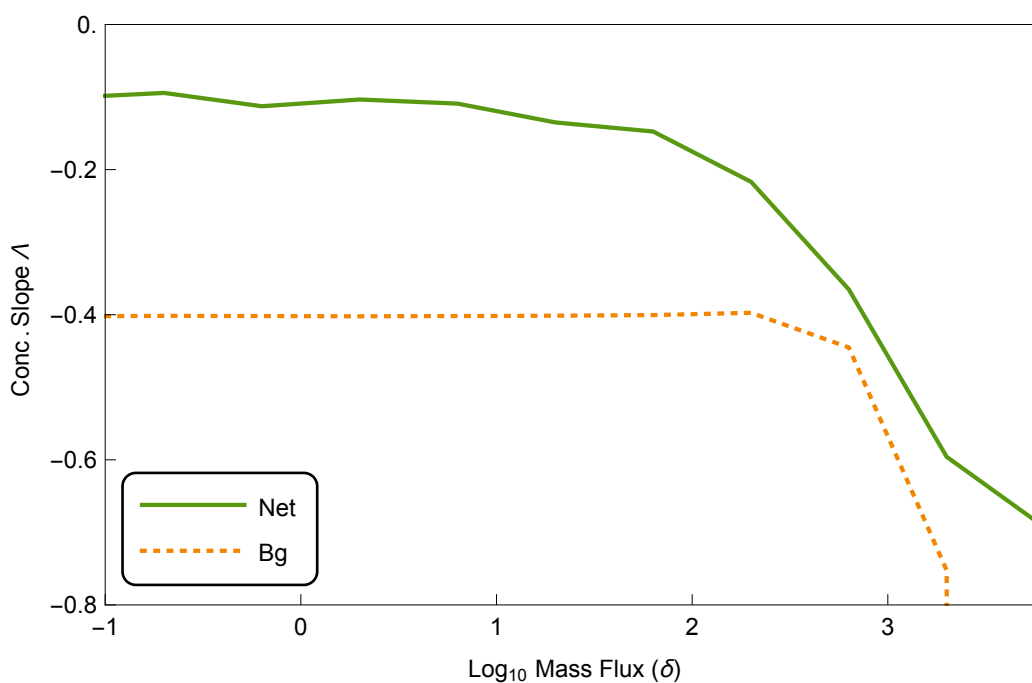
Figure 11. Sensitivity of concentration slope Λ to varying mass flux (DPD model).
Values are for the ARN (Net) and Background (Bg), averaged across 20 DPD simulations for each mass flux value. The Steady State Λ of approximately –0.4 is not pictured.

When a similar sensitivity analysis was executed with the DPD model, the results included the only substantial difference between the two models (Figure 11). As with the DE model, the highest values of mass flux are associated with a precipitous drop in Λ for both the ARN and the Background sets. However, unlike the DE model, the DPD model shows the ARN maintaining a Λ value well above the Background level, even at a very low mass flux. The reason for this is not certain, but investigation of these simulations suggests that it is due in whole or part to spatial aspects of the DPD model.

As noted previously, a system needs to be driven to have the potential for ARN emergence (among other conditions). In both the B-F models and the dissertation's models, the design is for the driving force to be an influx of mass. This ongoing influx makes it possible for catalyzed reactions to be more "productive" than non-catalyzed reactions, effectively changing the equilibrium point of those (catalyzed) reactions, since the system is being pushed away from equilibrium. In a standard chemical system (that isn't being driven), catalysis changes the rate of reactions but not the equilibrium point, and the system reaches the same dynamic steady state eventually, whether or not catalysis occurs. The understanding of this behavior is built into the design of the DE model and the research method around it, and it has been carried over into this dissertation. However, these features of chemical kinetics may rely on assumptions that have been violated in the process of translation from DE to DPD, or even in the creation of the DE model itself. This requirement that a system must be driven as prerequisite for ARN emergence may be based on assumptions that no longer apply. Specifically, it appears that the DPD model as parameterized in this experiment is very far from "well-mixed"

due to its flattened spatial geometry, so that it is effectively two-dimensional. Particles

are only able to react with those nearest them, and the supply of neighbors is not

replenished quickly (perhaps because mixing is slower than reacting), which means that

the increased reaction rate (or probability) of the catalyzed reactions has an effect that

lasts over time. It may be that the local region around a cluster of catalyzed reactions gets

"driven" by the wider neighborhood, as particles happen to migrate into range and enter

into the ARN.

This analysis is not complete, as the existing set of simulations doesn't provide all

of the data necessary to understand what happened, and running a complete new set of

simulations is beyond the scope of the dissertation. However, the hypothesis above does

suggest a quick test that could falsify it. If the elevated Λ at low mass flux is due to the

spread-out 2D spatiality of the DPD model, then a comparable 3D configuration ought to

behave differently. An *ad hoc* addition to the Research Method was devised to test this,

and a limited set of simulations was designed and run. The DPD model of Figure 11 has

dimensions of 6×6×1, though the 1 acts much like a 0 due to other parameters. Three new

configurations were generated from one of those used above (with mass flux = 0.2), but

with a proper third dimension and depth values of 2, 4, and 6. This progression allowed a

prediction that Λ should get smaller as geometry changes and mixing becomes more

prevalent, if the hypothesis is correct. If there were no change, or no clear progression,

the hypothesis would be less likely. The results of the *ad hoc* simulations are in the table

below, along with the comparable ARN value from Figure 11. The results do not

invalidate the hypothesis that spatiality is the cause, nor do they fully explain the

discrepancy, though they show a pattern of increasing Λ. The remaining difference may

be due to the density of the particles in the system, which decreased as depth increased

and other values remained constant. Exploration of this possibility is left for future work.

Table 5. Summary of Results for *ad hoc* Dimensionality Experiments.

Dimensions	Λ
6×6×1 (Figure 11)	−0.0943
6×6×2	−0.170
6×6×4	−0.199
6×6×6	−0.213
Non-catalyzed	−0.4

Result 1.4. Sensitivity of Mass Distribution to Mass Flux

The final result targeted for replication is also a sensitivity test as a function of

mass flux, this time characterizing the distribution of system mass as fractions of the

whole, by molecule set. The results for this experiment were extracted from the same

simulations as in the previous experiment, for both the DE and DPD models. The system

is measured this time in terms of mass, and how the total reactant mass is split between

three different portions: the Background, the Food Set, and the ARN, or Network.

Previously, when calculating Λ or defining the ARN's network graph, the Food set has

been defined to be part of the ARN set. They are counted separately here so that the

sensitivity to mass flux can be more clearly seen, especially at high or low flux. (These

results differ from the Part 2 analyses that use the metric R or *mass ratio*, which is a

scalar value that characterizes system mass in a different way.) The values in Figure 12

and Figure 13 are simply counts of the total number of monomers in each set, displayed

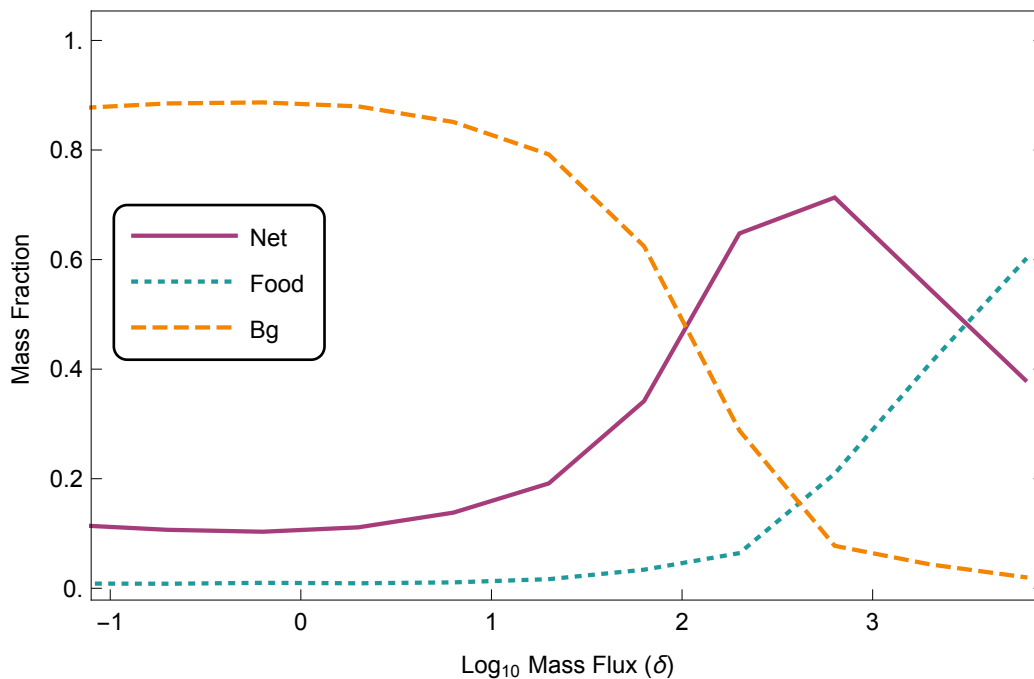as fractions of the (constant) total in the system.
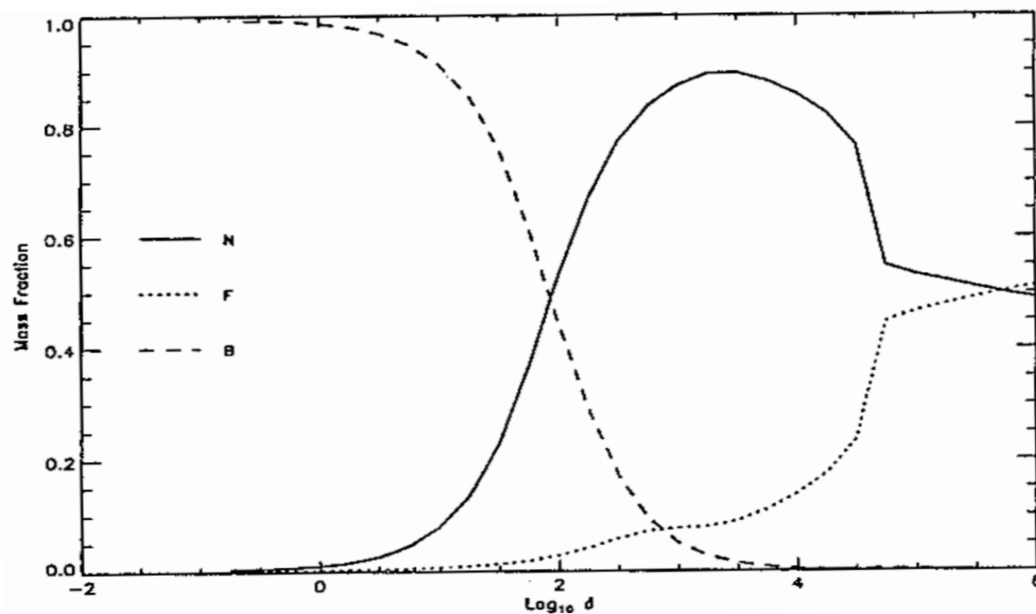
Figure 12. Sensitivity of mass distribution to varying mass flux (DE model).
Values are for the molecule sets of the ARN *minus Food* (N, solid), Background (B, dashed), and Food (F, dotted). The x-axis is Log_{10} mass flux, and the y-axis is mass fraction. (Reprinted from Figure 7 of Bagley and Farmer [1991].)

The mass distribution results for the DE model appear in Figure 12 as three curves, each tracking the fraction of mass held in a particular set. Since the three sets cover all possible molecules without overlapping, their curves sum vertically to 1.0 at all points. From left to right, the notable behavior includes the decline of the Background fraction from its maximum near 1.0, while the ARN set increases like a mirror image. Those two curves intersect, still mostly symmetric, but at higher flux values the Food set begins to rise noticeably. Continuing to the right, and even higher mass flux, the Background fraction bottoms out near zero mass. The ARN fraction peaks near 0.9 then drops, as the mass influx fuels the continued rise of the Food set. At the high end of the tested mass flux range, the Food set overtakes the Network set, as expected, and is likely to dominate mass distribution at all higher flux values. In this high-flux range, the inflow

of mass particles at every time step is approaching the order of magnitude of the system

population, so that the population is replaced too frequently for longer chains to have

time to form. Essentially, the inflow is flooding the system, washing the contents away

faster than reactions can happen.

The DPD model's mass distribution results (Figure 13) appear to replicate those

from the DE model very well. The qualitative behaviors described above are all present,

though with more variability and without the asymptotic approaches to the upper and

lower bounds. The two figures almost seem to be aligned on the mass flux scale, with

similar behavior near the middle (where Net and Background cross, near Log_{10} Mass

Flux = 2). The two sets of results are not aligned elsewhere, though, and there is little

reason to expect that they should match in this way given their many differences. One

feature that the results do have in common is their relation to the figures of Result 1.3.

For both the DE and the DPD models, the Λ value of the ARN drops at a lower mass flux

than the ARN's mass fraction does. In the DE model, Λ peaks at a flux just below 10^2,

while ARN mass peaks around $10^{3.5}$ – just as Λ is dropping most precipitously. The DPD

results are similar, with Λ declining rapidly between 10^2 and 10^3, the same interval in

which ARN mass rises to its peak. In both cases, this interval may mark a transition from

a region of lower mass flux in which the ARN is able to take in new particles and spread

them through its reactions, and a region of slightly higher flux in which the ARN is still

capable of absorbing the new particles, but perhaps cannot move them through the

network of reactions quickly enough to counteract the attrition to longer particles caused

by the increased out-flux. This is merely a hypothesis at this point, as this behavior was

not investigated in depth as part of the dissertation.

Figure 13. Sensitivity of mass distribution to varying mass flux (DPD model).
Values are for the molecule sets of the ARN *minus Food* (Net), Background (Bg), and
Food (F), averaged across several DPD simulations for each Mass Flux value.

In the results of Figure 13, there is again a slightly elevated strength to the ARN

at the lowest flux values, relative to the DE model's results. This appears to be the same

issue as in the previous experiment (Figure 11), which in that case caused Network Λ to

be elevated above Background Λ in the same range of low flux values. This issue is not

explored any further here, as it was addressed substantially in the previous results.

Effects of Environmental Dynamics

The second sequence of experiments was designed to answer the dissertation's

main question: can structured environmental dynamics promote the emergence of ARNs

in a simulation of artificial chemical reactions? The experiments focus on two types of

dynamic contributions to the chemical system by its environment: free energy and mass

flux. Both interactions are given the form of sine waves, and the system is tested for

sensitivity to variations in the amplitude or period of these waves. Those two parameters

are varied singly, in separate sets of simulations, for each of the two interaction types.

That is, either period or amplitude is varied while the other is held constant, for either

mass flux or energy, for a total of four sets of simulations. As a control and baseline, a

fifth set of simulations was run without variation in either free energy or mass flow, using

the parameter values common to the other simulations. For all of these simulations, all

other configuration parameters are identical and held constant. As described in RESEARCH

METHODS, these base settings were selected to be within a range where ARN emergence

occurs. However, the "optimal" parameters from earlier results were purposefully

avoided, so that these experiments did not begin at a (possible) local maximum, which

might bias the results toward lower metric values.

Sensitivity to each of the dynamic variations is described below, in terms of the

metrics R (mass ratio) and Λ (concentration slope). These two metrics are considered

separately for each set of simulations as well as together, to reveal which conditions

affect the metrics independently, and which seem to exhibit an interaction between R and

Λ. The sensitivity analysis is presented in overview for all tested cycle parameters, and

then focused attention is given to selected specific values. Throughout all four sets of experimental variations (free energy amplitude, free energy period, mass flux amplitude, and mass flux period), baseline values from the static (in energy and mass flux) control system are included as common reference points. These baseline values appear at energy amplitude of 1, and energy period, mass flux amplitude, and mass flux period of 0. (The value is '1' for energy amplitude because it has a multiplicative effect in the model, while the other parameters are additive.) Each of the dynamically varying contributions is approximately conserved over the length of each simulation, so that the total free energy and total mass flux is comparable between any of the simulations, including the non-dynamic baseline.

Table 6. Summary of Results for Experiments with Dynamic Environments.

| Experiment	Cycle	Varied parameter	Static cycle parameter
2.1	Energy	Amplitude	Period = 1000
2.2	Energy	Period	Amplitude = 10
2.3	Mass Flux	Amplitude	Period = 1000
2.4	Mass Flux	Period	Amplitude = 50

The four dynamic-environment experiments are summarized again in Table 6, with their varied parameters. Further details of the design are given in RESEARCH METHODS, and the results are explained in depth below. The default format for the measures R and Λ in the results below is as a mean over a time series of values and over many simulation repetitions, often with the standard error of the mean as an indicator of mean variability. The simulations for Part 2 were given at least 20 repetitions each, and values drawn from time series typically include the last 4000 time steps from a run of 10,000 steps. The number of time steps is modified to include a whole number of cycle periods when necessary, so that results are not biased by partial cycles. The repetitions

are necessitated by the randomness and variation inherent to the DPD modeling paradigm. Additional details on these decisions are provided in RESEARCH METHODS.

Assessing an ARN's deviation from the background can be ambiguous, as witnessed by the literature's definition of ARN emergence that is based on a "significant departure" from the non-catalyzed steady state, and is described as "subjective" and "vague" (see LITERATURE REVIEW section AUTOCATALYTIC REACTION NETWORKS; Bagley & Farmer 1991). The two metrics from the literature that have been selected for measuring the strength of ARN emergence—R and Λ—are not completely independent, and are of varying descriptive value depending on the system's state. Because of the lack of a clear and objective metric, the results presented below rely on comparisons of R and Λ values between configurations within an experiment and to non-dynamic baseline values. An increase in both measures could be considered the best signifier of improvement in ARN emergence, though the results more often exhibit an increase in only one metric. Therefore, simulations that show a substantial increase over the baseline in one or both metrics are taken to signify a positive result toward the dissertation's objective. Findings are also identified where the metrics demonstrate consistent trends across several sequential parameter values. Success for the environmental dynamics objectives requires only the demonstration of the existence of conditions that promote ARN emergence beyond what occurs at baseline levels, even if a detailed understanding of the distribution of such results and their causes is not readily apparent. Additional results of interest will be noted as appropriate, as sources of interesting hypotheses or as inspiration for future work.

*Result 2.1. Dynamic Free Energy, Varying in Amplitude*

In this experiment, the free energy of the system is varied cyclically throughout each simulation, and the amplitude of the cycle is varied between simulation sets. Of these four experiments in Part 2, this one features a slightly different method for applying the environmental dynamic to the system. Due to the internals of the DPD model's chemistry, changes in free energy must be made as multiplicative factors rather than additive offsets. The result of the multiplication is transformed into a reaction probability via an exponential formula similar to the Arrhenius equation (see Equation 3 and the surrounding explanation in RESEARCH METHODS).
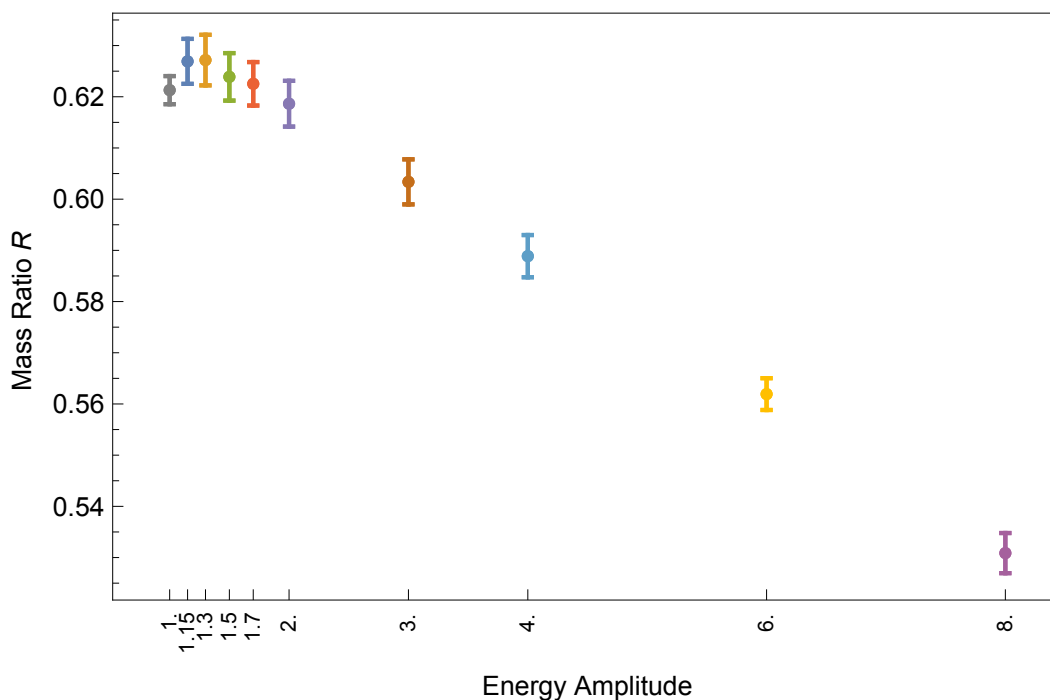
Figure 14. Sensitivity of *R* to varied amplitude of dynamic free energy.
Each bar represents a mean (at center) and its standard error for a set of simulations. Each of these sets shared identical parameters but some different random values, such as initial molecule positions. The amplitude of a cycle energy (heat) component was varied between the sets of simulations, and this value is marked on the horizontal axis. The bar at amplitude=1 shows the results for a system with a non-dynamic energy level, for reference. (Colors have no significance here, other than to differentiate the bars.)

The first positive finding appears when varying the amplitude of free energy (Figure 14), which gives the system a small positive increase in R over a small range near the baseline value of 1, and above that range, R declines consistently and smoothly.
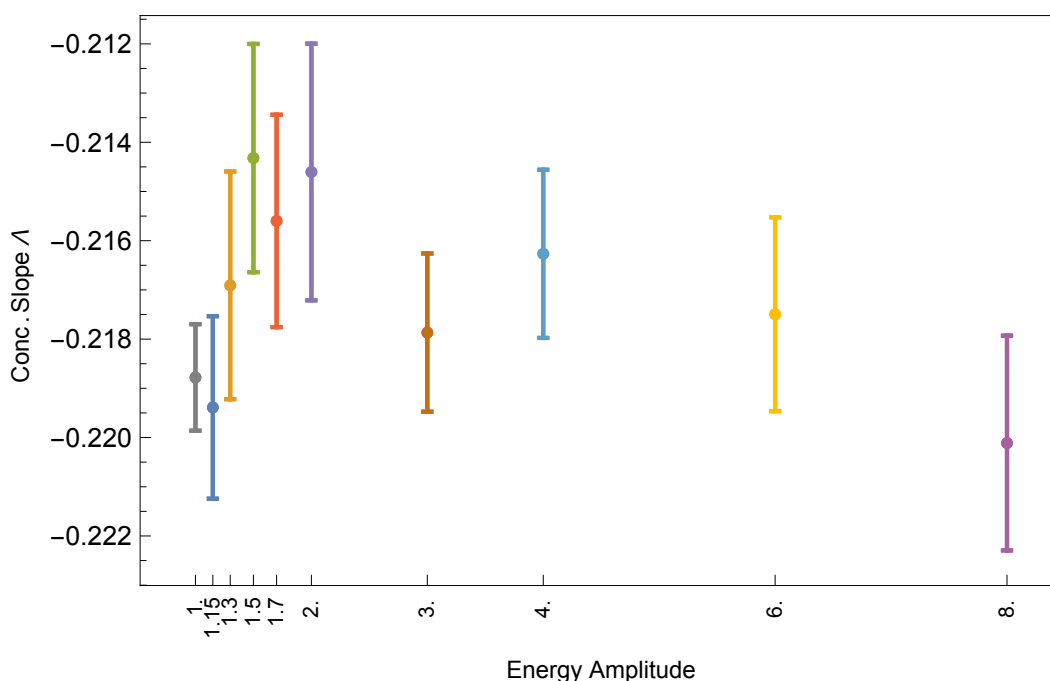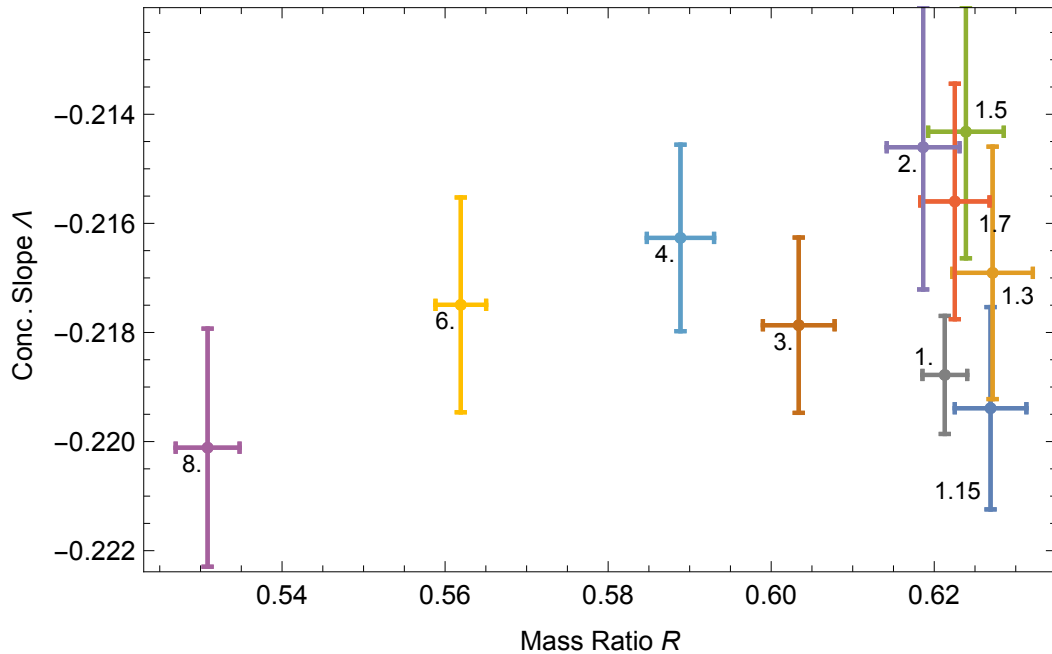

Figure 15. Sensitivity of Λ to varied amplitude of dynamic free energy.
As with Figure 14, each bar here is the mean and standard error of values from a set of simulations. Here, the values are the concentration slope Λ, which provide a measure of the ensemble concentrations for the ARN. A higher value for a system suggests that the ARN is maintaining concentrations at higher levels than for non-ARN molecule types. Variations that increase Λ over the non-dynamic reference value (shown at amplitude=1 here) are evidence that cyclic energy with those parameters is beneficial to the emergence of an ARN. As before, the colors of the bars are here only to help differentiate the results.

A second positive result is found in Figure 15, also among the smaller values for amplitude, but in this case correlating with an increase in Λ relative to the baseline value shown at amplitude=1. These results for Λ have much larger standard error values than were found for R, relative to the differences between means, and the trend among the

means is not as smooth. Throughout the results below, the slope metric frequently shows more variance than the mass metric, but rarely is the inverse true.

Figure 16. Sensitivity of *R* and *Λ* to varied amplitude of dynamic free energy. Extending Figure 14 and Figure 15, each crossbar here marks the means and standard errors for the same sets of simulations, but now combining the values from those previous figures. Each crossbar is labeled with an amplitude value, which was marked on the horizontal axis previously. The crossbar at amplitude=1 continues to mark the reference value of a system with a constant energy level, without cycles. This figure makes it easier to see which amplitude values gave the best results in both parameters. The colors of the bars are arbitrary, but consistent between like figures.

In Figure 16, the results in *R* and *Λ* have been combined into a single plot. This format can make it difficult to see the sequence in the varied parameter (energy amplitude), but it clarifies which of the values are improvements in one or both metrics relative to the baseline crossbar. In this case, no values show clear improvements in both metrics, though the values for amplitude 1.3, 1.5, and 1.7 come close.
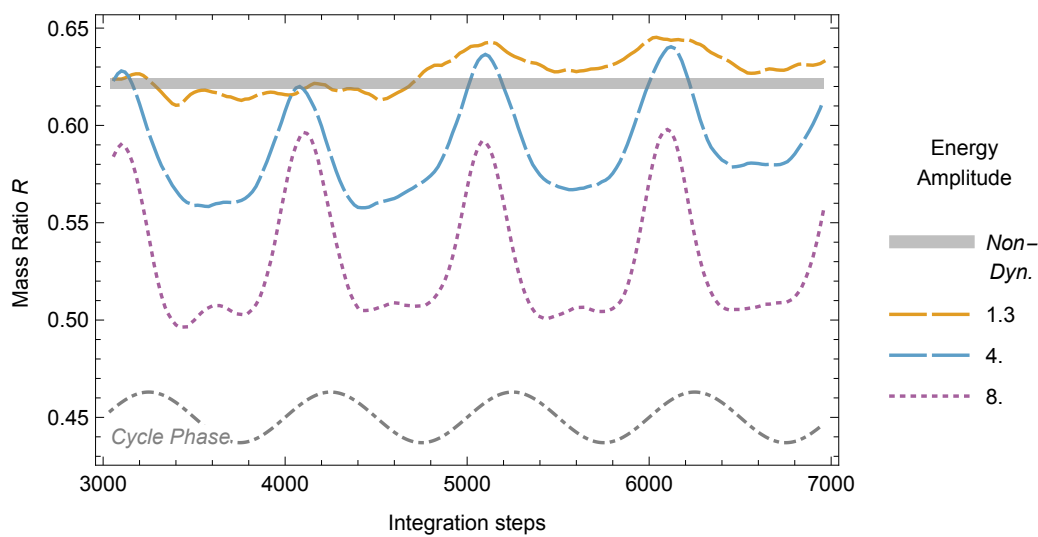
Figure 17. Time series of _R_ response to varied amplitude of dynamic free energy.
Each line is the mean of 20 simulations. The mean of the static baseline is included for
reference, with thickness showing ± standard error. The "Cycle Phase" at bottom shows
the phase of the dynamic energy.

Figure 17 and Figure 18 display time series for the metrics **_R_** and **_Λ_** respectively,

for a few select values of energy amplitude. They also include the baseline values for a

system with static (non-dynamic) energy, like the previous figures, but instead of a

crossbar, the standard error is now shown as line thickness. Also in Figure 17, a sine

wave of arbitrary amplitude is included to show the phase of the free energy cycle. These

figures don't necessarily reveal new positive findings, but they help to understand the

behavior of the system. For instance, it appears that the rising wave of the energy cycle

coincides with a sharp rise and fall in **_R_**, especially at higher amplitudes. In Figure 18, the

curves for amplitudes 4 and 8 experience sharp highs and lows, but it is difficult to

distinguish their mean behaviors from each other, or the mean of the static baseline. It

appears that the system with energy amplitude of 1.3 improves in both metrics when the

energy cycle is low, and is able to hold onto those gains more often than losing ground when energy is higher.
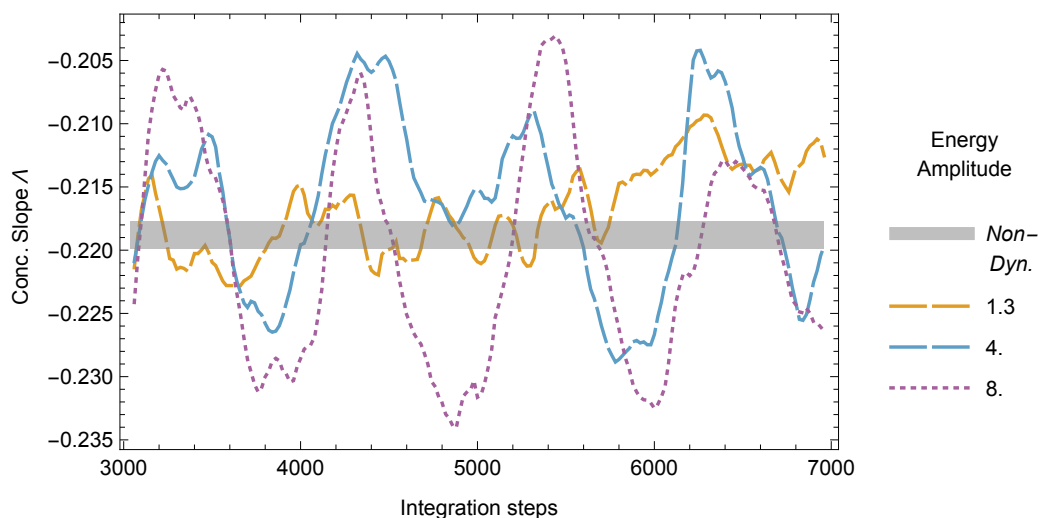
Figure 18. Time series of Λ response to varied amplitude of dynamic energy.
Each line is the mean of 20 simulations. The mean of the static baseline is included for reference, with thickness showing ± standard error. (The cycle phase, not shown, is the same as in Figure 17.)

Result 2.2. Dynamic Energy, Varying in Period Length

In this set of experiments, the varied parameter is the period of the dynamic energy interaction from the environment. It is varied over a wide range while the amplitude is held constant with a value of 10. Given that, and the results above that showed **R** only responded well to low amplitude free energy, the results in Figure 19 are not surprising. (The many replications for these experiments take a long time, and many were conducted in parallel; otherwise, the parameter ranges for the dynamic energy experiments might have been adjusted.) The results in **R** do not demonstrate improvement over the baseline here, but they do show a dramatic shift between very low

R values for short period energy cycles and the better (but still low) *R* values for longer

periods, suggesting a strong interaction effect or perhaps a type of phase transition.
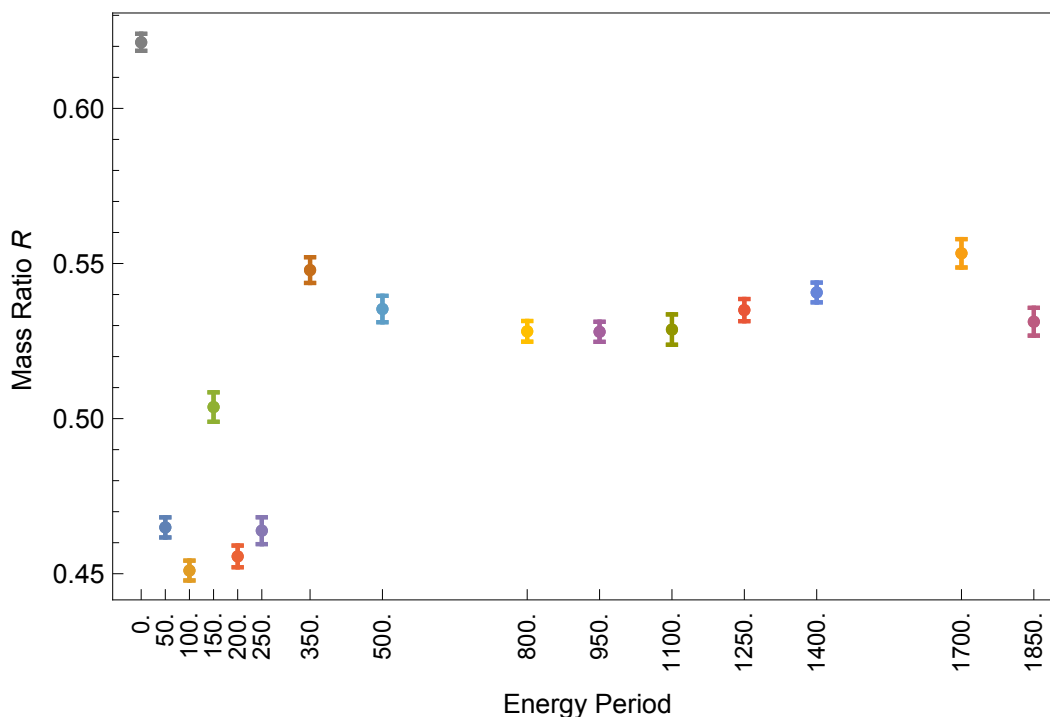
Figure 19. Sensitivity of *R* to variation in period of dynamic free energy.
Each bar represents a mean and standard error for a set of simulations. The bar at
period=0 marks the baseline value for a system with a static energy level. (Colors here
are only for differentiation.)

In Figure 20, the results show a definite positive finding for *Λ* for short periods in

the energy cycle, relative to the non-dynamic baseline value (at period 0). These

improved values for concentration slope occur for exactly the periods where the mass

ratio was lowest. In that range, for both *R* and *Λ*, periods of 50, 100, and 200 have the

most extreme values, while the period of 150 in the middle has a pronounced, less

extreme response. These oddities suggest an interaction between period and other time-

based parameters of similar scale, which could be worthy of future investigation.
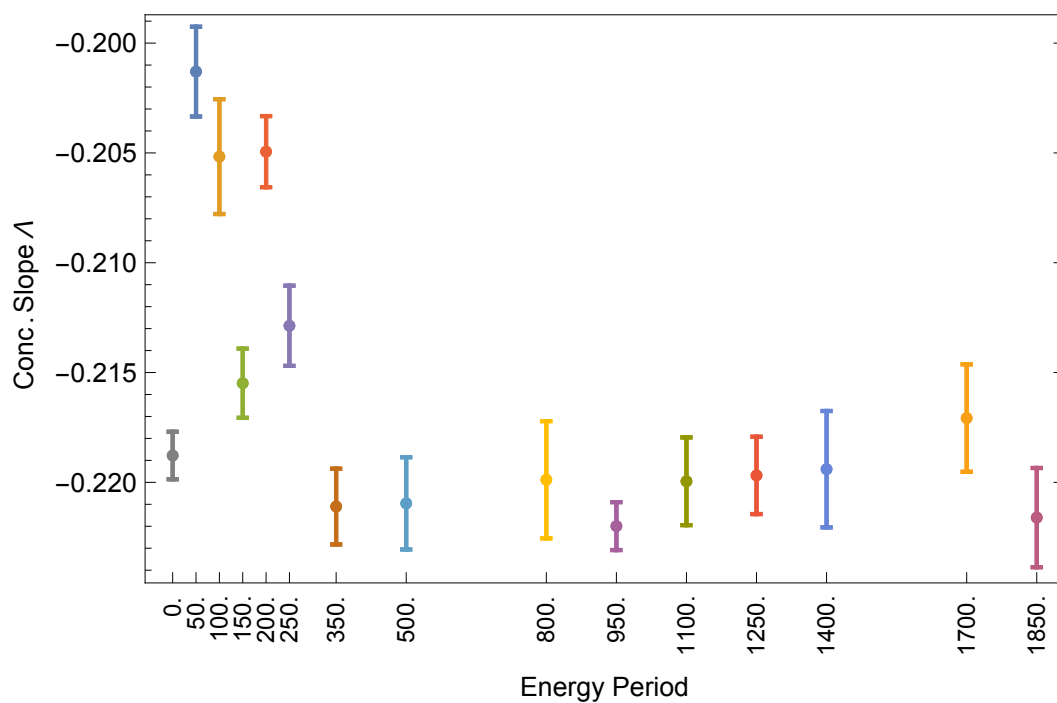
Figure 20. Sensitivity of Λ to variation in period of dynamic free energy.
Each bar represents a mean and standard error for a set of simulations. The bar at
period=0 marks the baseline value for a system with a static energy level. (Colors here
are only for differentiation.)

The plot of combined results in Figure 21 reveals a little more about the

interaction between **R, Λ,** and the energy cycle period, showing a tight cluster of

responses for the longer periods, a small cluster of short periods that behave very

differently (with improved concentrations slopes and lower mass ratio), and two results

(periods 150 and 250) that seem to bridge the other two clusters.
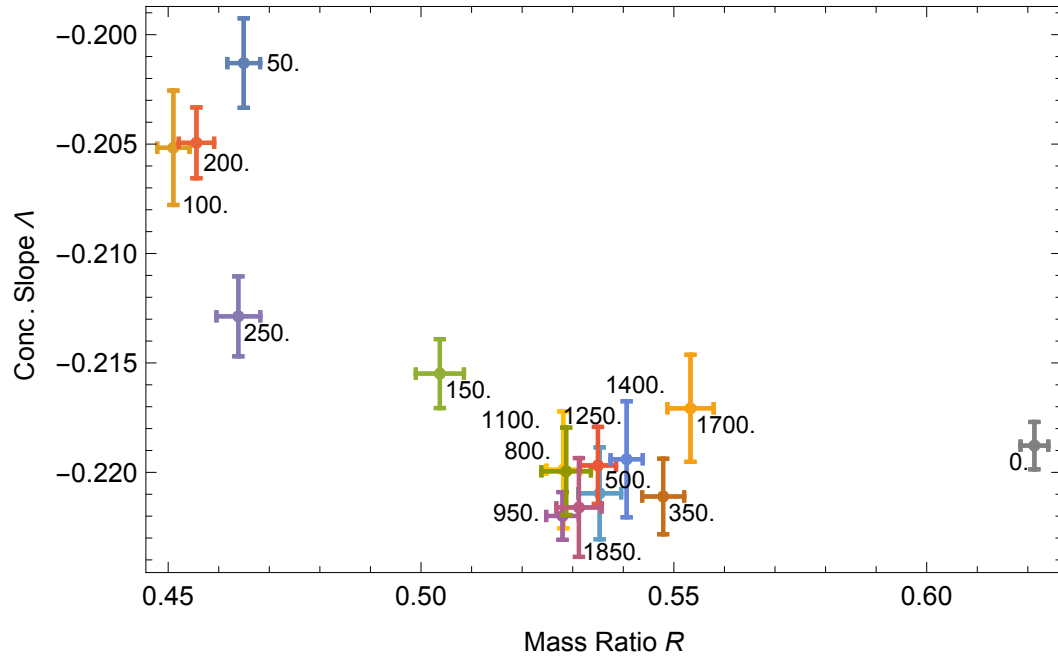
Figure 21. Sensitivity of *R* and *Λ* to variation in period of dynamic free energy.
Each crossbar represents means and standard errors for one set of simulations, labeled
with energy period value. The crossbar at period=0 marks the baseline value for a system
with a static energy level. (Colors here are only for differentiation.)

Result 2.3. Dynamic Mass Flux, Varying in Amplitude

The results in dynamic mass flux, both in amplitude and in period, show more

consistent positive findings than the variations in free energy did above. Of course, this

may be due to chance in the selection of parameters to explore, since the goal was to find

evidence of existence for the ARN-supporting dynamics, rather than surveying a wide

breadth of values or gaining a deep understanding of the interactions. Even so, the

following results for dynamic mass flux provide promising suggestions for future

research directions.
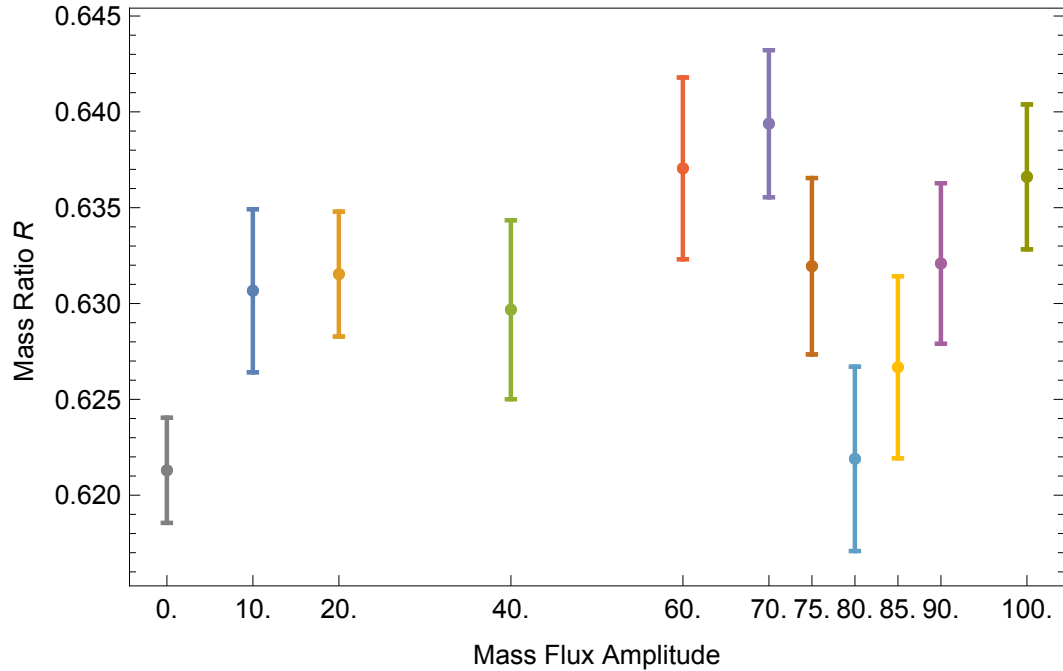
Figure 22. Sensitivity of *R* to varied amplitude of dynamic mass flux.
Each bar represents a mean and standard error for a set of simulations. The bar at
amplitude=0 marks the baseline value for a system with a static mass flux. (Colors here
are only for differentiation.)

The overall results for variations in amplitude of mass flux are presented in Figure

22 and Figure 23, for *R* and Λ respectively. Considering both figures, the main positive

finding seems to be that each of the tested variations in dynamic mass flux result either in

improvement or no change, for both metrics. This result might not be universal, of

course, but it is nonetheless very promising. There also appears to be some correlation

between the responses in *R* and Λ, suggesting that a mechanism for boosting ARN

emergence doesn't necessarily require a choice between the two metrics.
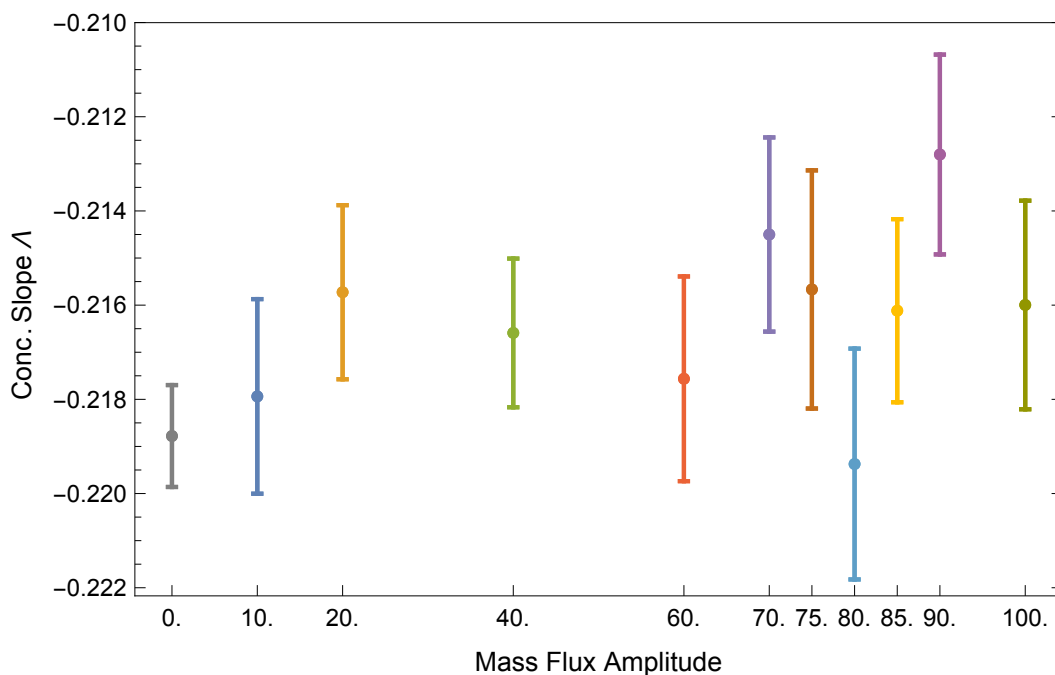
Figure 23. Sensitivity of Λ to varied amplitude of dynamic mass flux.
Each bar represents a mean and standard error for a set of simulations. The bar at
amplitude=0 marks the baseline value for a system with a static mass flux. (Colors here
are only for differentiation.)

Another finding in the results of Figure 22 and Figure 23 is the apparent pattern

around amplitude 80. For both metrics, but especially mass flux, the result for mass flux

amplitude of 80 is depressed relative to the values around it. This anomaly inspired the

addition of simulations to this experiment to explore the gaps between 80 and the

(original) neighboring values of 60 and 100. (Because the unusual value was noticed

early on, it was possible to include these simulations.) The additional results revealed a

distinct dip in the R values, which also appears in Λ but with a different shape. This

structure in the results suggests an interaction effect between mass flux amplitude and

another aspect of the model. The results also hint that further structure could be hiding

here. For instance, there may be additional complexity in R between the amplitude values

of 20, 40, and 60.
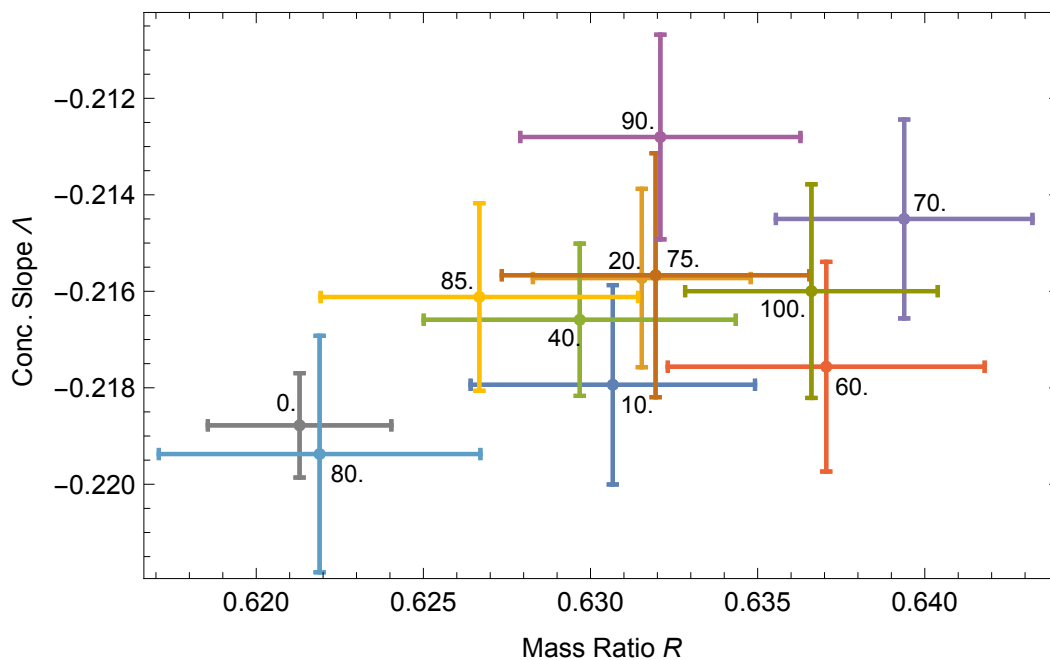
Figure 24. Sensitivity of *R* and *Λ* to varied amplitude of dynamic mass flux.
Each crossbar represents means and standard errors for one set of simulations, labeled
with amplitude values. The crossbar at amplitude=0 marks the baseline value for a system
with a static mass flux. (Colors here are only for differentiation.)

The combined plot of Figure 24 reinforces the earlier observation that dynamic

mass flux appears to have only positive or neutral effects upon the metrics of *R* and *Λ*.

While several of the results intersect with the non-dynamic baseline value in standard

error and the anomaly at amplitude of 80 is nearly identical to the baseline, it appears that

none of the tested values have a substantial negative impact on the ARN.

Figure 25 presents a detailed view of the mass ratio response to dynamic mass

flux over time. This graphic includes averaged data for 3 sets of simulations. The results

for mass flux amplitude of 10 remain consistently higher than the static baseline value,

but it isn't clear if that line is synchronized with the cycle's phase. The other two lines are

clearly synchronized with the mass flux cycle, and both seem to show the same offset to

the phase. The curves for amplitude 80 and 100 don't seem profoundly different, except

for the fact that the former stays close to the non-dynamic mean and the latter is mostly

above it, and perhaps is even on an upward trajectory.
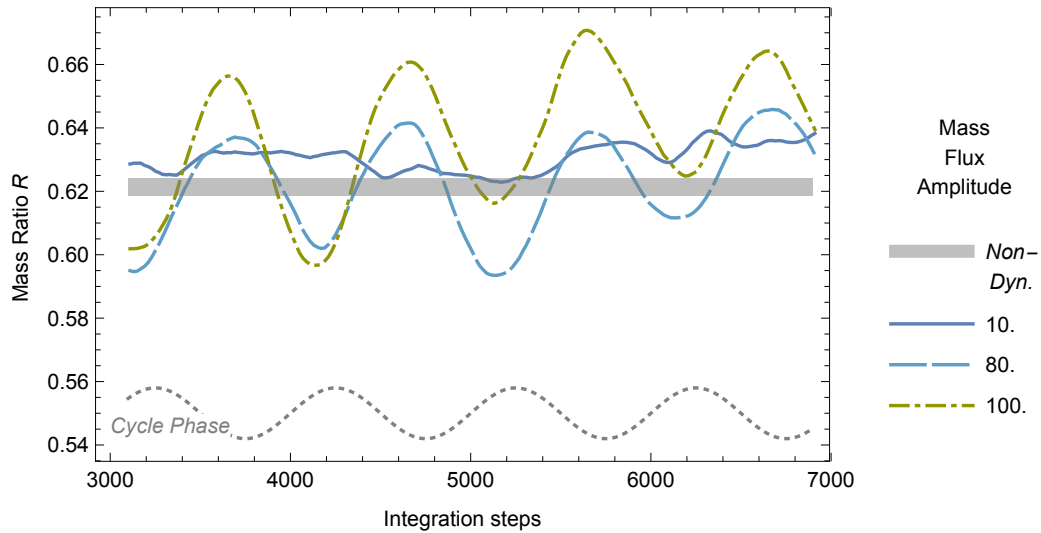
Figure 25. Time series of *R* response to varied amplitude of dynamic mass flux.
Each line is the mean of 20 simulations. The static baseline value and energy phase are
included for reference. The thickness of the non-dynamic bar shows ± standard error.

Result 2.4. Dynamic Mass Flux, Varying in Period Length

The final set of results presents the response of the DPD model to variations in the

period of dynamic mass flux. As described for the variations in mass flux amplitude, the

positive findings in these results are obvious and promising. Like those earlier results, it

appears that variations in the period of dynamic mass flux either have no effect or, more

commonly, provide a boost to both ARN metrics. Figure 26 and Figure 27 display the

simulation results in terms of R and Λ, respectively. In R, it appears that only a single

value for period (1400) caused a neutral response, with all others being positive. For Λ, a

few period lengths give results approximately equal to the baseline, and one comes close

to being worse (1550), but half of the tested periods yield improvements. A few of the

results in these figures hint at patterns like the sharp decrease seen in Figure 22 (at amplitude 80), but there is no example here of the same clarity. The odd R response near the period of 1250 is intriguing, for instance, and will likely be a target for future work, similar to the results behind Figure 22.
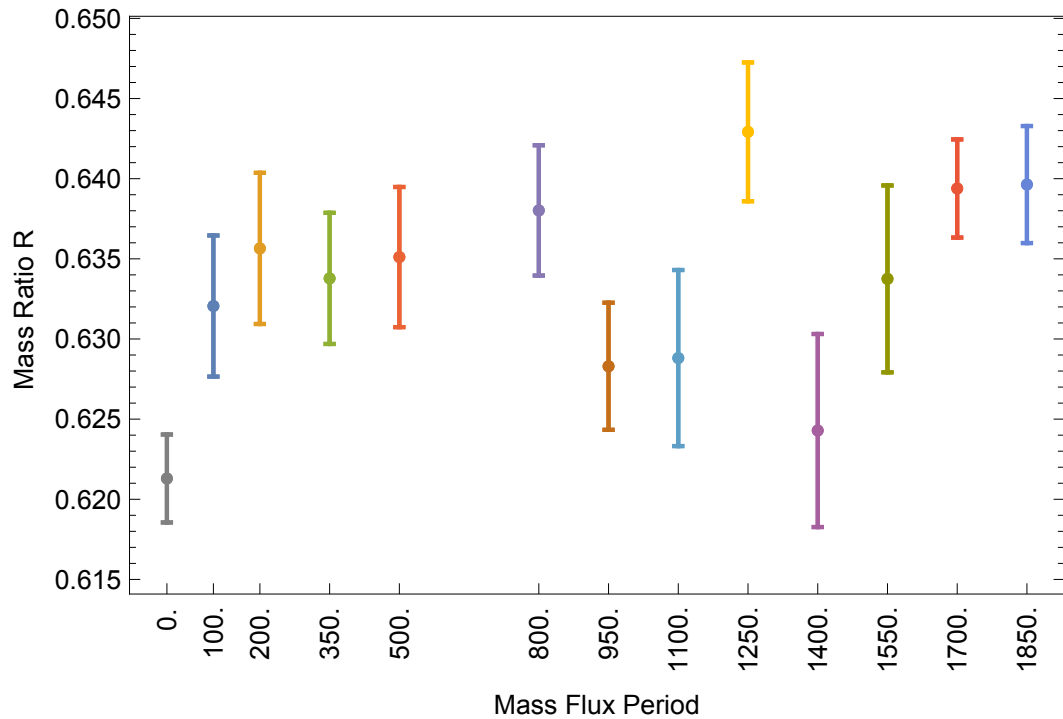
Figure 26. Sensitivity of R to variation in period of dynamic mass flux.
Each bar represents a mean and standard error for a set of simulations. The bar at period=0 marks the baseline value for a system with a static mass flux. (Colors here are only for differentiation.)
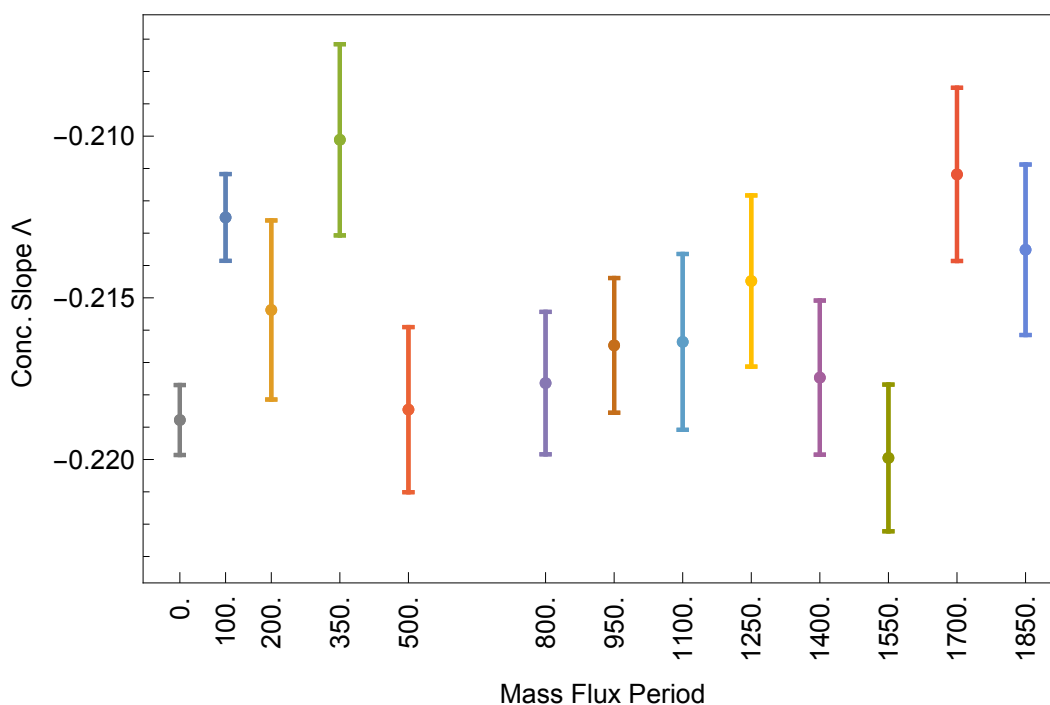
Figure 27. Sensitivity of Λ to variation in period of dynamic mass flux.
Each bar represents a mean and standard error for a set of simulations. The bar at
period=0 marks the baseline value for a system with a static mass flux. (Colors here are
only for differentiation.)

The combined results in Figure 28 are again similar to the mass flux amplitude

simulations (Figure 24). As with the earlier graphic, this plot provides clear evidence that

the effect of dynamic mass flux is predominantly positive, with essentially all of the

dynamic system values clustered together, above and to the right of the baseline value at

period 0. These results are further evidence that parameter variation can result in a boost

to both the mass ratio and the concentration slope of an ARN.
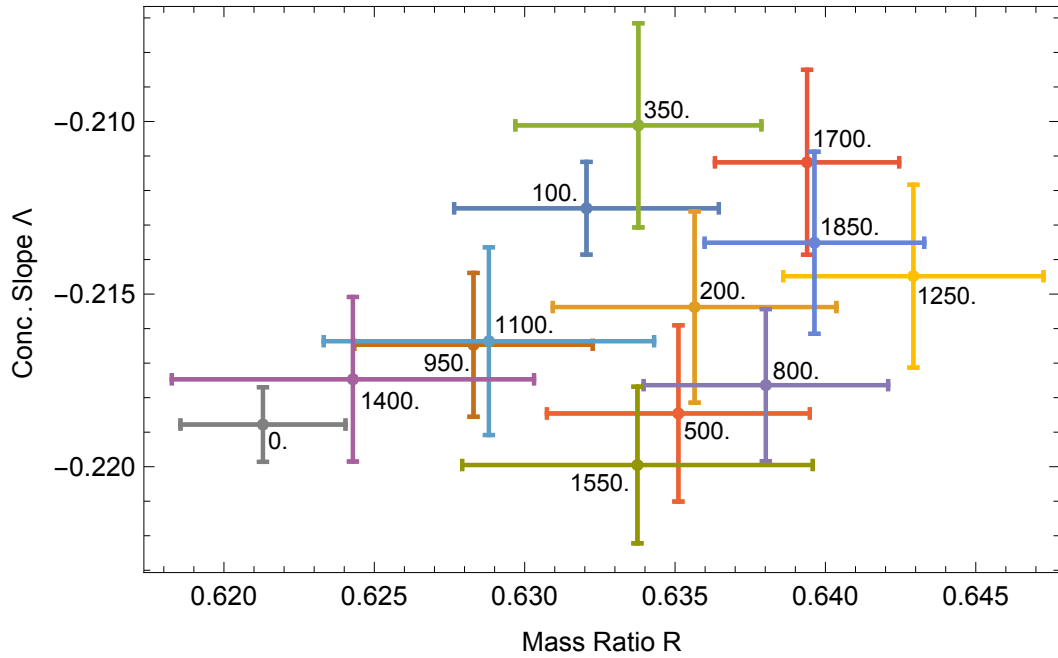
Figure 28. Sensitivity of *R* and *Λ* to variation in period of dynamic mass flux.
Each crossbar represents means and standard errors for one set of simulations, labeled
with period values. The crossbar at period=0 marks the baseline value for a system with a
static mass flux. (Colors here are only for differentiation.)

5. DISCUSSION

The most important finding of this dissertation is the clear, positive answer to the primary objective: yes, dynamically structured (specifically, cyclic) interactions with the environment can promote the emergence of ARNs. Although much investigation remains to be done into these interactions, the findings present multiple demonstrations of the existence of these effects, which successfully achieves the main research objective. Specifically, the findings include positive results for a wide range of dynamic mass flux parameters, to the extent that all of the parameter values attempted resulted in positive or neutral changes to both metrics. The results for dynamic free energy showed positive impact on the ARN only for a limited range of values. These results suggest multiple future research paths, including investigation into the mechanisms of the effects, and application of the findings to other models and domains.

The dissertation's other objective was to establish a comparative relationship between the DPD model and the DE models common in the literature, and this too was achieved. The DPD model was applied to four findings from the foundational work in Bagley & Farmer (1991). The results matched fairly well between the DPD model and the DE model, with very strong matches for two results, and interesting differences in the two that did not match as well.

Model Comparison

Result 1.2 (concentration profile analysis) and Result 1.4 (sensitivity of mass distribution to mass flux) showed strong agreement between the two models. In Results 1.2, the two figures showed nearly identical directionality of variation across polymer

concentrations, so that both ARN profile plots turned up or down at the same lengths. This seems to show that both models are implementing the network of relations in the same way, and that the chemical kinetics in both causes reasonably similar simulations. The comparison of polymer concentrations over time (Result 1.1) matched well overall, though the dissertation's DPD model results showed substantially more high-frequency variation in concentration. However, the difference between the two results was expected and even desirable, and served to demonstrate how the two model paradigms differ.

The remaining comparison, Result 1.3 (the sensitivity of Λ to mass flux) had the weakest match between the two models. The results included an apparent anomaly in the slope metric Λ for low mass flux values, with the ARN maintaining substantially higher levels in the DPD model. In the DE model of Bagley & Farmer (1991) the ARN becomes indistinguishable from the background in the analogous simulations. Closer examination of the DPD model's results suggested that the anomaly was not necessarily problematic and may even represent an unexpected finding. Working from the hypothesis that the anomaly was caused by the spatial parameters of the DPD model, a quick set of *ad hoc* simulations was run to test the idea. While not definitive, the simulation results (Table 5) were able to explain at least half of the anomaly in the low-mass-flux behavior. These simulations have only been carried out for a single value of mass flux, however, so further testing is required to confirm the hypothesis and understand the behavior.

Multiple possible explanations for the discrepancy were investigated without finding a definite cause. The leading hypothesis is that the difference is due (at least in part) to the geometry used in that DPD model, which was essentially 2-dimensional and

very different from the well-mixed conception of a DE model. A small number of *ad hoc*

tests were carried out, transitioning from the 2-D geometry to a cube in the zone of low

mass flux, where Λ was unexpectedly elevated. The results showed a substantial and

progressive reduction in Λ, reducing the discrepancy by about half. These tests were not

definitive, and further experiments will be needed across the range of mass flux values,

testing both for "3-D-ness" and for total system size (or density). Still, the *ad hoc* results

are promising for the suggestions of a possible explanation, and more so for the hint that

a 2-D geometry might have qualitatively different properties than 3-D geometries. This is

especially relevant to certain ALife/OOL models, such as those set within narrow mineral

cracks or thin films on surfaces. If this discrepancy turned into a finding that ARNs can

emerge from non-driven systems, it could be a very interesting surprise.

Environmental Dynamics

Each of the four experiments demonstrated findings that give a positive answer to

the dissertation's main objective. The experiments with mass flux dynamics showed good

results in both metrics that were considered, across many parameter values. The results

for energy dynamics were not as consistently positive, but still included positive findings.

The negative effects were expected to some extent, since increasing the energy in the

system would increase the rates of all reaction, catalyzed or note, much like an increase

in temperature. When all reaction rates increase, the relative advantage of catalyzed

reactions is weakened. Additional findings of interest in the results include apparent

interaction effects between some of the environmental dynamics and other aspects of the

reaction network. These appear as nonlinear changes in the result metrics, and they suggest the possibility of resonance between the multiple time-based aspects of the model.

The findings for mass flux dynamics showed very good results, with neutral or positive changes in both metrics over the entire range of parameter values. It is of course possible that this result is limited to the range that was tested, and the consistency of the results is due to chance. Nonetheless, the effect does exist in this range, proving that it is at least possible for a dynamic environmental interaction to be not only beneficial, but completely non-harmful within some domain. It also appears that the sensitivity of the two metrics to the parameters is not monotonic, and that some interaction causes the metrics to vary up and down over time. It appears possible that this interaction involves the time-based parameters of the model. Mass flux is itself temporal, even without the dynamic variation, and the ARN's potential to take advantage of it should depend in part on reaction rates, catalyst unbinding time, particle movement rates, and any other time-based behaviors in the system. The interaction of these time-based functions is perhaps the next step in the research that would be most relevant to extending and understanding the dissertation's results.

The findings for environmental dynamics of free energy in the system were not as uniformly positive as those for dynamic mass flux, but the positive results are still interesting and are proof of existence. In general, the effect of increased energy (in the form of temperature) on chemical reactions is to increase their rate, which should reduce the relative effect of catalysis. This is what the findings show for high-amplitude dynamic energy. The expectation for dynamic energy was that positive findings might arise from

parameters that describe conditions near the edge of emergence, where the potential

benefits of exploration could outweigh any ill effects from increased energy. This may be

what happened in Figure 14, for instance, for energy amplitudes between 1 and 2. The

positive results for varying amplitude occur at low levels, where they show a small but

consistent trend of improvement in R, and an increase in Λ over a wider range of values,

with a possible nonlinear interaction effect in the middle of that range. The findings for

variations in the period of dynamic energy show a different type of nonlinear behavior.

For a narrow range of values on the low end of those tested, the Λ metric sees a large

increase while R experiences a large decrease. Even within this narrow band, there

appears to be bump in the middle of the five data points, moving in the opposite direction

of the others. Beyond that range, in the longer periods, both metrics experience smaller

changes relative to the earlier variations. Within the longer periods, there may also be

interaction or resonance effects, but at a much smaller scale than in the short periods.

Relevance of Results

The findings of this dissertation, though limited in scope and not yet fully

explained, also have the potential to be generalized and applied to many domains.

Certainly, when the results are verified and better explored, they could be applicable to a

number of models within the field of ALife and the study of the Origin of Life. Several

theories for OOL seem ready for direct application of the findings, such as those with

unusual geometric features or regular exposure to external dynamics. Among these

settings are mineral surfaces (Koonin 2007; Orgel 1998), in subterranean minerals

(Koonin & Martin 2005; Wächtershäuser 2006), on the ocean surface (Dobson et al. 2000), and in deep-sea thermal vents (Corliss 1990; Russell & Hall 1997). Other models that feature polymer reactions may also find use for these results, including the many variations of The RNA World (Higgs & Lehman 2015). It is also possible that the findings could be useful to models with very different types of reactions, such as models of protocell formation (Fellermann et al. 2007; Solé 2009) and those in The Lipid World (Segré et al. 2001).

Another area of relevance might be in the area of ecosystem research, and the more general parallels between living systems at different scales, from ecosystems to animal groups to individual organisms to single cells. Ecosystem research has been one of the most active areas for studying energy and nutrient cycling, system-environment interdependency, and sources of variety, and it is a field with a strong component of computational simulation (DeAngelis 1980; Fletcher et al. 1997, 1998; Kaufmann 1993; Patten 1985). The parallels between living systems of various scales are somewhat widely known, at least as metaphors, but comparisons of experimental results are more rare. Such a line of research might one day validate the parallels between the real-world systems. Well before that, it could connect two very active areas of simulation research, perhaps enabling each to leverage the efforts of the other.

Limitations, and Issues of Validity

The present research primarily sought to demonstrate proof of the existence of certain behaviors, with the additional goal of providing a hint at the generality or frequency of those behaviors, at least by showing that the findings were not a "one-off"

or "lucky break." This was achieved, to an extent, by demonstrating multiple scenarios where the behaviors of interest emerge, and by showing that the behaviors can take on different, potentially interesting forms. However, this research also features uncertainty at nearly every level, and it is quite likely that some of the specifics of the findings will require revision or even discarding once validity is pursued more vigorously. For instance, the simulations were performed with software that had been recently acquired and modified, and little outside assistance was used in the process. Though the "replication" results came out quite well, the simulation systems differed so substantially that their behaviors may inhabit entirely different parts of state space, despite the apparent similarity in outputs.

Beyond the parameter translation issues, the space of possible system parameters is quite vast, and caution should be taken in drawing conclusions about the generality or frequency of the results. While it appeared to be relatively simple to find parameters that gave reasonable results, it could still be related to the limited number of experiments. In addition, researcher bias is always present, though much effort was undertaken to avoid bias in this process.

The use of custom software—built from source code—gave the research an impressive amount of flexibility, from simulation features to debugging. The unfortunate side effect of this is that the ownership and attribution of the code is still uncertain, as is the potential for posting the code online or otherwise sharing it. In order to be able to revisit the code at a future date, the source files and their surrounding project details will need to be retained and maintained. Operating systems, compilers, and graphics libraries

all might change, and yet the desire would be to keep this code in a format that is available to this researcher and to others for years to come.

Future Directions

These results have created many opportunities for future research. The most appropriate next steps might be those that provide a better understanding of the mechanisms underlying the dissertation's results, both for the main results of increased ARN emergence, and also for the ancillary findings, such as the apparent interaction effects between the environmental dynamics and the system's structure, and the not-yet-explained behavior in the DPD model at low mass flux. These secondary findings will first require additional simulations and analysis to establish their existence, given the smaller number of runs that went into those experiments. Even the main results would benefit from additional simulations, and this may be the best way to provide confidence in the findings. Doing more runs for existing parameter configurations could give a better understanding of consistency while reducing error, and runs with new parameters could allow for the confirmation of perceived trends and suspected interaction effects. Results such as those in Figure 14 would be improved by narrowing the parameter range to the region nearest the "baseline" value, in the small zone where the energy cycle appears beneficial and the metrics show an inflection. Other results, such as Figure 20 and Figure 23, could be improved by filling in gaps to understand nonlinearities, or by extending the parameter range further from the baseline, to determine if the interesting behavior extends into other regions.

Other directions for future research might start by identifying research partners for collaboration, particularly those who are more experienced with some model that could be used to verify, replicate, or extend the present findings. This could be very productive, whether it involved the same DPD model or a different implementation, or some other model for ARN research, or even a model in another domain. For instance, it would be interesting to apply the environmental dynamics to a DPD or ABM model whose behavior is better understood, and preferable well validated, regardless of its connection to the ALife field. This could help to begin determining how well the results on cyclic dynamics generalize. Working with a model that is better established might also make it easier to gain a deeper understanding of results such as the sensitivity analyses, if typical behavior is more well known, and the inner workings of the model are better understood.

This dissertation was designed to build upon and replicate the early work from a line of research that is now more than 25 years old. As described in the LITERATURE REVIEW chapter, this body of work contains numerous other results, and these could potentially be translated to the DPD model for further comparison. The dynamic interaction effects could also be translated to other existing models of autocatalysis, to confirm in which specific cases the differences in the modeling paradigms give different results. It may be that a DE model with stochastic elements and imposed cycling is able to exhibit some of the results observed in the present work, but not others, and the continuance of this comparison ought to be part of any future work.

There are a number of future directions for ARN research that will likely be much more difficult for models other than the type in the present research. Early in the course

of the present work, a wider literature review and unconstrained long-term vision together suggested many such research opportunities. These potential paths include multiple ways to connect the present work to the Origin of Life theories that have been mentioned above. The theories that best map to environmental dynamics—tide pools, deep-sea vents, geysers, mineral surfaces—could all be simulated with extensions of the DPD model, to various levels of fidelity. Many interesting variations of these scenarios could be simulated through a model with multiple chemical systems in parallel, separated by spatial structures that limit interactions in different ways. These interactions could model behaviors such as competition for resources, cooperation, adaptation to niches, and so on. With the spatial capabilities of the DPD model, these proto-life systems could be designed to simulate conditions such as: cracks or pores in minerals, deep-sea vents, puddles subjected to evaporation and condensation, or foam on the ocean surface. These are all potentially complex systems, and they will likely be a challenge to simulate and understand. The present work begins the task of making such work approachable, in the results that it has provided and in the methods used, from the automation of parallel execution, to the streamlining of the model code, to establishing structures for deriving analysis efficiently.

REFERENCES

Anet, F.A. "The place of metabolism in the origin of life." *Current opinion in chemical biology* (2004) 8 (6):654–659.

Ashby, W.R. *An Introduction to Cybernetics* (1956) London: Chapman & Hall.

Bagley, R.J. & J.D. Farmer. "Spontaneous emergence of a metabolism." *Artificial Life II* (1991) 10: 93–140.

Bagley, R.J., J.D. Farmer & W. Fontana. "Evolution of a metabolism." *Artificial Life II* (1991) 10: 141–158.

Benkö, G., F. Centler, P. Dittrich, C. Flamm, B.M.R. Stadler & P. Stadler. "A topological approach to chemical organizations." *Artificial Life* (2009) 15 (1): 71–88.

Bersini, H. "Software Replica of Minimal Living Processes." *Origins of Life and Evolution of Biospheres* (2010) 40 (2): 121–130.

Boerlijst, M.C., & P. Hogeweg. "Spiral wave structure in pre-biotic evolution: hypercycles stable against parasites." *Physica D: Nonlinear Phenomena* (1991) 48: 17–28.

Borshchev, A. & A. Filippov. "From system dynamics and discrete events to practical agent based modeling: reasons, techniques, tools." *Proceedings of the 22nd International Conference of the System Dynamics Society* (2004): 25–29.

Branciamore, S., E. Gallori, E. Szathmáry & T. Czárán. "The Origin of Life: Chemical Evolution of a Metabolic System in a Mineral Honeycomb?" *Journal of Molecular Evolution* (2009) 69 (5): 458–469.

Corliss, J.B. "Hot springs and the origin of life." *Nature* (1990) 347: 624.

Cromwell, P.R., E. Beltrami & M. Rampichini. "The Borromean Rings." *Mathematical Intelligencer* (1998) 20 (1): 53–62.

DeAngelis, D.L. "Energy flow, nutrient cycling, and ecosystem resilience." *Ecology* (1980): 764–771.

Dekking, F.M., C. Kraaikamp, H.P. Lopuhaä, & L.E. Meester. *A modern introduction to probability and statistics: understanding why and how* (2005) Springer-Verlag, London.

Dobson, C.M., G.B. Ellison, A.F. Tuck & V. Vaida. "Atmospheric aerosols as prebiotic chemical reactors." *Proceedings of the National Academy of Sciences USA* (2000) 97: 11864–11868.

Emery, F.E. & E.L. Trist. "The Causal Texture of Organizational Environments." *Human Relations* (1965) 18: 21–32.

Epstein, J. "Agent-based computational models and generative social science." *Complexity* (1999) 4 (5): 41–60.

Farmer, J.D., S.A. Kauffman & N.H. Packard. "Autocatalytic replication of polymers." *Physica D: Nonlinear Phenomena* (1986) 22 (1–3): 50–67.

Fellermann, H., S. Rasmussen, H.J. Ziock & R.V. Solé. "Life Cycle of a Minimal Protocell—A Dissipative Particle Dynamics Study." *Artificial Life* (2007) 13: 319–345.

Fernando, C., G. von Kiedrowski, & E. Szathmáry. "A Stochastic Model of Nonenzymatic Nucleic Acid Replication: "Elongators" Sequester Replicators." *Journal of Molecular Evolution* (2007) 64 (5): 572–585.

Fishkis, M. "Steps Towards the Formation of A Protocell: The Possible Role of Short Peptides." *Origins of Life and Evolution of Biospheres* (2007) 37 (6): 537–553.

Flekkøy, E. & P. Coveney. "From molecular dynamics to dissipative particle dynamics." *Physical Review Letters* (1999) 83 (9): 1775–1778.

Fletcher, J., M.A. Bedau & M. Zwick. "Effect of environmental structure on evolutionary adaptation." *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life* (1998): 189–198.

Fletcher, J. & M. Zwick. "Strong altruism can evolve in randomly formed groups." *Journal of Theoretical Biology* (2004) 228 (3): 303–313.

Fletcher, J. & M. Zwick. "The evolution of altruism: Game theory in multilevel selection and inclusive fitness." *Journal of Theoretical Biology* (2007) 245 (1): 26–36.

Fletcher, J., M. Zwick, & M.A. Bedau "Dependence of Adaptability on Environmental Structure in a Simple Evolutionary Model." *Adaptive Behavior* (1997) 4: 275–307.

Fraser, S. & C. Reidys. "Evolution of random catalytic networks." *ECAL 1997* (1997).

Freeman, S. & J.C. Herron. *Evolutionary Analysis* (2002) Prentice Hall, NJ.

Gánti, T. *The Principles of Life* (2003) Oxford University Press, UK.

Gazzola, G., A. Buchanan, N.H. Packard & M.A. Bedau. "Catalysis by self-assembled structures in emergent reaction networks." *Advances in Artificial Life* (2007): 876–885. Springer Berlin Heidelberg.

Gilbert, W. "The RNA World." *Nature* (1986) 319: 618.

Gordon-Smith, C. "SimSoup: Artificial Chemistry Meets Pauling." *ECAL 2009* (2009).

Hardin, G. "The Tragedy of the Commons." *Science* (1968) 162: 1243–1248.

Hart, S. "An interview with Robert Aumann." *Macroeconomic Dynamics* (2005) 9 (05): 683–740.

Hazen, R.M. *Genesis: the scientific quest for life's origin*. (2005) Joseph Henry Press, Washington D.C.

Higgs, P.G. & N. Lehman. "The RNA World: molecular cooperation at the origins of life." *Nature Reviews | Genetics* (2015) 16: 7–17.

Holland, J. *Adaptation in Natural and Artificial Systems* (1992) The MIT Press, Reprint edition.

Hordijk, W. & M. Steel. "Detecting autocatalytic, self-sustaining sets in chemical reaction systems." *Journal of Theoretical Biology* (2004) 227: 451–461.

Hordijk, W., N. Vaidya, & N. Lehma. "Serial transfer can aid the evolution of autocatalytic sets." *Journal of Systems Chemistry* (2014) 5 (1): 1–8.

Housecroft, C.E. & E.C. Constable. *Chemistry: An Introduction to Organic, Inorganic and Physical Chemistry* (4th ed.) (2010) Prentice Hall, Harlow, England.

Kauffman, S.A. "Autocatalytic Replication of Polymers." *Journal of Theoretical Biology* (1986) 119: 1–24.

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

Koonin, E.V. & W. Martin. "On the origin of genomes and cells within inorganic compartments." *TRENDS in Genetics* (2005) 21 (12): 647–654.

Koonin E.V. "An RNA-making reactor for the origin of life." *Proceedings of the National Academy of Sciences* (2007) 104: 9105–9106.

Langton, C.G. "Self-reproduction in cellular automata." *Physica D: Nonlinear Phenomena* (1984) 10 (1): 135–144.

MacArthur, R. & E.O. Wilson. *The Theory of Island Biogeography* (1967) Princeton University Press (2001 reprint).

McCaskill, J., N.H. Packard, S. Rasmussen & M.A. Bedau. "Evolutionary self-organization in complex fluids." *Philosophical Transactions of the Royal Society B: Biological Sciences* (2007) 362 (1486): 1763–1779.

McMullin, B. "Some remarks on autocatalysis and autopoiesis." *Proceedings of the Workshop on Closure: Emergent Organizations and their Dynamics* (1999) University of Ghent, Belgium.

Miller, S.L. "A production of amino acids under possible primitive earth conditions." *Science* (1953) 117 (3046): 528–529.

Miller, S.L., & H.C. Urey. "Origin of Life." *Science* (1959) 130 (3389): 1622–1624.

Minar, N., R. Burkhart, C. Langton & M. Askenazi. "The Swarm simulation system: a toolkit for building multi-agent simulations." *Working Paper 96-06-042* (1996) Santa Fe Institute, Santa Fe.

Morowitz, H.J., B. Heinz & D.W. Deamer. "The chemical logic of a minimum protocell." *Origins of Life and Evolution of the Biosphere* (1988) 18: 281–287.

Orgel, L.E. "Polymerization on the rocks: Theoretical Introduction." *Origins of Life and Evolution of the Biosphere* (1998) 28: 211–218.

Orgel, L.E. "The Implausibility of Metabolic Cycles on the Prebiotic Earth." *PLoS Biology* (2008) 6 (1): 5–13.

Patten, B.C. "Energy cycling in the ecosystem." *Ecological Modelling* (1985) 28 (1): 1–71.

Pepper, J.W. & B.B. Smuts. "Assortment through environmental feedback." *The American Naturalist* (2002) 160: 205–213.

Rasmussen, S., M.A. Bedau, L. Chen & D. Deamer. *Protocells: Bridging Nonliving and Living Matter*. (2008): 776.

Russell, M.J. & A.J. Hall. "The emergence of life from iron monosulphide bubbles at a submarine hydrothermal redox and pH front." *Journal of the Geological Society* (1997) 154: 377–402.

Schieritz, N. & P. Milling. "Modeling the Forest or Modeling the Trees." *Proceedings of the 21st International Conference of the System Dynamics Society* (2003).

Scholl, H. "Agent-based and system dynamics modeling: a call for cross study and joint research." *Structure* (2001) 2: 120.

Segbroeck, S.V., A. Nowé & T. Lanaerts. "Stochastic Simulation of the Chemoton." *Artificial Life* (2009) 15 (2): 213–226.

Segré, D. & D. Lancet. "Composing life." *EMBO Reports* (2000) 1 (3): 217–222.

Segré, D., D. Lancet, O. Kedem, & Y. Pilpel. "Graded Autocatalysis Replication Domain (GARD): kinetic analysis of self-replication in mutually catalytic sets." *Origins of Life and Evolution of Biospheres* (1998) 28 (4-6): 501–514.

Segré, D., D. Ben-Eli, D. Deamer & D. Lancet. "The lipid world." *Origins of Life and Evolution of Biospheres* (2001) 31 (1): 119–145.

Shenhav, B., A. Oz & D. Lancet. "Coevolution of compositional protocells and their environment." *Philosophical Transactions of the Royal Society B: Biological Sciences* (2007) 362 (1486): 1813–1819.

Solé, R.V. "Evolution and self-assembly of protocells." *International Journal of Biochemistry and Cell Biology* (2009) 41: 274–284 .

Stadler, P. & P. Schuster. "Mutation in autocatalytic reaction networks." *Journal of Mathematical Biology* (1992) 30: 597–632.

Stadler, P., W. Fontana & J. Miller. "Random catalytic reaction networks." *Physica D* (1993) 63 (3–4): 378–392.

Stassinopoulos, D., S.P. Colombano, J.D. Lohn, G.L. Haith, J. Scargle & S. Liang. "Spatial autocatalytic dynamics: An approach to modeling prebiotic evolution." (1998) Proceedings of the 1998 International Conference of Complex Systems.

Syvanen, M. "Cross-species gene transfer; implications for a new theory of evolution." *Journal of Theoretical Biology* (1985) 112 (2): 333–43.

Szathmáry, E. "The evolution of replicators." *Philosophical Transactions of the Royal Society B: Biological Sciences* (2000) 355 (1403): 1669–1676.

Szathmáry, E. "The origin of replicators and reproducers." *Philosophical Transactions of the Royal Society B: Biological Sciences* (2006) 361 (1474): 1761–1776.

Vasas, V., E. Szathmáry & M. Santos. "Lack of evolvability in self-sustaining autocatalytic networks constraints metabolism-first scenarios for the origin of life." *Proceedings of the National Academy of Sciences* (2010) 107 (4): 1470–1475.

Wächtershäuser, G. "From pre-cells to Eukarya – a tale of two lipids." *Molecular Microbiology* (2003) 47: 13–22.

Wächtershäuser, G. "From volcanic origins of chemoautotrophic life to Bacteria, Archaea and Eukarya." *Philosophical Transactions of the Royal Society B: Biological Sciences* (2006) 361 (1474): 1787–1808.

Wolfram Research, Inc. *Mathematica*, Version 10.0 (2014) Champaign, IL.

Wu, M. & P.G. Higgs. "Compositional Inheritance: Comparison of Self-assembly and Catalysis." *Origins of Life and Evolution of Biospheres* (2008) 38: 399–418.

Wu, M., & P.G. Higgs. "The origin of life is a spatially localized stochastic transition." *Biology Direct* (2012) 7 (42): 1–15.