

On the Origin of the Living State

by

Cole (Nicholas) Mathis

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved June 2018 by the
Graduate Supervisory Committee:

Sara Imari Walker, Chair
Paul CW Davies
Michael Lachmann
Ralph V Chamberlin

ARIZONA STATE UNIVERSITY

August 2018

©2018 Cole (Nicholas) Mathis

All Rights Reserved

ABSTRACT

The origin of Life on Earth is the greatest unsolved mystery in the history of science. In spite of progress in almost every scientific endeavor, we still have no clear theory, model, or framework to understand the processes that led to the emergence of life on Earth. Understanding such a processes would provide key insights into astrobiology, planetary science, geochemistry, evolutionary biology, physics, and philosophy. To date, most research on the origin of life has focused on characterizing and synthesizing the molecular building blocks of living systems. This bottom-up approach assumes that living systems are characterized by their component parts, however many of the essential features of life are system level properties which only manifest in the collective behavior of many components. In order to make progress towards solving the origin of life new modeling techniques are needed. In this dissertation I review historical approaches to modeling the origin of life. I proceed to elaborate on new approaches to understanding biology that are derived from statistical physics and prioritize the collective properties of living systems rather than the component parts. In order to study these collective properties of living systems, I develop computational models of chemical systems. Using these computational models I characterize several system level processes which have important implications for understanding the origin of life on Earth. First, I investigate a model of molecular replicators and demonstrate the existence of a phase transition which occurs dynamically in replicating systems. I characterize the properties of the phase transition and argue that living systems can be understood as a non-equilibrium state of matter with unique dynamical properties. Then I develop a model of molecular assembly based on a ribonucleic acid (RNA) system, which has been characterized in laboratory experiments. Using this model I demonstrate how the energetic properties of hydrogen bonding dictate the population

level dynamics of that RNA system. Finally I return to a model of replication in which replicators are strongly coupled to their environment. I demonstrate that this dynamic coupling results in qualitatively different evolutionary dynamics than those expected in static environments. A key difference is that when environmental coupling is included, evolutionary processes do not select a single replicating species but rather a dynamically stable community which consists of many species. Finally, I conclude with a discussion of how these computational models can inform future research on the origins of life.

DEDICATION

To my brothers, who made me curious about everything.

ACKNOWLEDGMENTS

It would be impossible to enumerate everyone who contributed to this dissertation but I'll try. First and foremost I'd like to thank my parents for their endless love, support and encouragement. Thanks to my high school physics teacher Mr Vining who got me started down this path by showing me that physics is cool. Taylor Womble-Dahl, Christian Hansen, Tim Gushe, and especially Stephanie Lauber for the late nights in PAC which got me through undergrad. Professor Joe Barranco for demonstrating that physics professors are bad-asses, and his guidance in applying to ASU. Professor Mithat Unsal for providing me with challenging theoretical problems and support in applying to graduate school. The Santa Fe Institute CSSS class of 2014 and the Winter Workshops in Madrid and Utrecht for bringing me into the complex systems community. The attendees AbGradCon 2015 and 2016 for bringing me into the Astrobiology community. Professor Everett Shock for taking me to Yellowstone National Park in 2016 and showing me how geobiochemists think. Thanks to all the past and present members of Elife. Special thanks to Alyssa Adams for slogging through the first year grad courses together, Harrison Smith for all the mornings at Royal and all the advice about plots, Tucker Ely for every beer and every point we disagreed on, and Doug Moore for the continuous programming support and for always playing the devil's advocate. I will always be indebted to Sara Walker for introducing me to the exciting and challenging field of the origin of life, for being the best advisor anyone has ever had and for being a limitless source of inspiration, advice and encouragement. This dissertation wouldn't have been possible without her. Of course I never would have made it through graduate school without the support of my future wife, Andrea Tidrow. Her love, patience and advice got me through the most trying aspects of life, without her I wouldn't have been able to realize my dreams.

TABLE OF CONTENTS

	Page
LIST OF TABLES	viii
LIST OF FIGURES	ix
CHAPTER	
1 ORIGIN OF LIFE BACKGROUND AND INTRODUCTION	1
1.1 The Greatest Mystery	1
1.2 Epistemology	2
1.3 Perspectives on the Origin of Life	5
1.3.1 Replication First	5
1.3.2 Metabolism First and Life as a Non-equilibrium system	15
1.4 Dissertation Outline	22
2 THE MEANING OF THE <i>LIVING STATE</i>	24
2.1 Introduction	24
2.2 A Brief Synopsis of Statistical Physics	27
2.3 Life as a state of matter	30
2.4 Conclusion	34
3 THE EMERGENCE OF LIFE AS A FIRST ORDER PHASE TRAN- SITION	35
3.1 Abstract	35
3.2 Introduction	36
3.3 Methods	38
3.3.1 Model Description and Motivation	38
3.3.2 Two Fitness Landscapes: Static and Dynamic	40
3.3.2.1 Static Fitness	41

CHAPTER	Page
3.3.2.2 Dynamic Fitness	42
3.3.3 Mutual Information as a Measure of the Transition from Non-life to Life	43
3.4 Results	43
3.4.1 Non-life and Life	43
3.4.2 A First Order Phase Transition from Non-life to Life	45
3.4.3 The Timescale for Life's Emergence	49
3.5 Discussion	52
3.6 Supplementary information	56
4 PREBIOTIC RNA NETWORK FORMATION: A TAXONOMY OF MOLECULAR COOPERATION	59
4.1 Abstract	59
4.2 Introduction	60
4.2.1 Nomenclature and chemistry	64
4.2.2 Classification and Taxonomy	66
4.3 Discussion	69
4.4 Materials and Methods	73
4.4.1 Kinetic Simulations	73
4.4.2 Fitting Growth Constants	74
4.4.3 Randomization Experiments	75
4.5 Conclusions	76
5 NOISY CHANNELS, ERROR CORRECTION AND THE ORIGIN OF LIFE	77
5.1 Abstract	77

CHAPTER	Page
5.2 Introduction	77
5.3 Model Description	80
5.4 Results	82
5.5 Discussion	89
6 CONCLUSION	92
NOTES	97
REFERENCES	98
APPENDIX	
A STATEMENT OF CO-AUTHOR PERMISSIONS	110

LIST OF TABLES

Table	Page
1 Autocatalytic Rate Constants for <i>Azoarcus</i> RNA Covalent Self Assembly.....	65

LIST OF FIGURES

Figure	Page
<p>1 Time Evolution of Replicators with Different Growth Laws. Parabolic (Orange), Exponential (Blue), and Hyperbolic (Green) Growth Laws Are Shown. In All Cases $A(0) = K = 1$</p>	8
<p>2 Ensemble Averaged Compositions of All Sequences with $L \leq 7$. The Distributions in the Top-Panel Characterize the Non-Life Phase, and the Bottom-Panel Characterize the Life Phase. Data Is Averaged over 100 Simulations.</p>	43
<p>3 Top: Time Series Showing the Mutual Information between Extant Replicator Composition and the Environment. The Phase Transition Is Clearly Evident in the Abrupt Shift from $\mathcal{I}(R; E) \sim 2.5$ in the Non-Life Phase to $\mathcal{I}(R; E) \sim 0$ in the Life Phase at $t \sim 5500$. Also Evident Are Several Failed Transitions, Including One that Nearly Runs to Completion before Reverting back to the Non-Life Phase near $t \sim 3000$. Bottom: The Frequency of Successful Phase Transitions Plotted against the Difference between the Distribution of Free Monomers and Replicator Composition, for an Ensemble Statistic of 256 Simulations.</p>	44
<p>4 Phase Trajectory for an Ensemble of 100 Systems Transitioning from Non-Life to Life, Moving from left to right. Axes Are the Mutual Information $\mathcal{I}(R; E)$ between Replicators and Environment (x Axis) and the Ratio of Formation (Polymerization and Replication) to Degradation Rates (y Axis).</p>	46

Figure	Page
5 Time Series for the Extant Species Population Size and Total Number of Sequences Explored by the System. Linear Fits to the Explored Species Are Shown. The Exploration Rate Is 75% Faster during the Life Phase Compared to the Non- Life Phase, and Is 2 Orders of Magnitude Larger during the Transition.	47
6 Series of Symmetry Breaking Transitions in the Selection of Fit, Homogeneous '0' and '1' Length $L = 7$ Replicators. Here, the Subscript Denotes the Number of '1' Monomers in the Sequence (<i>E.g.</i> x_0 Contains No '1's, x_1 Bins All Polymers with a Single '1' Monomer in Their Sequence, and x_7 Contains All '1's). Parameters Are $k_p = 0.0005$, $K_d = 0.9000$, and $k_r = 1.000$	49
7 Timescale for Completing the Phase Transition as a Function of Reaction Rate Constants for Replication k_r . Data from 25 Simulations Is Shown, All Data Points Are Included in the Box and Whisker Plots. The Center Line for Each Distribution Is the Median, the Boxes Contain Half the Data Points and the Bars Show the Range. Three Values of the Degradation Rate Constant k_d Are Shown, 5.0 (Blue), 1.0 (Purple), 0.5 (Green)	51
8 Timescale for Completing the Phase Transition as a Function of the Abiotic Distribution of Resources. Here, the Parameter R Is the Ratio of '1' Monomers (Which Confer Stability) to Total System Mass. Data from over 100 Simulations Is Shown, All Data Points Are Included in the Box and Whisker Plots. The Center Line for Each Distribution Is the Median, the Boxes Contain Half the Data Points and the Bars Show the Range.	52

Figure	Page
<p>9 The <i>Azoarcus</i> Ribozyme Self Assembly System. (a) The Reaction between the 148-Nt WXY RNA Fragment and the 55-Nt Z RNA Fragment. This Reaction Is Catalyzed by the Binding of the IGS of Another Catalyst (Either a Covalently-Contiguous WXYZ Molecule that Had Been Previously Assembled or a Non-Covalent WXY-Z Complex) to the Tag Sequence on the 3' End of the Substrate WXY Molecule. The Key Hydrogen Bonding Interactions Are Shown in Green (Invariant) and Red (Variable) Dotted Lines. (B) A Comparison of Selfish and Cooperative Assembly Reactions by the Use of Two Example Pairs. In Selfish Assembly, the IGS (Upper left of Each Molecule) of the Substrate and the Catalyst RNAs Match; in Cooperative Assembly They Do Not Match.</p>	62
<p>10 The 17 Possible Non-Trivial 3-Node, 3-Edge Network Topologies. Topology 2 Is the “Rock-Paper-Scissors” (RPS) Scenario. A, B, and C Denote Distinct WXY RNA Genotypes, While the Arrows Denote the Ability of One Genotype to Covalently Assemble Another. For Comparison, There Are Only Four 2-Node, 2-Edge Topologies, and We Explored These in Detail Previously [yeates2016dynamics].....</p>	63
<p>11 Taxonomy of Molecular Cooperation Using the Empirically Derived Rate Constants. The left Bar (Gray) Represents All Possible 560 Triplet Genotype Combinations. The Bars in the Center Correspond to the First Level Classifications, and the Height of the Bar Indicates Which Fraction of All Triplets Fell into that Category. The Bars on the right Show the Second Level Classifications and the Fraction of All Triplets Which Fall into that Category.</p>	69

Figure	Page
<p>12 The Effect of the Distributions of Ribozyme Self-Assembly Rates on Cooperation. Top, left: Flat Random <i>vs.</i> Real Rate Constants. The Empirically Derived Rate Constants Show an Enhanced Level of Cooperative Relative to a Completely Randomized Control. Top, right: Log Random <i>vs.</i> Real Rate Constants. The Empirically Derived Distribution of Cooperation Is Well Explained by a Logarithmically Distributed Random Distribution Bottom, left: WC Pairs <i>vs.</i> Real. The Empirically Derived Distribution of Cooperation Is Not Completely Explained by Including Two Values of Rate Constants. Bottom, right: WC with Three Wobble Pairs <i>vs.</i> Real. The Empirically Derived Rate Distribution of Cooperation Is Well Explained by Including Wobble Pairs in a Simplified Set of Rate Constants.</p>	70
<p>13 Average Abundance of Templates Changes as a Function of the Replication Rate k_r. The Time-Averaged, Total Template Abundance Is Shown for Different Values of k_r and μ Is Shown. The Value of $k_r = 5 * 10^{-5}$ Was Chosen for the following Results, and Is Shown Here with a Vertical Dashed Line.</p>	83
<p>14 Time Series of Templates for a Simulation with a Single Master Sequence Which Has a Factor of 10 Greater Replication Rate than All Other Sequences. The Colored Line Here Correspond to the Abundance Different Templates. The Abundances Are Plotted on a Logarithmic Scale. The Red Line Indicates the Abundance of the Master Sequence at Each Time Step. The Grey Line Represents the <i>Average</i> Abundance of Species Which Differ from the Master by a Single Mutation, the Pink Line Represents the <i>Average</i> Abundance of Species Which Differ from the Master by Two Mutations and So on.</p>	84

Figure	Page
15 Selection of the Master Sequence Occurs Stochastically in This Model. Top: The Mass Fraction Is the Time Averaged Ratio between the Mass in the Master Sequence against the Mass in All Templates. The Mass Fraction of the Master Sequence for 1127 Simulations Is Shown as a Function of the Mutation Rate. Points Are Colored Based on Whether or Not Selection Occurred in that Simulation (See Main Text). Bottom: The Estimated Probability of Selection Occurring Is Shown as a Function of the Mutation Rate μ (Black Line). We Find that in Contrast to Previous Models of Molecular Replication, the Error Threshold in This Model Does Not Represent a Strict Limitation. Selection of the Master Sequence Occurs with a Probability that Decreases as the Mutation Rate Increases. Selection Can Occur above the Error Threshold.	86
16 The Probability of the Master Sequence Fixating Is Shown as a Function of the Mutation Rate μ , for Two Different Models. In Orange, the Probability for a Model with a Static Mutation Rate Is Shown. In This Case the Error Threshold Is a Sharp Transition, as Expected from Quasi-Species Theory, When the Mutation Rate Is above a Critical Value the Master Sequence Cannot Be Selected. In Green, the Same Probability Is Shown a Model with a Dynamic Mutation Rate. The Model with a Static Mutation Rate Has a Sharp Error Threshold, While the Model with a Dynamic Mutation Rate Has a Stochastic Error Threshold.	87

17 Functional Templates Which Catalyze Changes in the Environment Can Enable Selection in the Face of High Mutation Rates. The Probability of Selection Is Shown as a Function of the Catalytic Rate Constant Associated with Functional Templates. Left: The Master Sequence (<i>BBBBBBBB</i>) Is Also a Catalyst Which Converts <i>A</i> Monomers into <i>B</i> Monomers. As the Catalytic Rate Constant f_c , Is Increased the Probability of Selection Increases. For Higher Mutation Rates, Greater Catalytic Function Is Needed to Achieve Selection. Right: The Master Sequence Can Be Different. Here the Probability of Selection Is Shown the Cases Where the Catalyst Is the Same as the Master Sequence, as Well as When the Catalyst Is a Mutant of the Master Sequence (See Text for Particular Sequences).	89
---	----

Chapter 1

ORIGIN OF LIFE BACKGROUND AND INTRODUCTION

1.1 The Greatest Mystery

The origin of Life on Earth is the greatest unsolved mystery in the history of science. In the last hundred years our contemporary theories of matter have reached an unprecedented degree of predictive capability and our cosmological models continue to provide ever richer explanations for our position in the universe [1, 2]. Meanwhile our mechanistic understanding of biological processes is becoming clearer every day [3]. Large scale patterns in ecosystems and societies are emerging [4]. However, in spite of progress in almost every scientific endeavor, we still have no clear theory, model, or framework to understand the processes underlying the emergence of the first living organisms [5–7]. Understanding such a process would provide key insights into Astrobiology, planetary science, geochemistry, evolutionary biology, physics, and philosophy.

In this chapter I outline some of the core concepts related to studies on the origin of life. I begin with a brief summary of the historical origins of the work. I point out some issues unique to the origin of life community as an interdisciplinary scientific endeavor. Then I delve into two distinct schools of thought which have dominated the development of theoretical models, the replication or genes first camp, and the metabolism first camp. Both of these schools of thought assert that different kinds of organization explain the essential features of biology. I conclude with an overview

of the rest of this dissertation and outline where this work fits within the larger framework surrounding the origins of life.

1.2 Epistemology

“If we want to understand nature, if we want to master our physical surroundings, then we must use all ideas, all methods, and not just a small selection of them.”

- Paul Feyerabend, Against Method

The study of the origin of life on Earth has a rich and complex history. The beginning of modern scientific work on the topic can be traced back to Louis Pasteur’s seminal work which refuted the (then popular) idea of *spontaneous generation* [8]. Interestingly, prior to Pasteur’s work, many philosophers and scientists subscribed to the idea that, “living organisms were the historical outcome of gradual transformation of lifeless matter [8].” Investigations and refinements of that idea have transformed into the study of prebiotic chemistry, and represent a view known as the *material* basis of life. After Pasteur publicly refuted spontaneous generation, many philosophers and theologians used his work to justify the more vitalist views of biology [8]. The vitalist interpretation suggests that there are distinct principles which govern living matter but do not constrain non-living matter, while the material basis for life asserts that life can be completely explained and understood in of physical and chemical terms. Contemporary origin of life research can often be understood as an attempt to refute the vitalist perspective. Typically this involves demonstrating that a fundamental principle of biological organization can be understood in physical terms. Practically speaking, this means that solving the problem of the origin of life on Earth will, in

part, consist of demonstrating that what was impossible in Pasteur's sterilized flask, was possible on a sterilized planet.

Understanding the origin of life on Earth is complicated in part because of the number of different scientific disciplines which are interested in the problem [6]. This naturally leads to conflicts within the community because different disciplines tend to have different technical languages, different publication practices, and distinct standards of evidence [9]. In many scientific endeavors, conflicts of this nature can lead to new paradigms and even revolutions [10]. However it appears that progress on the origin of life may be stuck in a *pre-paradigm* phase, perhaps due to the fact that research on the origin the of life is composed of a number of different disciplines rather than defining a discipline in its own right. The historian of science Thomas Kuhn characterized pre-paradigm phases as a time when there is no consensus on any particular theory, and there is no body of knowledge that can be taken for granted [10]. A survey of literature on the origin of life reveals a plethora of different frameworks which are all based on different facts, and represent different perspectives on the same phenomena. Some paradigms are based on extrapolating backwards in time from contemporary biology to primitive life and its environment [11–15], while some focus on the dynamical and statistical features of living systems [16–21] and still others emphasize informational or computational features of biology [22–25].

Any contribution to understanding the origin of life, must be placed within the context of a particular conceptual framework. The search for an optimal, or rational means to decide between competing research programs was a significant area of study for contemporary philosophers of science [26], but the overwhelming conclusion has been that paradigms are adopted by scientific communities for sociological and historical reasons, rather than logical or empirical reasons [10, 27]. One way to

categorize paradigms is based on the way they provide explanations for phenomena. For example, paradigms which originated in the tradition of physics offer explanations in terms of first principles [28] -as opposed to historical explanations which are common in biology and geology [29]. As an example, consider the fact that all the planets in the solar system are in nearly the same plane of rotation relative to the sun. A historical explanation of this would involve recounting the steps which lead to the creation of the solar system and explaining how the planets happened to stabilize in the same plane. A first principle explanation would invoke conservation of angular momentum to demonstrate why it must have been the case that large bodies in the solar system be confined to a single plane. Candidate first principles have been suggested in various sub-fields of biology [30–32], although the ubiquity and utility of these principles remains to be tested. Paradigms which attempt to explain the origin of life on Earth in terms of first principles have an advantage relative to historical accounts, because such principles will be testable within laboratory settings, and provide possible principles of universal biology, which would be testable in exoplanetary systems. While it is not yet clear that biological theories will ever be entirely based on first principle explanations [33], in the discussion that follows paradigms which emphasize first principle explanations will be prioritized.

There is no universally accepted definition of life. Disagreements about definitions of life and the nature of living systems are a constant source of conflict between competing paradigms within the community [34–37]. While any proposed universal definition of life will not be tested until a second sample of life is discovered, there is undoubtedly value in attempts to describe life as we know it. This is exemplified by the work of Schrödinger in his lecture and book, “What is life?” By developing a rich *description* of the physical requirements of any genetic molecule, Schrödinger was

able to predict that it must have been composed of an aperiodic crystal [38], prior to the discovery of DNA. Schrödinger was able to provide such a specific prediction because he focused on developing a thorough description of the physical characteristics of genes. One of the biggest conceptual hurdles in understanding the origin of life lies in developing a language which can effectively describe many forms of biological phenomena. Therefore developing coherent descriptions of biological organization should be prioritized. If clear descriptions can be made, they lead to predictions, and eventually explanations for phenomena. Indeed many attempts to understand the origin of life begin with describing particular features of living systems. In the next section the results of starting with two very different descriptions of life are explored.

1.3 Perspectives on the Origin of Life

1.3.1 Replication First

One of the most distinctive features of life on earth is the capacity for biological entities to reproduce. The central role of self-replication in Darwinian evolution has led many to suggest that replication was essential in driving the transition from non-life to life. Many definitions of life include the capacity to engage in Darwinian evolution as a necessary requirement, and some authors have even suggested that the origin of replication is coincident with the origin of life [17]. One of the most well studied frameworks in origin of life literature is the so called “replication first“ perspective. The key assumption of this perspective is that at some point in the process of the origin of life on Earth, individual molecules with the capacity to catalyze their own reproduction emerged from a chemical milieu. Here I describe the dynamical, logical,

statistical and thermodynamics constraints on such a system. These constraints define the theoretical framework of the replication first camp.

The primary motivation for the replication first framework is that the process of replication allows for Darwinian evolution. Darwinian evolution over the course of ~ 4 billion years is then taken as the explanation for the diversity and complexity of the contemporary biosphere. However, “*survival of the fittest*” requires particular dynamics in order to occur in theoretical models [39]. To see this, consider a well mixed bath of supplies and molecular replicators. Here we ignore the effect of mutations and spatial diffusion. The role of spatial diffusion in these systems is difficult to understated, particularly in the face of mutations and parasites, although it will not be considered here. The simplest means to achieve survival of the fittest in a dynamical model like this is with strictly exponential growth, which can be described as, $A + S \rightarrow 2A + W$, where A is the replicator, S is supplies and W is waste. The scheme leads to the following growth law,

$$\dot{A} = k_A A, \tag{1.1}$$

which yields the solution [39]

$$A(t) = A(0) \exp(k_A t). \tag{1.2}$$

Two replicators in the same environment with different growth rates (*e.g.* different values of k in the growth law), will compete with each other, and in the limit of $t \rightarrow \infty$ the one with the larger growth rate will out compete the other. If fitness here is represented by the growth rate, then this situation typifies *survival of the fittest*. If the growth law is not strictly exponential this is not guaranteed [40]. Situations of this type might seem exotic, but they are easily imagined, for example if the dynamics

were instead, $2A + S \rightarrow 3A + W$, the associated growth law would be of the form,

$$\dot{A} = k_A A^2, \quad (1.3)$$

which is known as hyperbolic growth, because it is faster than exponential. In this situation competing replicators would grow to an infinite concentration in a finite time [40], which is given by

$$t_\infty \approx \frac{1}{A(0)}. \quad (1.4)$$

This situation is known as *survival of the common*, because the replicator with the highest initial concentration will ultimately out-compete others. Similarly, parabolic growth can occur when the dynamics follow the form of, $A + S \rightarrow AA$ and $AA \leftrightarrow A + A$ where the reaction rate constant of the re-association is much faster than the dissociation[40]. This leads to a growth law of the form,

$$[\dot{A}] = k_A [A]^{\frac{1}{2}}. \quad (1.5)$$

Competing replicators with parabolic growth laws will coexist indefinitely, in proportion to the square of their growth rate k and in the absence of external perturbations. This situation is known as *survival of everyone* [40]. A comparison of all these dynamics over a short time interval is shown in Figure 1.

Early in evolutionary history, replication rates may have been much slower than in contemporary biology. If replication rates were sufficiently slow, they may have not contributed significantly to the dynamics of prebiotic systems [41, 42]. Nowak and collaborators developed a mathematical model in an attempt to understand how evolutionary dynamics could have emerged from a prebiotic polymer system [17]. They utilized a binary polymer model, where every chemical species can be represented




Figure 1. Time evolution of replicators with different growth laws. Parabolic (orange), Exponential (blue), and Hyperbolic (green) growth laws are shown. In all cases $A(0) = k = 1$

as a binary string. In their model, strings could grow through the addition of a monomer to one end of the sequence, allowing any reaction of the form $i + 0 \rightarrow i0$ or $i + 1 \rightarrow i1$, where i is any binary string. Models of this form are common in models of prebiotic chemistry because they allow for an unlimited number of distinct chemical species, and they have a natural measure of “complexity” which is the length of the sequences [17, 18, 43]. Nowak *et al* demonstrated that by varying the rates at which monomers were added to different sequences, “selection” could occur without replication, a situation they deemed *pre-life*. They then demonstrated that if sequences of a given complexity (*e.g.* length) could replicate themselves selection could occur on the replicating sequences, which they called *life*. They go on to provide the conditions under which the dynamics of life dominated the dynamics of pre-life. They show that the dynamics depend on the prebiotic reaction rates, the replication rates and the minimal replicating length (complexity). Nowak *et al* did not investigate the role of finite resources in nucleating the transitions they observed in their dynamical model.

It is often assumed that Darwinian evolution is responsible for one of the most

striking features of biology: is its apparent capacity for open-ended evolution [44]. In the 1940s von Neumann considered what kind of machine would be capable of sustaining this type of evolution [45]. To formalize “open-ended evolution” von Neumann considered what would be required to build a machine capable of creating any other machine, including itself. Following Turing, von Neumann separated the machine into a “vehicle” and an “instruction tape” [45]. The process of interest can be represented as $VT \xrightarrow{VT} 2VT$,

where ‘V’ represents the vehicle or machine and ‘T’ represents the instruction tape, and the ‘VT’ over the arrow indicates that the vehicle tape combination facilitates (catalyzes) the process. The bulk process can be viewed in atomic steps as follows:

- 1) the vehicle uses the instruction tape to create a new vehicle, $VT \xrightarrow{VT} VT + V$
- 2) the vehicle makes a copy of the instruction tape, and $VT \xrightarrow{VT} VT + T$
- 3) the new vehicle and new instruction tape are joined together $V + T \xrightarrow{VT} VT$

von Neumann recognized that the instruction tape would need to play two roles, illustrated by steps one and two. First the tape must be read as instructions, and then the tape must be blindly copied. The decision of how to treat the tape cannot be contained within the tape, because the tape would need to contain a representation of itself. von Neumann argued that such a representation would cause a logical infinite regress and would not be possible in a physical system. Instead he proposed a third element to the constructor, which he labeled the “supervisory unit,” which would control how the constructor would respond to receiving a tape. Although von Neumann’s goal was to develop a physical universal constructor, he was unable to realize it. He settled instead for a toy model in a 29 state cellular automata, which was implemented after his death [45].

Von Neumann’s work on self-reproduction preceded the discovery of DNA, although

his manuscripts were published much later. Biologists now refer to the dual roles of the tape as the processes of *translation and transcription*. The lasting effect of his work on the matter was to demonstrate that this vehicle-tape logic was a *sufficient* condition for open-ended evolution [45]. In 2013, some 60 years after von Neumann’s work on the matter Chiara Marletto published *Constructor Theory of Life*, demonstrating that the vehicle-tape logic was a *necessary* condition for open-ended evolution [25]. Marletto’s motivation was distinct from von Neumann’s in that she did not intend to demonstrate how a self-reproducing machine could be created, but rather how the appearance of biological design is consistent with the known laws of physics. She formalized her argument using the newly proposed framework of *constructor theory*, and by assuming that there are “no design laws,” contained within the laws of physics. Marletto then argues that the appearance of design requires that replication can be achieved with arbitrary accuracy, which can only occur with error-correcting codes, which must be digital. The process of error correction can only occur if the tape can be copied blindly, which leads again to the arguments of von Neumann. Thus the requirement of “no design laws of physics,” necessitates a particular logical architecture for systems capable of open-ended evolution [25].

von Neumann and Marletto are some of the only authors who have focused on the logical consequences of self-reproduction. Most studies have instead emphasized the dynamical consequences of such systems. Investigations in this area originate in the traditions of mathematical biology and statistical physics. The role of thermodynamics and statistical physics in describing biological systems has been debated widely. Schrodinger argued that living systems must thrive on some type of “negative entropy [38].” Unfortunately the appropriate extensions of equilibrium thermodynamics into non-equilibrium systems remained elusive for many years. This prevented

the development of thermodynamic descriptions of the non-equilibrium patterns in biological systems and testing of these ideas, until recently. The development of several theorems about non-equilibrium statistical mechanics, which are now called *fluctuation theorems* has allowed for new investigations into the statistical physics of non-equilibrium biological processes [46]. A key result of these new investigations is a generalization of the second law of thermodynamics to include the process of replication [20].

England's generalized second law is developed directly from the Crooks Fluctuation Theorem, which states [47]

$$\frac{\pi(I \rightarrow II)}{\pi(II \rightarrow I)} = \exp(\beta Q_{I \rightarrow II}). \quad (1.6)$$

Here $\pi(I \rightarrow II)$ represents the probability of transitioning from macrostate I to macrostate II in some finite time, while $Q_{I \rightarrow II}$ is the heat generated in the transition from I to II , β is the inverse temperature. The first step in the derivation is to recognize the identities,

$$\pi(I \rightarrow II) = \int_{II} dj \int_I di p(i|I) \pi(i \rightarrow j) \quad (1.7)$$

$$\pi(II \rightarrow I) = \int_I di \int_{II} dj p(j|II) \pi(j \rightarrow i), \quad (1.8)$$

where i is any microstate in macrostate I , and j is any microstate in macrostate II .

By substituting these into Eqn 1.6, some algebra yields,

$$\left\langle \exp^{-\ln\left[\frac{\pi(I \rightarrow II)}{\pi(II \rightarrow I)}\right]} \exp^{\ln\left[\frac{p(j|II)}{p(i|I)}\right]} \langle \exp^{-\beta Q_{i \rightarrow j}} \rangle_{i \rightarrow j} \right\rangle_{I \rightarrow II} = 1. \quad (1.9)$$

England claims that this immediately yields the generalized second law,

$$\beta \langle Q \rangle_{I \rightarrow II} + \ln \left[\frac{\pi(II \rightarrow I)}{\pi(I \rightarrow II)} \right] + \Delta S_{int} \geq 0, \quad (1.10)$$

where the fact that $\exp(x) \geq x + 1$, has been used to replace the $\exp\left(\ln\left[\frac{\pi(II \rightarrow I)}{\pi(I \rightarrow II)}\right]\right)$ term. ΔS_{int} is defined as the difference between the Shannon entropy of macrostate II

and macrostate I . England then argues that macrostate I and II can be defined in an arbitrary way. It is important to note that it is currently unclear whether, arbitrary macroscopic descriptions of systems provide accurate descriptions of microscopic dynamics [21, 48, 49]. He chooses to define I as the state where there is a single replicator and raw materials, and II as the state where there are two replicators. He formalizes this using the master equation approach and saying that the growth rate of the replicator population is defined by a parameter g , while the degradation rate is given by a parameter d . He argues that the ratio $\frac{d}{g}$ should be the transition probabilities given by $\frac{\pi(II \rightarrow I)}{\pi(I \rightarrow II)}$. Substituting this into Eqn 1.10 gives

$$\ln \left[\frac{g}{d} \right] \leq \beta Q + \Delta S_{int}. \quad (1.11)$$

This can be used immediately to bound the maximum growth rate of a replicator, given by

$$g_{max} - d = d(\exp(\beta Q + \Delta S) - 1). \quad (1.12)$$

Eqn 1.12 defines the thermodynamic limits on self-replication, although it is not clear that any replicators operate near this limit, England provides a few possible examples in [20]. One of the more interesting aspects of Eqn 1.12, is that it makes Schrodinger's intuition, that replicators must feed off some "negative entropy," an explicit statistical statement. Replicators which produce a larger amount of heat (Q) in their reproductive process, can have a higher reproduction rate. This result supports the conjecture that biology emerged on the early Earth as a dissipative system, in response to far-from-equilibrium boundary conditions.

Perhaps one of the cleanest unions of theoretical biology and statistical physics comes from Gregry Karev, in an article titled *Replicator Equations and the Principle*

of *Minimal Production of Information* [50]. Karev investigated a very general version of the replicator equation,

$$\frac{dP_t(\mathbf{a})}{dt} = P_t(\mathbf{a}) (F(t, \mathbf{a}) - E_t[F(t, \mathbf{a})]), \quad (1.13)$$

where $P(\mathbf{a})$ represents the frequency of individuals with characteristics defined by the vector \mathbf{a} , and F is some measurable function of \mathbf{a} and t , often referred to as the *Malthusian fitness*, and $E_t[F(t, \mathbf{a})]$ is its expected value [50]. He was able to demonstrate that any solution to this equation can be written as a generalized Boltzmann distribution, which takes the form

$$P_t(\mathbf{a}) = \frac{\exp \Phi(t, \mathbf{a})}{\mathbf{Z}} P_0(\mathbf{a}), \quad (1.14)$$

and further, he was able to demonstrate that any generalized Boltzmann distribution solves Eqn 1.13. Karev goes on to demonstrate that the solutions to 1.13, and indeed all generalized Boltzmann distributions, obey the principle of maximum entropy production, also called the principle of minimum discriminate information.

The principle of maximum entropy production, and related principles, are grounded in the principle of maximum entropy, also known as MaxEnt [51]. MaxEnt is a guiding principle in statistical inference, which states that when inferring a distribution from limited information, the distribution should be as unbiased as possible to unknown information [51]. This principle is implemented mathematically by maximum entropy methods, which maximize the uncertainty in a probability distribution subject to some known averages over the distribution (which constitute the known data). The principle of maximum entropy production (MEP) states that at all times in the evolution of a probability distribution, the entropy of the distribution is maximized given the constraints at that time [52]. The principle of minimum discriminate information (MinxEnt) states that when updating a probability distribution given new information,

the updated distribution should contain as few discriminating features as possible relative to the previous distribution [50]. MinxEnt is equivalent to MEP in situations where the prior distribution constrains the future (*e.g.* where the distribution has some kind of memory), which is the case in replicating populations [50].

The principle of Maximum Entropy Production (MEP) has been applied to a large number of dynamical systems with varying degrees of success [52, 53]. It has been suggested as a general principle of non-equilibrium statistical systems [54], but it is not clear that it necessarily generalized the features which made equilibrium thermodynamics successful[21]. If it turns out that MEP constitutes a general principle of non-equilibrium statistical systems, Karev's result would have important implications for understanding what biological systems are. If all dynamical statistical systems obey MEP in one way or another, then all such systems would yields distributions which take the form of the generalized Boltzmann distribution, which solve the replicator equation. This would imply that our models biological evolution may be equivalent to our models thermodynamic optimization.

The replication first paradigm invokes Darwinian evolution as an explanation the complexity of the biosphere. This perspective essentially represents an out growth of evolutionary biology, and an attempt to extend neo-Darwinism all the way to the beginning of biology [25]. However, the experimental difficulties in finding a simplified replicator have led many to question this perspective [55]. Similarly the emphasis on a single molecule replicator may be misplaced given the logical considerations of von Neumann and Marletto. In spite of being richly developed, the replication first paradigm has some fundamental issues, which may not be addressable from within that framework. The solution to many of the issues is contingent on discovering a primitive biopolymer capable of exponential amplification. However even if such a

molecule was found, it has never been demonstrated how genes could have become embedded within the larger context of a cell, or even a primitive metabolism.

In spite of these problem, two chapters of this dissertation are dedicated to models of molecular replication. In contrast to other models [17, 43], the studies in this dissertation focus on the impact that replicating molecules have on their environmental conditions. In chapter 3 I will demonstrate how the emergence of replication causes a dramatic reorganization in the entire system, redistributing mass between replicators and the environment. In chapter 5 I demonstrate how environmental feed backs can dynamically reduce the mutation rate in replicators, enabling the stable transmission of genetic information in the face oh high mutation rates. These models characterize previous unexplored dynamics which only emerge in closed systems.

1.3.2 Metabolism First and Life as a Non-equilibrium system

While the replication first paradigm views living systems as essentially Darwinian systems, there is another perspective which emphasizes the physical and chemical characteristics of living systems. This prospective is often framed as a response to the replication first perspective, and therefore is less unified in its guiding principles. Some authors maintain the assertion that the capacity for evolution by natural selection is the essential feature of biology, while accepting the lack of empirical evidence for the existence of self-replicating biopolymers [55–57]. These authors therefore focus on understanding how distributed systems of reactions could evolve in a Darwinian sense. Other authors set aside the Darwianian nature of living systems and instead focus directly on the thermodynamic, and nonequilibrium nature of living systems

[15, 16, 21, 58]. Both of these approaches have been referred to as the *metabolism first* paradigm.

The first perspective, which maintains a Darwinian paradigm, originated with the work of Stuart Kauffman, with a paper from 1986 in which he demonstrated that random reaction networks were asymptotically guaranteed to have *reflexively autocatalytic* sets [59]. Kauffman defined chemical reaction systems as collections of reactions (defined by the products and reactants), and associated catalysts. He then defined autocatalytic sets to be sets of reactions where each member of the set was the product of a reaction which was catalyzed by another member of the set [60]. The central idea motivating this work was that these autocatalytic sets would be units of reproduction prior to self-replication or genes, which would be able to compete with other sets in a Darwinian sense. To formalize his ideas, Kauffman considered the set of molecules X_n , represented by binary sequences up to length n . For a given length n there are [60]

$$|X_n| = \sum_i^n 2^i \approx 2^{n+1}, \quad (1.15)$$

molecules. If each molecule can be formed by attaching two smaller molecules together then the number of possible reactions to make molecules of up to length n , the number of reactions is [60]

$$|R_n| = \sum_i^n (i-1)2^i \approx (n-2)2^{n+1}. \quad (1.16)$$

Kauffman's result is the consequence of two observations about this type of system. First the number of possible reactions grows linearly with the maximum size of the molecules because $|R_n|/|X_n| \approx (n-2)$, meaning that in the large n limit there will be many more possible reactions than molecules. Second, Kauffman noticed that the number of possible reactions will overwhelmingly involve the smaller sized molecules

in the system [60]. This is due to the fact that the number of reactions of which would make a sequence of length L_* in a set up to length N is given by [60]

$$R_{L_*}^n = (L - 1) + \sum_{i=L+1}^N 2 * 2^{i-L}. \quad (1.17)$$

Thus for a set of sequences up to length n , the ratio between the number of ways to make sequences of length L and sequences of length $L + 1$ is

$$\frac{R_L^n}{R_{L+1}^n} \approx 2. \quad (1.18)$$

With these two facts in mind, we can ask the question, if there is a uniform probability p that any given reaction is catalyzed by any given molecule, what is the likelihood of finding an autocatalytic set within the set of reactions? This is a particular type of percolation problem. Kauffman noted that the existence of an autocatalytic set is effectively assured if most of the longest molecules are produced by a catalyzed reaction. If that is was the case, the increased density of reactions creating smaller molecules would assure that the rest of reactions needed to close the set would also be catalyzed. These arguments led the the probability of **not** finding an autocatalytic set,

$$\bar{P} = (1 - p)^{(N-1)2^{N+1}} \quad (1.19)$$

$$\bar{P} \approx \exp(-p * (N - 1)2^{N+1}). \quad (1.20)$$

In the large polymer limit, Eqn 1.20 shows that the probability of catalysis p can be made arbitrarily low and the appearance of autocatalytic sets is still assured. This result was widely criticized by organic chemists and biochemists. They noted another implication of this model was that the average number of of catalyzed reactions per molecule grew exponentially with N , which is not observed in empirical data. This

issue and almost every other complaint leveled at Kauffman’s original result have been addressed and resolved [61–63], without negating the original conclusion.

Once it had been shown that these binary polymer models of chemical reaction systems could in principle contain autocatalytic sets, the question became whether or not such sets could engage in Darwinian evolution [56, 59, 64]. The evolvability of autocatalytic sets has become the source of great confusion [64–66]. In an attempt to understand the problem, Vasas *et al* constructed a kinetic model of the growth and competition of autocatalytic sets based on the GARD model of Lancet *et al* [65]. They imagined an ensemble of molecular species which could catalyze the production of other molecules in the ensemble. The catalytic efficiency for each reaction was drawn from a log-normal distribution known to model peptide catalysts. They then considered the ensembles of molecules to be contained within many different compartments. The total number of compartments was held constant throughout the experiment. When the concentration of molecules in one of these compartments reached a critical value, the compartments were split in two, at which point each molecule was transferred to one of the two new compartments, since the total number of compartments was held constant, one of the new ‘daughter’ compartments replaced the previous one and a different random compartment was replaced by the other new daughter. This scheme mimics the well known *Moran Process*, which models evolution [67]. The authors first ran this simulation to determine the steady state abundances of different compositions in the absence of a “fitness landscape.” With the control case established, they then imposed a fitness advantage to a particular composition, by increasing the catalytic rates in each compartment by $fH(m)$, where $H(m)$ was a measure of similarity to a ‘master composition’, and $f > 1$. The authors claim that if these systems were evolvable, the master composition would be more common in the selection experiment

than in the control case [64]. They found that the changes in the distribution of composition abundances were not statistically significant, which led them to claim that this type of composition based inheritance was not evolvable [64].

To confuse matters, the very same authors published a paper almost exactly two years later entitled, *Evolution Before Genes*, which claimed to elucidate the necessary conditions for a system of catalyzed reactions to engage in Darwinian evolution [56]. In this model, the authors abandoned the previous GARD framework, and instead adopted the binary polymer model proposed by Kuaffman. The authors tested the evolvability of these networks by using the same Moran style selection experiments as the previous paper. They found that some randomly generated binary polymer reaction sets were capable of Darwinian evolution [64]. They posited that the ability of networks to evolve was contingent on the presence of a particular kind of network motif, which they labeled, *viable cores*. These viable cores are essentially small collections of molecules where each molecule in the core is catalyzed by and catalyzes another molecule in the core [56]. This is a stricter definition than the autocatalytic sets proposed by Kauffman, in Kauffmans autocatalytic sets each molecule needed to be catalyzed by another molecule in the set, but there could be many *periphery* reactions which did not produce catalysts. Viable cores are the strongly connected components of an autocatalytic set (see additional document 1) [56]. The authors argue that the viable cores of the reaction network can act analogously to geneotypes while the various periphery reactions can be viewed as distinct phenotypes [56]. They admit that such a system would have only a limited form of heredity, but that it would be capable of evolution.

In chapter 4 I characterize the dynamics of an autocatalytic set based on a real RNA system known as the *Azoarcus* ribozyme. This ribozyme is approximately 250

nucleotides long, and when it is broken at 3 different points, it can reassemble itself. The fragments of the *Azoarcus* form the substrates in a catalytic set. The assembled ribozyme can catalyze its own formation as well as the assembly of mutants. Using a model of this system I demonstrate how the energetic properties of Watson-Crick pairing dictate the macroscopic features of the entire population. I show that the real system is intrinsically bias towards promoting the emergence of the autocatalytic sets. This work is an example of how some of the organizational properties of living systems may “come for free” from the properties of chemical systems.

While Kauffman’s original inspiration came from the notion of self-organization, the emphasis of most work on autocatalytic sets and related concepts has been on the idea that distributed systems of chemical reactions can be considered units of Darwinian evolution. In spite of not emphasizing the replication of single genes, this school of thought still assumes the process of Darwinian evolution as the explanation for the complexity of the biosphere. However, there are some authors who have approached the problem from a radically different perspective. This perspective views living systems, first and fore most as non-equilibrium systems. Advocates of this position tend to be a strange combination of physicists and geochemists [15, 21], although they sometimes find support from evolutionary biologists [68]. This perspective does not attempt to explain the complexity of the biosphere from Darwinian evolution alone, but rather as the manifestation of many different kinds of collective phenomena.

A great example of this school of thought came from a collaboration between condensed matter physicist Nigel Goldenfeld and famous evolutionary biologist Carl Woese [69]. The premise of their work was to demonstrate that a non-Darwinian mechanism could optimize the robustness of the genetic code. The genetic code is universal to all organisms and it had previously been discovered that it was a nearly

optimal error correcting code [31]. The non-Darwinian mechanism they employed was based on the concept of *horizontal gene transfer*, which is known to occur in microbial communities [68]. The term horizontal is intended to be contrasted with the more commonly understood mechanism of vertical gene transfer. While vertical gene transfer passes genetic information down, through lineages of organisms, horizontal gene transfer passes genetic information *between* lineages. Goldenfeld and Woese demonstrated that a mechanism of horizontal gene transfer prior to the origin of heredity would have outperformed a Darwinian system at such optimizing the genetic code. The horizontal transfer of information prevented the genetic system from getting trapped in metastable (sub-optimal) configurations [69]. Goldenfeld and Woese's work on the genetic code is an excellent example of how non-Darwinian processes can lead to more complex phenomena. Indeed many other authors have begun to suggest that other non-Darwinian means of increasing complexity should be explored [70, 71].

Another dramatic departure from the Darwinian paradigm comes from some work by Harold Morowitz and Eric Smith [58]. In their paper *Universality in Intermediate Metabolisms*, they posit that the *reductive tricarboxylic acid cycle* (rTCA) is a possible primitive metabolic core. Their arguments are based several well established facts. First, the rTCA is an anabolic cycle, meaning that it synthesizes larger organic compounds from smaller ones. The rTCA is found universally throughout life on earth, although in modern life most organisms use the process "in reverse," to gain energy by breaking larger compounds into smaller one. The rTCA is an autocatalytic cycle which is along a relaxation pathway for oxidation-reduction couples in disequilibrium. This means that in a reducing environment the steps in the cycle proceed in the thermodynamically favorable direction. It is generally accepted that the atmosphere of the early earth would have extremely reducing. The rTCA produces the organic

compounds which can be used to synthesize all known classes of biomolecules. Finally, every reaction in the rTCA is first order in its reactants, meaning that it would be reliable even in dilute solutions, which would have been the case prior to the development of cellular compartments [58]. These facts taken together can be viewed as evidence that the rTCA was a type of ordered geochemical process prior to its incorporation into the biosphere. These arguments are perhaps the strongest arguments in favor of viewing life as a non-equilibrium system. This perspective truly emphasizes the thermodynamic and statistical features of living systems, not as consequences of Darwinian evolution, but rather as enablers of Darwinian evolution.

1.4 Dissertation Outline

The proceeding discussion, while sprawling, serves to frame the context of the remainder of this dissertation. In the next chapter, I will detail the history of the term *the living state* and elaborate on its connection to advances in statistical physics. I argue that the epistemological framework of statistical physics can be productively applied in astrobiology. In chapter 3 I put this argument into action and demonstrate the existence of a phase transition in a computational model of molecular replication. In chapter 4 I develop a computational model of the *Azoarcus* ribozyme, and I use that model to identify the microscopic properties of the real system drive the macroscopic properties observed in the lab. In chapter 5 I develop a model of non-enzymatic replication in which mutation rates are dynamically coupled to a changing environment. Using this model I demonstrate that replicators which modify their environmental condition are able to transmit information reliably in the face of high mutation rates.

I conclude with a discussion of the future of origin of life research and how the results of this dissertation fit into that future.

Chapter 2

THE MEANING OF THE *LIVING STATE*

2.1 Introduction

Astrobiology is the study of life in the universe [6]. However, in spite of rigorous debate, the astrobiology community does not have an agreed upon definition of life [35, 36]. To make progress in the face of this conceptual issue, astrobiologists focus on specific properties of living systems, such as replication or cellular respiration [12, 19]. This has allowed researchers to make progress in limited domains, such as characterizing the emergence of Darwinian evolution, or quantifying the detectability of biosignatures [17, 72]. Unfortunately, without a consistent definition of life, there is no clear way to integrate the progress from these domains into a better understanding of life in the universe or its origin on Earth. Here I elaborate on the emerging concept of the *living state*, which may provide a framework to enable such integration.

References to a *living state*, can be found throughout origin of life and astrobiology science [16, 21, 55, 70, 73, 74]. For different authors, the *living state* often has different meanings and connotations associated with it. For some, this term appears to be a convenient linguistic tool, used to describe the phenomena associated with biology [55, 70]. For others, this concept is intended to characterize life as a unique class of non-equilibrium processes [16, 21].

Perhaps the earliest mention of the *living state*, was by the Nobel Laureate biochemist, Albert Szent-Gyorgyi. In 1941, he wrote two very similar manuscripts, one for Nature and one for Science [73, 74]. In both he argues that to make progress

biochemists must probe the submolecular structure of biomolecules [73, 74]. In particular, he drew inspiration from the electronic properties of crystals and semi-conductors which were just becoming clear thanks to advances in statistical and condensed matter physics [73, 75]. Szent-Gyorgyi was struck by the collective behavior of electrons in semi-conductors, and hypothesized that similar principles were at play in the function of biomolecules [73, 75, 76]. He suggested that the deepest mysteries in biochemistry would only be explained by appealing to submolecular considerations. He went on to posit that certain features of the *living state* may be consequences of quantum mechanical laws [73, 74]. For Szent-Gyorgyi the *living state* could be distinguished from non-living state based on the collective behavior of electrons. Interestingly, he rejected this idea later in his career but it has recently seen renewed interest from other researchers [77, 78].

Since this first use by Albert Szent-Gyorgyi, the term has been used by many more authors [16, 21, 55, 70]. Most of these authors use the *living state* when discussing the origin of life on Earth. These authors chose to investigate the origin of the *living state* rather than the origin of living cells, or organisms. The adoption of this term may be due to the realization that the ‘atoms’ of biology, cannot exist in isolation, physically or conceptually [79]. The description of living systems requires a specification of a macroscopic (or at least mesoscopic) system, which not only contains individual components (such as cells or organism) but also the nature of their interactions and their environment [79]. Therefore, the *living state*, is used to refer to the essential features of biological processes which are not strictly contained within individual objects, but rather, manifest in the interactions between objects.

This use of the term can be found in a review of the progress on the RNA world hypothesis by Higgs and Lehman [55]. The RNA world posits that RNA played a

crucial role in origin and early evolution of life on Earth [13]. In an RNA world scenario RNA molecules are assumed to have been, at some point, the primary information carrying molecules required for primitive genetics, as well as the primary enzymatic molecules required for primitive metabolisms. In [55] the authors describe the evidence for an RNA world as well as the processes which would be required for it to exist. They report progress on RNA nucleotide synthesis, describe various models of RNA polymerization, and explore the concept of molecular cooperation [55]. In that review the authors define the *living state* to mean a state of the world in which the processes of enzymatic nucleotide synthesis, polymerization and recombination are coordinated in a such a way that RNA molecules are reliably and robustly produced. This *living state* is contrasted to the *dead state* where all those processes may exist in an uncoordinated or unorganized manner (see specifically box 3) [55]. Thus these authors use the *living state* to identify the global scale organization necessary for the persistence of the RNA world.

Other researchers have used the concept of the *living state*, to explicitly place biological phenomena within the epistemological scope of statistical physics [21, 58]. This conception of the *living state* may enable astrobiologists to integrate progress from different disciplinary perspectives into a quantitative theory of life. Living systems are influenced by many different processes, such as geological, geochemical, atmospheric, and astronomical processes [15, 80–82]. Understanding biological organization through the lens of the *living state* doesn't attempt to reduce all of these processes to physics, but rather generalizes the approach of statistical physics to accommodate the diversity of phenomena seen in the biosphere. To understand how the tools of statistical physics can be used in this way, it is important to understand the history of that field.

2.2 A Brief Synopsis of Statistical Physics

The goal of statistical physics is to reconcile the microscopic behavior of atoms or molecules, with the macroscopic properties of materials. In the late 19th century the foundations of statistical mechanics were developed by Ludwig Boltzmann, Josiah Willard Gibbs, and James Clerk Maxwell [83]. At the time, the laws of thermodynamics were still being established but the primacy of thermodynamic descriptions of natural and artificial systems were widely accepted [83]. By contrast, there were still debates about the legitimacy of atomic theory [83]. Boltzmann's goal was to advance atomic theory by showing it was consistent with the known laws of thermodynamics [83]. To that end, Boltzmann calculated the average properties of particles interacting according to Newtonian mechanics. By taking the limit where the number of particles gets very large, Boltzmann proved that his formalism reproduced the second law of thermodynamics. In essence he demonstrated that the second law of thermodynamics was a statistically guaranteed consequence of Newton's laws of motion applied to a very large number of particles. This was the first explicit demonstration that a macroscopic theory (thermodynamics) could emerge from coarse-graining (in this case by averaging) a microscopic theory (Newtonian Mechanics).

The emergence of a macroscopic theory from a microscopic theory can be understood from the example of the ideal gas law. Gases are composed of a very large number of molecules. Each one of those molecules obeys Newton's laws of motion and therefore can be described by its velocity and position. If the number of particles in the gas is N , the number of parameters required to describe the gas using Newton's laws would be $6N$, because each molecule has components of its velocity in 3 dimensions, similarly for its position. For any large number of particles the information required

to describe the dynamical properties of a gas could become huge. However, it turns out that as a larger and larger number of particles are considered, the statistical properties of the gas become highly constrained [49]. These statistical constraints guarantee that the system will have certain features [49]. In the case of gases, those features are the pressure, temperature, and volume of the gas. In the thermodynamic limit, where the number of molecules goes to an arbitrarily large number, these features completely characterize the entire gas system [48].

By the early 20th century, atomic theory was widely accepted, thanks in part to Boltzmann and the development of quantum mechanics [75]. Around this time, research in statistical physics became organized around the concept of phase transitions [75, 84]. Some examples of phase transitions are the familiar phenomena of the melting of solid materials and the evaporation of liquids. Prior to the development of statistical physics, certain features of phase transitions were well understood phenomenologically [75]. For example, it was well known that pure metals had very specific melting points, thanks to the many industrial uses of metallurgy. However, experimental and theoretical interests in phase transitions were reinvigorated in the 1930s thanks to the discovery of superfluid helium and superconducting metals [75]. While the foundations of statistical mechanics and thermochemistry provided by Boltzmann and Gibbs had demonstrated that microscopic laws of motion acting on Newtonian particles could give rise to the macroscopic properties of materials, the study of phase transitions attempted to understand how the same microscopic laws applied to the same particles could give rise to such a diversity of macroscopic phenomena [85]. How was it that water molecules, subjected to the same microscopic laws of physics could collectively exhibit the properties of a solid, liquid or gas? The empirical facts provided by new phases of matter would elude theoretical explanation for most of the century [75, 86].

Quantum mechanics had provided a description of single (or few) electrons and their interaction with hydrogen nuclei, but these new phases of matter presented novel patterns in large systems with many electrons. These phenomena were some of the first examples of collective behavior in physics [86]. Understanding these processes required a set of theoretical tools known as the Renormalization Group (RG) [86]. The RG was developed simultaneously in statistical physics and quantum field theory [84, 86]. Initially these techniques were implemented in an ad-hoc manner in order to deal with infinities which emerged in quantum field theories. However, the subsequent formalization of RG thanks to Freeman Dyson and later Kenneth Wilson demonstrated that RG techniques need not be ad-hoc. The modern understanding of the RG is that it represents a set of tools to describe how different theories transform into each other when viewed from different perspectives [84–86].

Distinct states of matter emerge from similar microscopic systems because the collective behavior of the microscopic parts changes as larger and larger systems are considered [85, 86]. For example, the key difference between steam and liquid water is that individual molecules in steam have velocities seemingly independent from one another, whereas in the liquid state they are strongly correlated. This difference is not obvious at the microscopic scale. When observing a single molecule, whether in the gas or liquid its motion will be correlated with the other molecules nearby, due to inter-molecular forces. However as you consider more particles, the effect of this correlation tends towards zero in the gas because particles rarely interact in gases, due to their low density. Meanwhile in the liquid, with its higher density and therefore the higher interaction frequency, the effect of these correlations tends to increase. This qualitative difference (between zero and non-zero correlation) emerges as a consequence of quantitative differences in the microscopic dynamics, and is

responsible for the different macroscopic properties of the two phases [85]. In the study of phase transitions, these qualitative differences are usually tracked using *order parameters*, which are macroscopic properties that distinguish between different states. Usually order parameters will take on a value of zero in one phase and a non-zero value in the other [84, 86].

The history of statistical mechanics is a story of reconciling different descriptions of nature. Equilibrium statistical mechanics was successful because Boltzmann demonstrated that the laws of thermodynamics emerge as a consequence of the dynamics of many-particle systems [83, 85]. Those properties which are statistically guaranteed by the microscopic dynamics end up defining thermal states at the macroscopic scale [49, 86]. The renormalization group demonstrated how systems with similar microscopic dynamics can result in different macroscopic states by formalizing how descriptions of those microscopic dynamics change as they are probed at different sizes or scales [85]. In summary, as a scientific enterprise, statistical physics in the 20th century provided answers to two very general questions [85], 1) What are statistically guaranteed consequences of a given set of dynamics, and 2) under what circumstances do those consequences change? As a conceptual framework, the *living state* attempts to leverage these theoretical advances to integrate progress from many different fields into a coherent theory of living systems.

2.3 Life as a state of matter

Using the theoretical approach of statistical physics to investigate biological phenomena provides an opportunity to reconceptualize our understanding of biology. The notion of the *living state* emerges in the attempt to realize that theoretical

approach. The *living state* is defined by the collection of all statistically guaranteed properties associated with the biosphere, in the same way that the gaseous state is defined by the pressure, temperature and volume of the container. Adopting this view doesn't necessarily propose a definition of life, but rather a description of the features of life on Earth which are relevant at the large scale. This prospective has led many researchers to re-evaluate established empirical data, as in [58], and it has led to new scientific questions [87, 88].

In the study of thermal states, the relevant properties emerge as the number of particles approaches 10^{23} (one mole). It is still not clear how to determine the appropriately "large" scale at which the relevant features of the *living state* emerge. Biologists study living systems at a number of different length and energy scales, from the molecular to ecological. Recent advances in DNA sequencing, metagenomic analysis and information sciences have enabled scientists to develop databases which span all of these scales [87]. These global databases have opened the possibility of studying life on Earth at the scale of the entire biosphere [21, 87, 89–91]. These studies have led some authors to suggest that the relevant features of the *living state* only emerge at the scale of the entire planet.

This view represents a radical departure from many traditional perspectives in biology [92]. For example, in [89] Falkowski *et al*, argue that one of the most important features of the *living state* is the way in which it facilitates global scale cycling of material and energy by interacting with geological and atmospheric processes [89]. They argue that these processes emerge not due to the dynamics of individual organism or even species. Instead they suggest that horizontal gene transfer is one of the key dynamical processes which statistically guarantees those features of the biosphere [89]. Prioritizing the role of horizontal gene transfer stands in stark contrast

to most work in the biology, which emphasizes the role of evolutionary dynamics by vertical descent in shaping the relevant features of living systems [33, 92].

While some researchers have suggested that the defining characteristics of the *living state* emerge at the scale of the biosphere, others (including myself) have suggested that the defining features of the *living state* emerge at many scales, not just one [5, 87]. We recently demonstrated this concept using biochemical reaction networks, which were constructed using genomic data [87]. We analyzed over 28,000 networks across three different scales of organization. We used individual genomes to construct networks for organisms, metagenomes to construct networks for ecosystems, and every known biochemical reaction to construct a network for the entire biosphere. By comparing the statistical features of these networks, we found that they shared certain properties which could not be explained simply by the shared rules of biochemistry [87]. These features appeared in genomes from different evolutionary domains, in metagenomes from different environments, and in the biosphere as a whole. The ubiquity of these features suggest that there may be underlying dynamical laws out of which these emerge as a statistical guarantee.

As the essential features of the *living state* are better characterized and understood they will help inform our understanding of the origin of life. In the context of the *living state*, the origin of life has a natural interpretation as a phase transition [21, 88]. Just as in thermal states where the laws of physics are the same for molecules in a gas or a liquid, the laws of organic chemistry are the same for carbon in the *living state* or in the non-living state, but the macroscopic consequences of those laws are very different. Understanding how these macroscopic differences manifest will require identifying the relevant order parameters for distinguishing the living and non-living states.

Contemporary research in the origin of life suggest a few candidate order parameters [21, 88]. Smith and Morowitz have argued extensively that the origin of life on Earth emerged as a response to planetary scale disequilibria [21]. This chemical disequilibria is due to the extremely different oxidation states of the Earth's mantle and atmosphere, where the relatively reduced mantle is much richer in electrons than the relatively oxidized atmosphere. They argue the biosphere dissipates this disequilibria by facilitating the flow of electrons from reduced sources in the mantle to oxidized sinks in the ocean and atmosphere. Accordingly, they suggest that the flow of electrons through organic carbon may be a key order parameter for the *living state* [21]. This conclusion is remarkably similar to Szent-Gyorgyi's original hypothesis that the collective behavior of electron is responsible for the *living state*.

In my own work I have demonstrated that the origin of life-like properties may be effectively tracked using information theoretic quantities [88]. We developed a chemical kinetic model of primitive replicators which are strongly coupled to a dynamic environment. In that model we observed two stable states, which dynamically emerged. In the first state, labelled the non-life state, few replicators exist, and they are not selected according to their fitness. By contrast, the life state is dominated by replicators that were dynamically selected according to their fitness. To characterize the relationship between replicators and their environment we employed mutual information, which is a non-linear measure of correlations. We saw that the transition from the non-life to life state was tracked by these correlations, consistent with the idea that the *living state* is characterized more by the relationship between individual components rather than the components themselves [88].

2.4 Conclusion

Both features of the biosphere discussed here- life's interface with geochemical processes and the universal features of biochemical networks- may be independent of the particular details of terrestrial biochemistry alone. Any living system would be expected to interface with it's planetary environment, and the universal features of biochemical networks can not be explained by their shared biochemistry. Similarly, the two candidate order parameters for the living state discussed here do not require specific information about life on Earth. Accordingly, these features should be of great interest to astrobiologists who seek to understand life as it could be, not life as it is on Earth. The key to understanding the relevance of these features to biological organization lies in viewing life as a state of matter.

The *living state* is defined by the collection of statistically guaranteed properties associated with the biosphere. This concept emerges when scientists attempt to apply theoretical concepts from the field of statistical physics to characterize biological systems. Adopting this prospective leads to new scientific hypotheses regarding the nature of life on Earth as well as it's origins. These new research directions assume that "life" is a phenomena that manifests at a macroscopic scale and attempt to identify the key parameters characterizing that phenomena. These features may be independent of particulars of Earth life's chemistry and would therefore be useful in guiding searches for life beyond our planet. Thus, the concept of the *living state* may prove fundamental in the future of astrobiology.

Chapter 3

THE EMERGENCE OF LIFE AS A FIRST ORDER PHASE TRANSITION

This chapter was written in collaboration with Dr Tanmoy Bhattacharya and Dr Sara Imari Walker. It published in Astrobiology in 2017 [88]

3.1 Abstract

We demonstrate a first-order phase transition from non-life to life, defined as non-replicating and replicating systems respectively, and characterize some of its dynamical properties. The model differs from those described previously in that we explicitly couple replicators to their environment through the recycling of a finite supply of resources. Correspondingly, we find that the environment plays an active role in the dynamics of the transition. The phase transition corresponds to a redistribution of matter in replicators *and their environment*, driven by selection on replicators. In the absence of successfully repartitioning system resources, the transition fails to complete, leading to the possibility of many frustrated trials before life first emerges. The mutual information shared between replicators and environment accurately tracks the progress of the phase transition. The phase transition is marked by a sequence of abrupt transitions whereby replicators become increasingly distinct from their environment. Often, the replicators that nucleate the transition in the non-life phase are not those that are ultimately selected in the life phase. During the phase transition the system experiences an explosive growth in diversity. We discuss the implications

of these results for understanding life’s emergence and evolutionary transitions more broadly.

3.2 Introduction

Life is a state of matter characterized by stable patterns of non-equilibrium behavior. It has been suggested that dynamics of self-replication and/or autocatalysis are responsible for the persistence of such far-from-equilibrium patterns [93], a phenomena known as *dynamic kinetic stability* (DKS). Accordingly, numerous theoretical studies have focused on the emergence of the first replicators, including identifying the conditions under which replicators can be selected. Of note is the transition from “pre-life” to “life” observed by Nowak and collaborators, where pre-life is defined as a generative chemistry with no replication, to be contrasted with life, which replicates and evolves [17, 42, 94]. By tuning model parameters in those studies a transition from ordinary, thermodynamic stability to DKS is observed to occur. Similar features have been noted by Wu *et al.* [43, 95] and others [93]. While these studies have posited necessary conditions for this transition to occur, it is unclear what sufficient conditions are required for a system to prefer DKS over thermodynamic stability. It is well known that natural selection acting on replicators can have important geochemical consequences [96]. We therefore propose that a complete understanding of this transition requires a description of what effect primitive replicators can have on the chemical environment which generated them.

In the present study, we demonstrate a spontaneous, first-order phase transition from non-life to life, defined as non-replicating and replicating systems respectively. Further, we establish the essential role of dynamic environmental conditions in defining

the properties of this transition. A prominent feature of prebiotic chemistry, as evidenced by more than sixty years of experiments in prebiotic synthesis, is that it is difficult to abiotically synthesize biomolecules with high yield under early Earth conditions [14]. Accordingly, we regard resource limitation to be a common, and likely important, feature of prebiotic environments. Indeed, previous studies have indicated that finite resources could have been an important factor in *driving* the emergence of life [18, 97–99]. We therefore consider a model prebiotic chemistry with a finite supply of monomers, which must be recycled through polymer degradation to replenish resources available for synthesis of new polymers. We note that the majority of theoretical models for the emergence of replicators, by contrast, implement reactor flows with a constant flux of monomers into the system and removal of chemical species via dilution (see *e.g.* [17, 39, 95]). We are motivated by empirical evidence that recycling, mediated by coupled polymerization and degradation reactions, may have been an important mechanism driving the early evolution of biopolymers [100–102]. The explicit incorporation of non-linear feedback between environment and replicators via degradative recycling presented here therefore distinguishes our computational model from most studied previously.

The environmental feedback drives the most notable feature of this model, the transition corresponds to a redistribution of matter in replicators *and the environment*. This redistribution of matter can be tracked by measuring mutual information shared between replicators and environment. The transition has a higher probability to occur when replicators and environment share relatively high mutual information. Selection on the fitness of replicators, including their replicative efficiency and stability as studied here, occurs only in the life phase. The first replicator(s) that appear typically match the bulk composition of their environment. Since the composition of

the environment does not in general match that of “fit” replicators (the environment is not fine-tuned for life), the replicators that nucleate the transition in the non-life phase are often not those that are ultimately selected in the life phase. Due to resource constraints the phase transition proceeds through a sequence of abrupt transitions whereby the composition of replicators becomes increasingly distinct from that of their environment. During the phase transition the system experiences an explosive growth in diversity, and concomitantly a massive extinction of extant chemical species to accommodate restructuring driven by the selection on replicators. Extant diversity and the rate of exploration of novel diversity is higher after the transition than before. In the absence of successfully repartitioning system resources, the transition fails to complete, leading to the possibility of many frustrated trials before life first emerges. This feature suggests that presence replicating entities is a necessary but not sufficient condition for a dynamical system to support DKS. A further requirement may be a dynamic, and malleable environment. We discuss the implications of these results for understanding the emergence of life and evolutionary transitions more broadly.

3.3 Methods

3.3.1 Model Description and Motivation

We model the emergence of replicators in an artificial prebiotic chemistry consisting of two monomer types denoted by ‘0’ and ‘1’. Polymerization occurs via addition of monomers to the end of growing strings. Sequences can degrade into shorter sequences, which can occur at any bond within a given sequence with equal probability. We assume that the inverse process of two short but non-monomeric sequences ligating to

produce a longer polymer is sufficiently rare to be neglected. Sequences of length $L \geq r$ can self-replicate. This minimal replicator length approximates a minimal complexity for self-replication in our simplified model, and provides a clear distinction between replicators (sequences of length r or greater), and their environment (sequences of any length shorter than r). We note that the properties reported here are general and qualitatively similar for any r , where r primarily determines the relative timescale for discovering replicators, and thus for the phase transition to occur (described below). In this study, we set $r = 7$, such that the appearance of the first replicators is rare, but not so rare that we never observe it [95]. We expect that changing r will change the timescale of the transition but will not qualitatively effect the results presented here.

In the absence of an imposed fitness landscape (defined below), the properties of our chemistry are fully specified by the rate constants k_p , k_d and k_r for polymerization, degradation and replication, respectively. We regard these parameters, along with the abundances of '0' and '1' monomers, as fully specifying the prebiotic environment in our model chemistry. Thus, for example, a difference in temperature defining two different prebiotic environments would correspond in our system to simulations conducted with two different k_d values, that is, if temperature affects the degradation rate of polymers for the particular chemistry of interest (*e.g.* which might perhaps depend on the backbone chemistry of biopolymers, see [18] for discussion).

Since we are interested in the dynamics of replication in this work, and specifically the transition from non-life to life, we do not include the effects of mutation, which is well known to play an important role in evolution once life has already emerged [19]. Therefore a simplification in our model is that replication only functions to copy extant sequences, and does not produce any novelty. Novelty is introduced through

the prebiotic processes of polymerization and degradation. As long as mutations also obey the principles of resource constraints, we expect that their primary effect would be to increase the search rate for replicator(s) that match their environment, which, as we discuss below, nucleates the phase transition described here.

Simulations were implemented using a kinetic Monte Carlo algorithm [103, 104]. For more detailed discussion of the implementation of that algorithm in prebiotic recycling chemistries we refer the reader to [18] or [101]. In what follows, the polymerization, degradation, and replication rate constants were set to $k_p = 0.0005$, $k_d = 0.5000$, and $k_r = 0.0050$ respectively, and the system was initialized with 500 monomers each of ‘0’ and ‘1’ with no polymers present, unless otherwise noted. It is important to note that since this is a closed system, the initial conditions specify the bulk composition of the system for all time.

3.3.2 Two Fitness Landscapes: Static and Dynamic

To explicitly couple the properties of replicators to that of their environment, we model the fitness of replicators as determined by two factors:

1. a *static* fitness associated with a trade-off between stability and replicative efficiency intrinsic to individual replicators , and
2. a *dynamic* fitness associated with resource availability in the environment .

The first introduces a non-dynamic component of the fitness landscape associated with the properties of individual replicators that is a common feature of origin of life models. The latter environmentally-dictated fitness is dynamic and a unique feature of the resource-dependent replication model presented here (see also [18] or [101]).

3.3.2.1 Static Fitness.

The trade-off between stability and replicative efficiency of sequences captures features of nucleic acid systems believed to play an important role in early evolution—in particular that molecules that fold well are typically not good templates and conversely that good templates often do not fold well and are thus less resistant to degradation [20, 105]. The mathematical form this trade-off implemented here is inspired by that of [105], in that we encode both the stability and replicative efficiency of replicators utilizing a sigmoid function:

$$f(n) = 0.5 + \frac{n^2}{2(10 + n^2)}. \quad (3.1)$$

The replicative fitness of a sequence x_i with length L is quantified by scaling its replication rate by $1 + f(\text{zeros}(x_i))$. Similarly, the stability of sequence x_i is determined by scaling the degradation rate by $1 - f(\text{ones}(x_i))$. Thus, for sequences of a given length $L \geq 7$, the homogeneous string of ‘1’s is the most stable sequence of that length, the homogeneous string of ‘0’s is the fastest replicator, and sequences with a roughly equal number of ‘0’s and ‘1’s best balance stability with replicative efficiency. This landscape is constructed to reflect two features we regard as important to the generally unknown shape of biopolymer fitness landscapes: the fittest sequences are much rarer than less fit sequences and their composition is not necessarily reflective of the bulk composition of their environment. Sequences with $L < 7$ do not replicate, so only the stability landscape is relevant for short sequences. This trade-off establishes a fitness landscape intrinsic to a polymer’s specific sequence composition that is fixed within a given environmental context.

3.3.2.2 Dynamic Fitness.

Replication rates are also dynamically determined by the availability of free monomers in the environment. The replication rate for sequence x_i is weighted by a factor $\sum_{n_i}^{L-1} y_{n_i} y_{n_i+1}$, where y_n is the abundance of the monomer species at position n in sequence x_i . This term yields a resource-dependent replication rate that is also sequence dependent, similar to that implemented by one of us (SIW) in [101] where analogous dynamics to the phase transition reported herein were observed, although not characterized as such. The functional form of the resource dependence is an approximation intended to model nucleation of the first bond formed on a template as the rate-limiting step, while simultaneously capturing sequence information by summing over all possible nucleation events on the template, as introduced in [101]. Thus, in our model the replication rate of a given sequence x_i depends in part on how well its sequence composition matches the relative abundances of ‘0’ and ‘1’ monomers in the environment. Since the abundances of ‘0’ and ‘1’ monomers change over time as monomers are consumed via polymerization and replication and generated via degradation, this creates an environmentally dictated fitness landscape that is a central feature of any resource-constrained dynamics. We expect qualitative features of the dynamics observed here to be a general feature of sequence-specific resource-dependent replication that does not depend specifically on the form of dynamic landscape chosen for this study and will be qualitatively similar for other replicator models where stoichiometry plays a role in setting the efficiency of replication.




Figure 2. Ensemble averaged compositions of all sequences with $L \leq 7$. The distributions in the top-panel characterize the non-life phase, and the bottom-panel characterize the life phase. Data is averaged over 100 simulations.

3.3.3 Mutual Information as a Measure of the Transition from Non-life to Life

To characterize the dynamics of the phase transition, we employ mutual information, a common tool in information theory, which measures the mutual dependence of two variables within a dynamic time series. We use mutual information to measure the extent to which the composition replicators is determined by their environment, and visa-versa. We define the sets R and E which contain ordered pairs that track the number of zeros and number of ones in both replicators (R) and in free monomers in the environment (E), allowing us to measure $\mathcal{I}(R : E)$.

3.4 Results

3.4.1 Non-life and Life

Two long-lived states are observed in our model: “non-life” and “life”, which are dominated by polymerization and replication, respectively. While the non-life phase here shares features in common with “pre-life” as previously characterized [17], it also has some striking differences. We therefore use “non-life” rather than “pre-life” as it

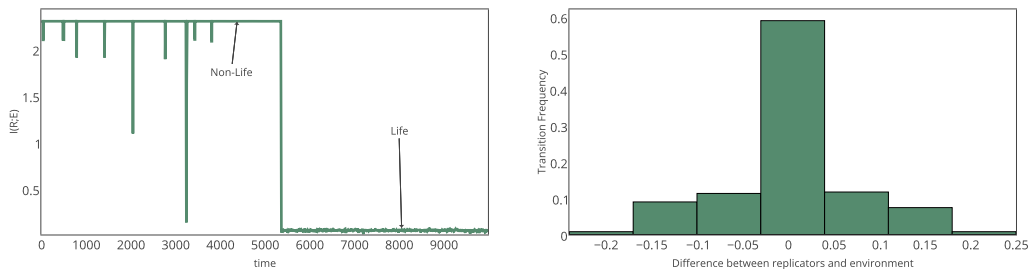


Figure 3. Top: Time series showing the mutual information between extant replicator composition and the environment. The phase transition is clearly evident in the abrupt shift from $\mathcal{I}(R; E) \sim 2.5$ in the non-life phase to $\mathcal{I}(R; E) \sim 0$ in the life phase at $t \sim 5500$. Also evident are several failed transitions, including one that nearly runs to completion before reverting back to the non-life phase near $t \sim 3000$. Bottom: The frequency of successful phase transitions plotted against the difference between the distribution of free monomers and replicator composition, for an ensemble statistic of 256 simulations.

does not imply life will inevitably emerge since in our system many transitions fail to complete.

In the non-life phase, long sequences are exponentially rare, and the majority of system mass is in monomers and dimers (not shown). Sequences of all lengths have relatively similar composition, as shown in Fig. 2. The composition of extant polymers is reflective of the abiotic availability of resources and the stability landscape established by Eq. 3.1. In contrast to other models [17, 43, 93], here *replicators exist in the non-life phase*, albeit at exponentially low abundance. However, selection on replication cannot overcome environmental constraints. Replicators in non-life are not the most fit in terms of stability or replicative efficiency (*e.g.* homogenous ‘0’ or ‘1’ sequences, respectively) but instead are predominately heterogenous with compositions determined by the bulk distribution of resources (right, top panel in Fig. 2).

In the life phase, the symmetry of the environment, constituting an equal number

of ‘0’ and ‘1’ monomers, is broken in the composition of replicators due to selection on the static fitness landscape, a feature which is not observed for non-life (compare $L = 7$ compositions, Fig. 2). In the life phase, replicators are primarily homogeneous ‘1’s or ‘0’s. The asymmetry imposed by selection on replicators is exported to shorter sequences, which have the opposite compositional signature than that of the replicators. The compositional reversal is seen only below $L = 6$. Although $L = 6$ sequences cannot replicate, they are formed primarily via degradation of $L = 7$ replicators and thus their formation is dominated by templated assembly. In the life phase, we observe that replicators are selected based on their intrinsic fitness and not strictly their composition. The defining feature of the life phase is therefore not necessarily the presence of replicators, which exist in both phases. Instead, the defining characteristic of “life” in this model is that the distribution of resources is dictated by selection on the properties of replicators.

3.4.2 A First Order Phase Transition from Non-life to Life

For fixed values of k_p , k_d and k_r , the system exhibits a spontaneous and abrupt phase transition from non-life to life (no external tuning), as shown in Fig. 3. The time of transition is exponentially distributed (not shown), indicative that the transition is first order [84]. Often there are many frustrated transitions prior to a successful phase transition (top, Fig. 3). The frequency that the transition will occur is dependent on how well the composition of the replicator(s) matches the environment (bottom, Fig. 3). This result is distinct from other models that do not account for environmental feedback [17, 95] – here the transition is *not* coincident with the first ‘discovery’ of a sequence capable of replication, since replicators exist in non-life. Instead, the




Figure 4. Phase trajectory for an ensemble of 100 systems transitioning from non-life to life, moving from left to right. Axes are the mutual information $\mathcal{I}(R; E)$ between replicators and environment (x axis) and the ratio of formation (polymerization and replication) to degradation rates (y axis).

transition occurs when the dynamic fitness of replicators is high, as occurs when the composition of replicators and environment synchronize their composition. In our example, since both monomer species are equally abundant in the initial distribution of resources, the nucleation event is typically mediated by heterogeneous replicator(s) composed of a roughly equal number of ‘0’s and ‘1’s. These are not the sequences that are ultimately selected in the life phase, which include the homogeneous, fit sequences. Thus, in resource-limited models like ours, the replicator(s) that nucleates the transition will often not be that which is ultimately selected.

However, just as in models without resource restrictions, selection ultimately leads to the fixation of fit sequences. The transition is accurately tracked by measuring




Figure 5. Time series for the extant species population size and total number of sequences explored by the system. Linear fits to the explored species are shown. The exploration rate is 75% faster during the life phase compared to the non- life phase, and is 2 orders of magnitude larger during the transition.

the mutual information $\mathcal{I}(R; E)$ between the composition of extant replicators and free monomer resources (top panel, Fig. 3). Prior to the transition, replicators and monomers share a high degree of mutual information and the composition of replicators generally matches that of their environment. However, once life emerges and selection reconstitutes the allocation of resources, replicators no longer match the information content of the environment and $\mathcal{I}(R; E) \rightarrow 0$ with small fluctuations. This behavior clearly illustrates how environmental (dynamic) selection, with a high degree of mutual information between environment and replicators dominates in the non-life phase, whereas functional (static) selection, which is not dependent on the information content of the environment, dominates in the life phase.

An interesting feature of this phase transition is that, due to resource constraints, the selection of replicators coincides with dynamic restructuring of the environment (including both monomer and non-replicating ($L < 7$) sequence populations, see *e.g.* Fig. 2). Fig. 4 shows an ensemble averaged phase space trajectory through this restructuring, which shows that the phase transition moving from non-life to life equilibria is highly unstable and dominated by degradation. In both the non-life and life phases, polymer formation rates (polymerization and replication) balance rates for polymer degradation, with ratios of formation/degradation ~ 1 . However, the life and non-life phases are clearly distinguished in phase space by very different values for the mutual information: $\mathcal{I}(R; E) \sim 3.0$ for non-life and $\mathcal{I}(R; E) \sim 0.25$ for life, for results in Fig. 4 (see also time series with different model parameters, top Fig. 3). The rampant degradation through the phase transition results in a rapid and dramatic restructuring of the extant polymer population and a steep slope in the rate of sequence exploration, as observed in Fig. 5. The extant diversity and the rate of introduction of new sequences are both higher in the life phase than the non-life phase (Fig. 5), attributable to the higher turnover rate of resources in the life phase (due to the higher assembly rate of polymers via replication).

Shown in Fig. 6 is the time evolution for all sequences with $L = 7$, binned by sequence composition, for a set of parameters where the transition is prolonged enough to resolve details of the restructuring. Resources constraints enforce selection of sequences in complementary pairs that maintain the symmetry of the bulk resource distribution of the environment (50% '0's and 50% '1's). The system subsequently undergoes a series of abrupt transitions associated with increasing sequence homogeneity, where replicator composition increasingly breaks the symmetries imposed by the environment. The asymmetry introduced by replicators is exported to their




Figure 6. Series of symmetry breaking transitions in the selection of fit, homogeneous '0' and '1' length $L = 7$ replicators. Here, the subscript denotes the number of '1' monomers in the sequence (*e.g.* x_0 contains no '1's, x_1 bins all polymers with a single '1' monomer in their sequence, and x_7 contains all '1's). Parameters are $k_p = 0.0005$, $k_d = 0.9000$, and $k_r = 1.000$.

environment, as evidenced by the compositions of shorter sequences $L < 7$ and the distribution of monomers (see *e.g.* Fig. 2).

3.4.3 The Timescale for Life's Emergence

The phase transition from non-life to life described here is a robust feature of the dynamics, observed over a large range of parameters values with qualitatively similar features. Quantitative differences arise in the final abundances of replicators and in the timescale for the transition to occur, which are both sensitive to the specific details of the prebiotic chemistry under consideration. Fig. 7 shows the average time

to complete the phase transition as a function of the degradation and replication rate constants, k_d and k_r , which as noted earlier may be viewed as specifying different environmental contexts within which life might emerge. For the results presented, the transition was identified as complete when 75% of the total replicating mass was allocated in homogeneous (fit) sequences.

One might *a priori* expect the transition to be most rapid (favored) for fast replication (high k_r) and slow degradation (low k_d), however this is not what is observed. For high degradation rate $k_d = 5.0$, the time to the transition is largely independent of k_r (top panel, Fig. 7). Lowering the degradation rate ($k_d = 1.0$ and $k_d = 0.5$, bottom two panels in Fig. 7) increases the dependence of the transition time on k_r , which, on average, occurs most rapidly for relatively low k_r . This counterintuitive behavior arises as a result of the resource constraints. For high degradation rates, there is a high rate of turnover increasing the likelihood of discovering functionally fit sequences, but the probability of survival is low, so the transition time is long regardless of replicative efficiency. For lower degradation rates, high replication rates lock resources in less fit sequences, frustrating the system's restructuring, also leading to long transition times.

The rate of degradative recycling seems to be the primary factor in determining the transition timescale. Fig. 8 shows the transition time observed for different abiotic resources abundances, quantified by the ratio R of the total number of '1' monomers to total system mass. The transition timescale is not expected to be symmetric with respect to the relative abundance of '0' and '1' monomers. For large values of R (environments rich in '0' monomers that confer stability), where recycling is inherently slower, the average transition time may be much longer than in environments with fewer stable polymers. Our data supports this expectation although the variation in




Figure 7. Timescale for completing the phase transition as a function of reaction rate constants for replication k_r . Data from 25 simulations is shown, all data points are included in the box and whisker plots. The center line for each distribution is the median, the boxes contain half the data points and the bars show the range. Three values of the degradation rate constant k_d are shown, 5.0 (blue), 1.0 (purple), 0.5 (green)

transition times is large. These features suggest that environments which engender degradative recycling at a moderate rate may be the most conducive to nucleating the origin of life under resource-limited conditions.

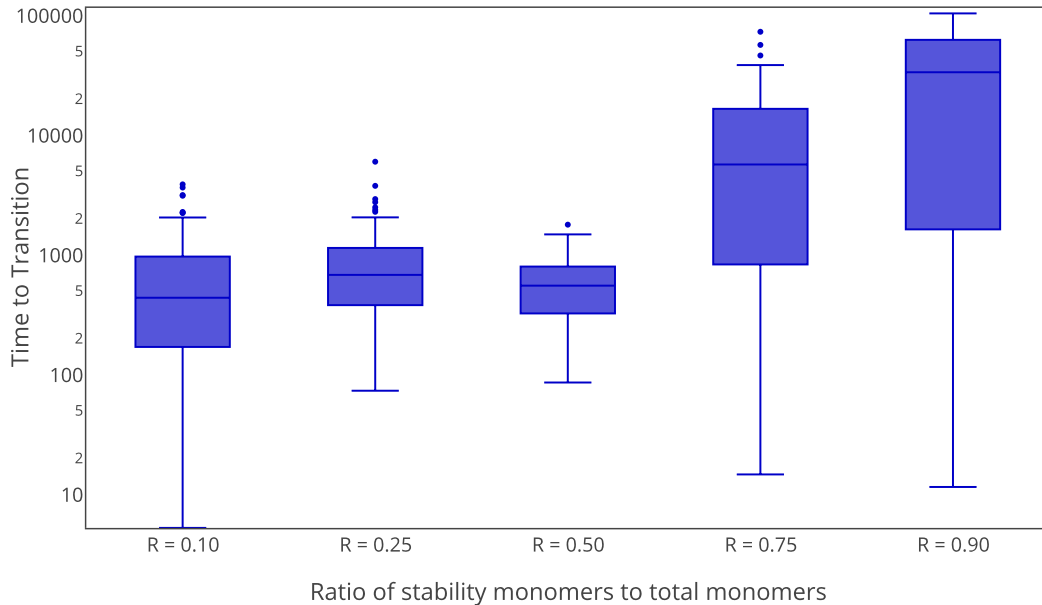


Figure 8. Timescale for completing the phase transition as a function of the abiotic distribution of resources. Here, the parameter R is the ratio of '1' monomers (which confer stability) to total system mass. Data from over 100 simulations is shown, all data points are included in the box and whisker plots. The center line for each distribution is the median, the boxes contain half the data points and the bars show the range.

3.5 Discussion

We have demonstrated the existence of a spontaneous, first order phase transition from non-life to life that demonstrates many features not previously observed, which arises due to explicit incorporation of environmental feedback. It might be argued that the dynamics reported here do not represent a true phase transition. In the study of equilibrium physical systems, free energy is the quantity which is minimized to determine the state of the system [84]. Typically, this involves a play-off between minimizing total energy and maximizing entropy. When these two favor different results, a system is expected to exhibit a first order phase transition from order to

disorder. Here, in our dynamical scenario, a similar tradeoff happens between two processes that consume and try to minimize the number of free monomers (which may be related to the minimization of free energy [106]). These two different ways—*viz.*, maximizing the number of bonds via polymerization, or by maximizing the number of polymers via replication—yield distinct results with a sharp boundary between them, which motivates the classification of the observed dynamics as a phase transition.

Due to the explicit coupling between replicators and environment, the phase transition reported here displays many features one might expect for a newly emergent biosphere that are not observable in open-flow reactor models. In particular, restructuring during the phase transition drives a vast increase in extant diversity and in the rate of exploration of novel diversity. This indicates that the emergence of life should coincide with an explosive growth of novelty in resource limited systems. Concomitantly, during the phase transition the system is dramatically restructured, indicating that the emergence of life should have significantly altered the environment of the early Earth. It is well known that biology alters its environment over generational and geological timescales, and that the presence of life defines many features of the Earth-system. It is interesting to observe this as a generic feature of life, characteristic of life as a phase of matter from its first inception, even in an abstract and simplistic model as presented here. In fact, we observe that if replicators do not transform their environment to enable their selection, the phase transition fails to complete. The model therefore includes the possibility of many frustrated trials before life first emerged (see *e.g.* Fig. 3), with success entailing a transformation of the environment as a necessary component of the process of biogenesis. These features indicate that it should be difficult to retrace the precise history of the origin of life: the replicators that are ultimately selected will, in general, neither be reflective of

the ancient planetary environment nor be representative of the replicators that first nucleate the phase transition. Thus, in general the conditions favoring the emergence of life may not be the same as those favoring its subsequent evolution.

Interestingly, the features characteristic of the phase transition are heavily dependent on degradative recycling of finite resources. This suggests new perspectives regarding the role of degradation in the origin of life, which is typically viewed as an impediment in prebiotic chemistry, rather than a process central to early evolution [13]. Cast under new light in the resource-constrained dynamics observed here, it is perhaps not a coincidence that RNA, as a biopolymer that played a prominent role in early evolution, is highly susceptible to hydrolysis, perhaps resolving an apparent paradox in the origin of life[7]. The properties of this phase transition are in principle testable in the laboratory in experimental systems that permit recycling of biopolymers, for example as reported in [101]. In particular, the observed dynamics should place further constraints on the kinds of chemistries (defined by relative rates k_d , k_r and k_p and resources abundances) that are most conducive to mediating the transition to living matter (see also *e.g.* [18]).

Importantly, the most distinctive feature of the life phase in our model is not necessarily the presence of replicators, since these also exist in non-life. Instead it is *selection* on the properties of replicators, such as replicative efficiency and stability in the example presented here. Selection in turn necessitates a redistribution of matter due to resources constraints. This restructuring is coincident with a sharp transition in the mutual information between the composition of replicators and the distribution of free monomers, which accurately tracks the phase transition. Previous work connecting information theory to life's origins reported that the probability to discover a self-replicator by chance should depend exponentially on the availability of its composite

monomers [107]. Our results demonstrate an additional necessary feature is that, in the case of resource constrained replication, replicators and environment synchronize their composition (share high mutual information). Such synchronization enables exponential growth of the replicator population based on high dynamic fitness, which in turn enables selection on the properties of new replicators discovered. When these replicators do not match the bulk composition of the system, they force restructuring to accommodate their selection. We further note that very few measures have been proposed to explicitly quantify the origin of life transition. Here, mutual information between replicators and environment accurately measures the progress of the phase transition reported, perhaps acting as an order parameter. Future work should identify how broadly applicable this approach is, by applying mutual information to other candidate scenarios for the origin of life, such as in the formation of autocatalytic sets.

We have identified replicators with “life” in this simple model. However, the definition of life is an important open philosophical and scientific question [35]. The observed information-theoretic properties of this transition are consistent with proposals that life is most defined by its informational properties [23] (here, replicators might be interpreted as driving the dynamics of the entire system in a “top-down” manner due to adaptive selection [108]). The life phase may be interpreted a state where the kinetics of individual replicators (*e.g.* as quantified by their replicative efficiency and stability) dictate the behavior of the entire system, consistent with the notion that life is a kinetically driven state of matter [93]. Although our motivation is to understand the origin of life utilizing this model system, we note that the model is sufficiently abstract to capture features that may be universal to a broader class of evolutionary transitions. In particular, the dynamics could be universally characteristic of the discovery of novel, selectable patterns in the distribution of

resources among replicating populations. For example, the abrupt nature of the transition shares features in common with punctuated equilibrium [109]. The dynamics of this phase transition also demonstrate behavior that may be characteristic to mass extinctions: the system’s restructuring necessitates a period of instability driven by rampant destruction of extant diversity, which is followed by an explosion in novel diversity. The relationship to the phase transition reported here could be tested, for example, by analyzing the connection between resource distribution patterns and abrupt evolutionary transitions.

Finally, we point out that from the perspective of stably propagating informational patterns (replicators) that are decoupled from those imposed by the bulk environment, simple replicators such as those presented here may not be the most effective architecture for a self-reproducing system. In this model, the total composition of the system remains fixed, what life does is restructure the distribution of matter within the system, due to the propagation of selectable replicating resource allocation patterns. An interesting open question is how this phase transition might play out for more life-like replicative systems, such as those with the architecture of a von Neumann self-reproducing automata [23, 25, 45], a subject we leave to future work.

3.6 Supplementary information

Herein, we explicitly measure the mutual information between two variables *as a time-series variable* itself to track the progress of the phase transition from non-life to life. To generate a time series for mutual information we use the *pointwise mutual information*, \mathcal{P} . Given two random variables $X = \{x_1, x_2 \dots x_n\}$ and $Y = \{y_1, y_2 \dots y_m\}$,

\mathcal{P} is quantified as:

$$\mathcal{P}(x_i : y_i) = \log \frac{p(x_i, y_i)}{p(x_i)p(y_i)}, \quad (3.2)$$

where $p(x_i)$ and $p(y_i)$ are the probabilities of observing the event where X is in state x_i and Y is in state y_i , respectively, and $p(x_i, y_i)$ is the joint probability of this event occurring. We generated probability distributions by counting the frequency of a given event (*e.g.* abundance of '0' and '1' monomers and of replicators of a given sequence composition) in our time series data. In the results presented here, the distributions were generated using time series data from an ensemble of 100 experimental runs over 10,000 time steps each. To ensure that the frequency based probability distributions were not biased by counting states from different phases of the system (see below), the frequencies were generated from data that sampled equally from both phases. The probabilities of different states therefore represent ensemble statistics which do not depend on time, while the particular ordering of states in a time series is used to determine the time series \mathcal{P} . In stochastic systems, such as ours, \mathcal{P} will fluctuate rapidly in time and is unlikely to yield useful insights. We therefore sum \mathcal{P} over a fixed time window to yield the mutual information for that window. Explicitly, for a window size of w , the mutual information at time t is defined as the average of $\mathcal{P}(x_i : y_i)$, and is given by:

$$\mathcal{I}(X(t) : Y(t)) = \sum_{i=t-(w/2)}^{t+(w/2)} p(x_i, y_i) \log \frac{p(x_i, y_i)}{p(x_i)p(y_i)}. \quad (3.3)$$

[110]. This value will depend on time, not because the probabilities of different states will depend on time, but rather the realization of different states is time ordered. Determining an appropriate size for w is important. If w is too large, the entire measurement collapses into one value yielding no insights into how the system is evolving in time. By contrast, if w is too small, fluctuations wash out interesting larger

scale structure. We chose w heuristically, such that the value of $\mathcal{I}(R; E)$, tracking the mutual information between replicators and environment, was relatively constant but large fluctuations could still be resolved. For the results presented $w = 100\Delta t$, where $\Delta t = 0.1k_h^{-1}$ is the resolution of the time series data in natural units. We note that different values of w change the results quantitatively, but not qualitatively: the system still maintains a non-zero value of the mutual information in the non-life phase which tends toward zero in the life phase.

Chapter 4

PREBIOTIC RNA NETWORK FORMATION: A TAXONOMY OF MOLECULAR COOPERATION

This chapter was written in collaboration with Sanjay N. Ramprasad, Dr. Sara Imari Walker and Dr. Niles Lehman. It was published in Life in 2017 [111]

4.1 Abstract

Cooperation is essential for evolution of biological complexity. Recent work has shown game theoretic arguments, commonly used to model biological cooperation, can also illuminate the dynamics of chemical systems. Here we investigate the types of cooperation possible in a real RNA system based on the *Azoarcus* ribozyme, by constructing a taxonomy of possible cooperative groups. We construct a computational model of this system to investigate the features of the real system promoting cooperation. We find triplet interactions among genotypes are intrinsically biased towards cooperation due to the particular distribution of catalytic rate constants measured empirically in the real system. For other distributions cooperation is less favored. We discuss implications for understanding cooperation as a driver of complexification in the origin of life.

4.2 Introduction

Recently, there is renewed interest in the importance of cooperation and collective behavior for both the emergence and evolution of biological complexity [112, 113] and, importantly, for the origins of life itself [55]. It is becoming increasingly clear that one can trace the biological roots of cooperation down to the chemical level [114]. Molecular behaviors can be described in terms of game theoretic analyses with kinetic fitness payoffs that have direct relevance to the evolutionary fate of molecular populations [115, 116]. It has also been suggested that cooperative effects may be responsible for the long-term stability of the living state [21].

The *Azoracus* catalytic RNA covalent self-assembly system [117] allows us to explore these concepts directly both experimentally and computationally. In this system, cooperation arises due to the interaction among distinct genotypes leading to collective fitness gain. Briefly, in this system RNA molecules form an interaction network enabling the autocatalytic assembly of similar – and dissimilar – genotypes [118, 119]. Genotype identity can be embodied in two three-nucleotide sequences at the 5' and 3' ends of the ca. 200-nucleotide ribozyme. The 5' triplet is termed the internal guide sequence (IGS). IGS binding to a psuedo-complementary triplet at the 3' end of another RNA determines the rate at which one ribozyme genotype assembles another from smaller fragments. A simplified version of this system in which only two fragments need to be covalently recombined to produce the ribozyme is depicted in Figure 1a. Here we describe the reaction as $\mathbf{WXY} + \mathbf{Z} \rightarrow \mathbf{WXYZ}$, where \mathbf{WXYZ} represents the full-length ribozyme. Cooperation, in this molecular context, is the event whereby one \mathbf{WXY} genotype assembles another ribozyme with a distinct genotype; in other words, when the IGS sequences of the catalyst and its

substrate do not match. Conversely, selfish assembly occurs when the IGS triplets of the catalyst and its substrate are the same [119] (Figure 1b). The balance between cooperative and selfish strategies is a key determinant in population evolution [41].

The origins of life likely required a kinetic competition among informational units prior to the advent of Darwinian evolution seen in contemporary life [57]. In a chemical soup this competition would have been played out in a network setting among a wide variety of molecular species. One approach to decomposing such network dynamics into smaller mechanistic elements is to use game theoretical methods to understand how pairs or triplets of genotypes interact in isolation. Game theory, in this context, is a mathematical framework in which the dynamical behavior of alternative reproductive strategies compete for limiting resources can be quantified through payoff matrices. The dynamics of such subnetworks can then be used to understand how larger chemical networks may behave. But before this can be done, it is necessary to have a complete understanding of the rules that govern subnetwork growth and interaction.

In the *Azoarcus* ribozyme system, evolvable networks can be composed of as few as three **WXY** genotypes. Previously we have explored the joint growth dynamics of a particular 3-membered network of such genotypes [116, 118]. This network embodied the “rock-paper-scissors” (RPS) game theoretic dynamic, which has been also explored at a variety of biological levels [120, 121]. In the RPS scenario, one genotype is superior to another by direct comparison, and so on around the cycle. However there are far more possible topologies than just RPS, depending on the directionalities and strengths of the various inter-genotype node interactions. In fact, there are 17 distinct 3-node, 3-edge topologies (including edges that are a self-loop), with RPS being but one of them (Figure 2).

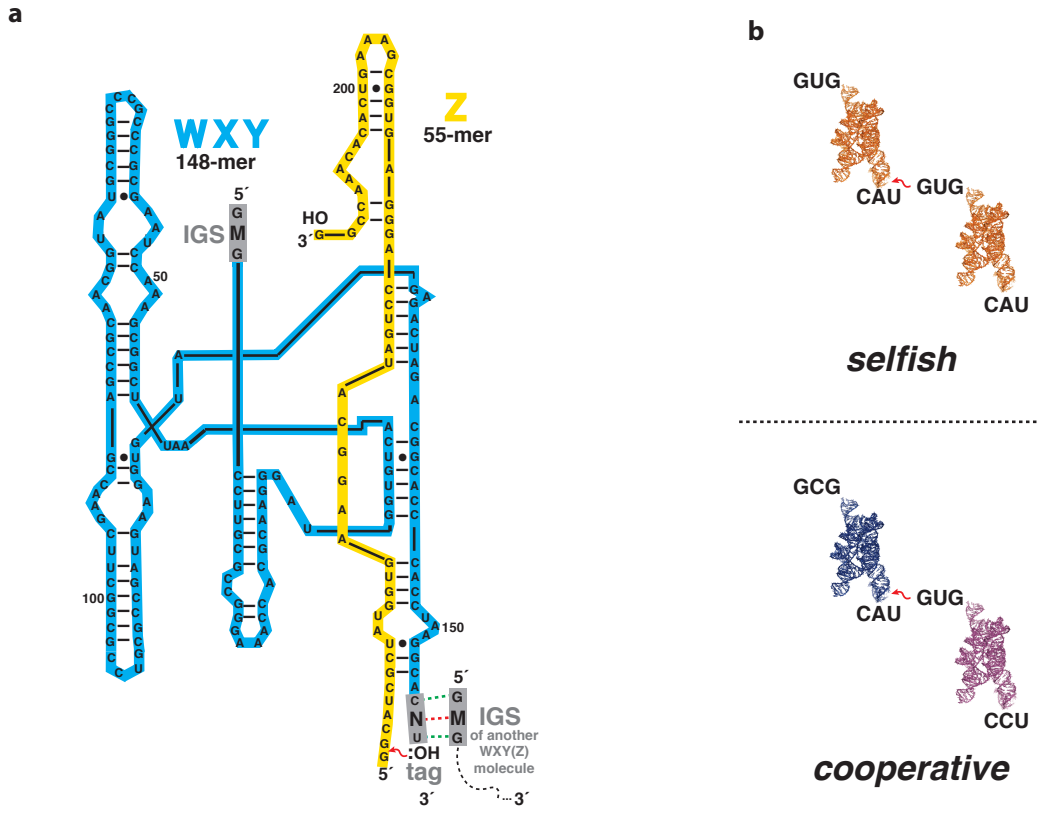


Figure 9. The *Azoarcus* ribozyme self assembly system. (a) The reaction between the 148-nt **WXY** RNA fragment and the 55-nt **Z** RNA fragment. This reaction is catalyzed by the binding of the IGS of another catalyst (either a covalently-contiguous **WXYZ** molecule that had been previously assembled or a non-covalent **WXY-Z** complex) to the tag sequence on the 3' end of the substrate **WXY** molecule. The key hydrogen bonding interactions are shown in green (invariant) and red (variable) dotted lines. (b) A comparison of selfish and cooperative assembly reactions by the use of two example pairs. In selfish assembly, the IGS (upper left of each molecule) of the substrate and the catalyst RNAs match; in cooperative assembly they do not match.

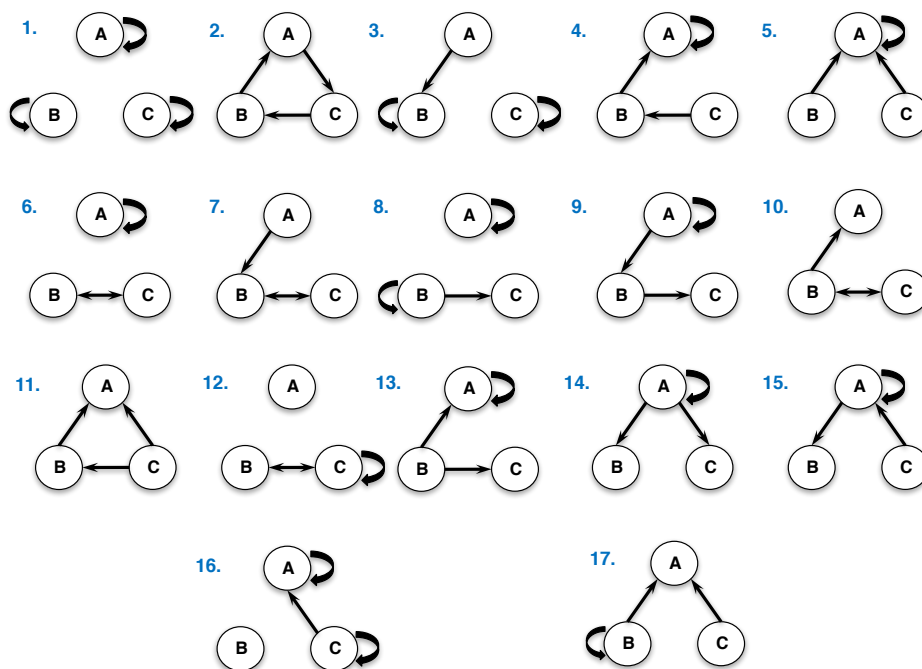


Figure 10. The 17 possible non-trivial 3-node, 3-edge network topologies. Topology 2 is the “rock-paper-scissors” (RPS) scenario. **A**, **B**, and **C** denote distinct **WXY** RNA genotypes, while the arrows denote the ability of one genotype to covalently assemble another. For comparison, there are only four 2-node, 2-edge topologies, and we explored these in detail previously [116].

The complexity of network dynamics is manifest by the constraints imposed on the processes that define the pairwise interactions of the members of the network itself. Understanding the emergence of these dynamics is of significant interest in prebiotic chemistry [57]. The topological artifacts of network graphs, e.g. edges and nodes, inform both the direction and magnitude of these interactions. Figure 2 represents a global overview of the classes of pairwise interactions possible in such 3-member RNA networks. All possible triplet genotype combinations are represented by one of these 17 classes of global network topologies. In these networks, the constraints imposed

are the real rate constants of self-assembly derived from previously demonstrated empirical data [116].

In this manuscript, we explore the hierarchical relationships among the types of molecular behavior (i.e., cooperative or selfish) in these 3-genotype networks. Specifically, we are interested in the emergence of inter-genotype dynamics at the network level due to bond type and strength at the chemical level. Using empirically determined rate constants for the various IGS-substrate interactions, we computationally model the landscape of all possible genotypic interactions and explore RNA growth dynamics for each, producing a full picture of the dynamics of small networks that would be intractable to explore in the lab. By systematically modeling all 1-genotype, 2-genotype, and 3-genotype fitnesses within these simple topologies, and categorizing the various possible fitness benefits to individuals versus groups, we construct a taxonomy of cooperation that provides new insights into the molecular dynamics of cooperation that could have driven the emergence of life.

4.2.1 Nomenclature and chemistry

As in previous work [116, 118, 119, 122], we designate each **WXY** RNA genotype by a two-letter name, MN . The first letter, M , denotes the middle nucleotide of the IGS, while the second letter, N , denotes the middle nucleotide of the target triplet on the 3' end of the **WXY** fragment (Figure 1a). We term this latter triplet the “tag” sequence. The IGS-tag triplet-triplet binding through nucleotide pairing is the key chemical and informational interaction that determines the rate at which one genotype will catalyze the covalent assembly of another ribozyme [122]. The wildtype IGS sequence of the *Azoarcus* group I intron ribozyme is 5'-GUG-3' [123], and the

Table 1. Autocatalytic rate constants for *Azoarcus* RNA covalent self assembly.

WXY genotype	self-assembly rates (k_a); min^{-1}	nucleotide pair type
C•G	0.0415	strong (Watson-Crick)
A•U	0.0319	
U•A	0.0197	
G•C	0.0125	
G•U	0.0091	intermediate (wobble)
A•C	0.0069	
U•G	0.0049	
U•C	0.0038	weak
U•U	0.0022	
C•A	0.0020	
C•C	0.0016	
G•G	0.0006	
G•A	0.0005	
A•A	0.0004	
C•U	0.0004	
A•G	0.0001	

psuedocomplementary sequence to this in the tag would be 5'-CAU-3'. Although in principle all three positions could be varied, proper group I intron activity is highly dependent on a G•U wobble at the splice site [124, 125], meaning that the first (5') nucleotide of the IGS was best fixed as a guanosine, and the third (3') nucleotide of the tag was best fixed as a uridine (Figure 1a). Moreover, we have found that variation at the middle nucleotide of the IGS-tag triplet is better tolerated than variation at either end [126]. Thus we focused only on variation in the middle position of the IGS-tag interactions (the red dotted line in Figure 1a) such that there are 16 possible **WXY** genotypes, i.e., 5'- GMG **WXY** CNU -3'. These can be abbreviated MN , such that, for example, the genotype 5'- GUG **WXY** CAU -3' can be simply denoted UA.

In previous work, we measured all possible self-assembly rates for the 16 MN **WXY** genotypes when incubated with equimolar **Z** in 100 mM MgCl_2 at 48°C [116, 122]. The rates are listed in Table 1 and serve as the empirical basis for all the estimated

genotype-genotype edge strengths in the networks we examine herein. These rates vary over three orders of magnitude and exhibit a specific distribution of values, owing to the energetics of IGS-tag interactions. The goal of the current work is to test the causal chain of events from the H-bonding patterns that exist in the IGS-tag interaction to the patterns of cooperative or selfish behavior that manifest at the population level. To our knowledge this has never been explicitly done before. Specifically, we explored the differences in growth rates, as a measure of fitness, for genotypes when they are isolated, exist in pairs, or exist in triplets.

4.2.2 Classification and Taxonomy

There are 16 possible *Azoarcus* **WXY** genotypes, therefore there are $16\text{Choose}3 = 560$ possible triplet groups, and within each of those 560 triplets there are $3\text{Choose}2 = 3$ unique pairs. Each of these triplets can be classified according to the degree of cooperation exhibited by the triplet. One type of classification is simply to count the number of growth rates which were improved in the triplet relative to the individual rates. If all the genotypes in a triplet had a growth enhancement relative to their isolated growth rates, $R_m^{ijk} > R_m^m$ for $m \in (i, j, k)$, we say that triplets has 3 cooperators and we call that triplet *cooperative*. If two genotypes improve their growth rates relative to the isolated condition, we say that triplet has 2 cooperators and we refer to that triplet as *semi-cooperative*. If only one growth rate improves relative to the isolated cases, it has one cooperator and we refer to that triplet as *selfish*, and if all the rates either remained constant or decreased in the triplet relative to the isolated rates, $R_m^{ijk} \leq R_m^m$ for $m \in (i, j, k)$, it has 0 cooperators, and we call that triplet *antagonistic* (or in the case that all rates remain constant *neutral*). These classifications describe

the degree to which all three genotypes gain a fitness advantage by being in a triplet as opposed to existing as isolated individuals. We refer to this as the first level of classification in our taxonomy.

There is another, finer, level of classification based on the number of growth rates which improved in the triplet relative to the pair-wise interactions. For any one of the 560 triplets, there are 3 possible pairs (ij) , (ik) , and (jk) , the growth rate of each genotype in each pair (e.g. R_j^{jk}) can be compared to the growth rate of that triplet in the triplet (R_j^{ijk}), allowing for 6 comparisons per triplet. We can count the number of times a triplet growth rate is greater than a pair-wise growth rate for a given triplet, giving a number between 0 and 6. This number describes the degree to which pairs of genotypes gain a fitness advantage by being in a triplet as opposed to an isolated pair. For example, if a triplet scores a 0 on this level of classification, that means that the growth rates for all the genotypes in the triplet would have been higher in either of their two possible pairs, as compared to their growth rates in the triplet. While if a triplet scores a 6 on this level, every genotype has a higher growth rate in the triplet, as compared to any of its two possible pairs. Many different configurations can generate scores between 0 and 6; for example a score of 3 could be caused by one genotype in each possible pair having an increased rate in the triplet relative to those pairs, or by one pair having both members increase their growth rates in the triplet and another individual benefit from the triplet as compared to its possible pairs. We refer to this as the second level of classification.

The overall “taxonomy” that results by considering these two levels of classification can be seen in Figure 3. One immediate result is that the taxonomy is skewed towards cooperation at both levels of analysis. There are more semi-cooperative and cooperative cases combined ($256 + 167 = 423$) than there are antagonistic and selfish

($14 + 123 = 128$). Likewise there are more cases where 4, 5, and 6 pairs are improved (261) compared to cases where 0, 1, or 2 pairs are improved (156) upon group inclusion. This bias is a combination of two factors. First, there are simply more combinatorial ways to be cooperative. This is due to the fact that none of the IGS-tag interactions are inhibitory in nature, so the presence of any other genotype in principle can contribute to an enhanced growth rate. In the finite resource simulations presented here, this effect is mitigated because the improvement gained by a small contribution to the catalytic rate is off-set by more rapidly decreasing resource abundances.

Secondly and more critically, there is an influence based on the distribution of rates as seen in Table 1. We randomized the rate distributions in various ways to compare them against the empirical distribution of rates (Figure 4). In short, we tested a completely random (flat distribution), a heterogeneous distribution (random on a logarithmic scale) a distribution which accounts for the discrepancy between Watson-Crick pairs and all other base pairing (WC distribution), as well as one which includes the effects of ‘wobble’ pairs (WCW3 distribution). For a complete description of the different distributions see *Materials and Methods*. The degree of cooperation seen in the empirically derived system is much greater than would be expected if the catalytic rate constants were homogeneous (e.g., drawn from a random flat distribution). We also found that while the WC distribution was unable to produce the degree of cooperation seen in the empirically derived system, the WCW3 distribution agreed well with the empirically derived system. The heterogeneous distribution agreed well with the empirically derived system also, suggesting that the heterogeneity of the distribution, and not the particular order of the base pairs drives the observed cooperation. In particular we note that both the WCW3 and Log

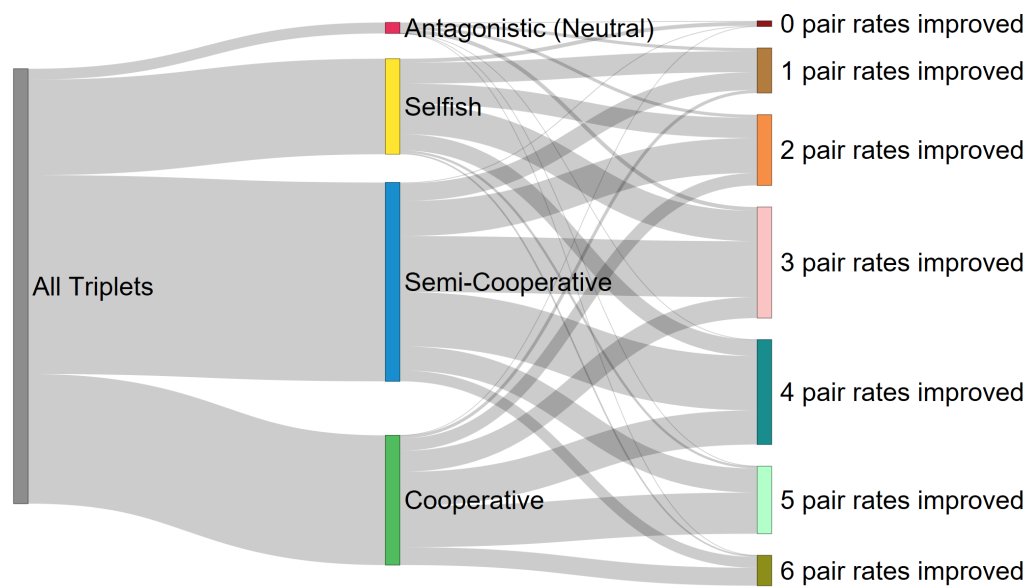


Figure 11. Taxonomy of molecular cooperation using the empirically derived rate constants. The left bar (gray) represents all possible 560 triplet genotype combinations. The bars in the center correspond to the first level classifications, and the height of the bar indicates which fraction of all triplets fell into that category. The bars on the right show the second level classifications and the fraction of all triplets which fall into that category.

Random distributions yielded many more cooperative triplets relative to the WC and flat random distributions.

4.3 Discussion

By building a kinetic model of the *Azoracus* ribozyme **WXY** genotype system, we were able to investigate the distribution of cooperative effects in that system, as well as the drivers of that distribution. We found that the real system demonstrated an enhanced level of cooperation relative to a randomized control. We are also able to observe a diverse array of evolutionary trade-offs, as the number of unique genotypes

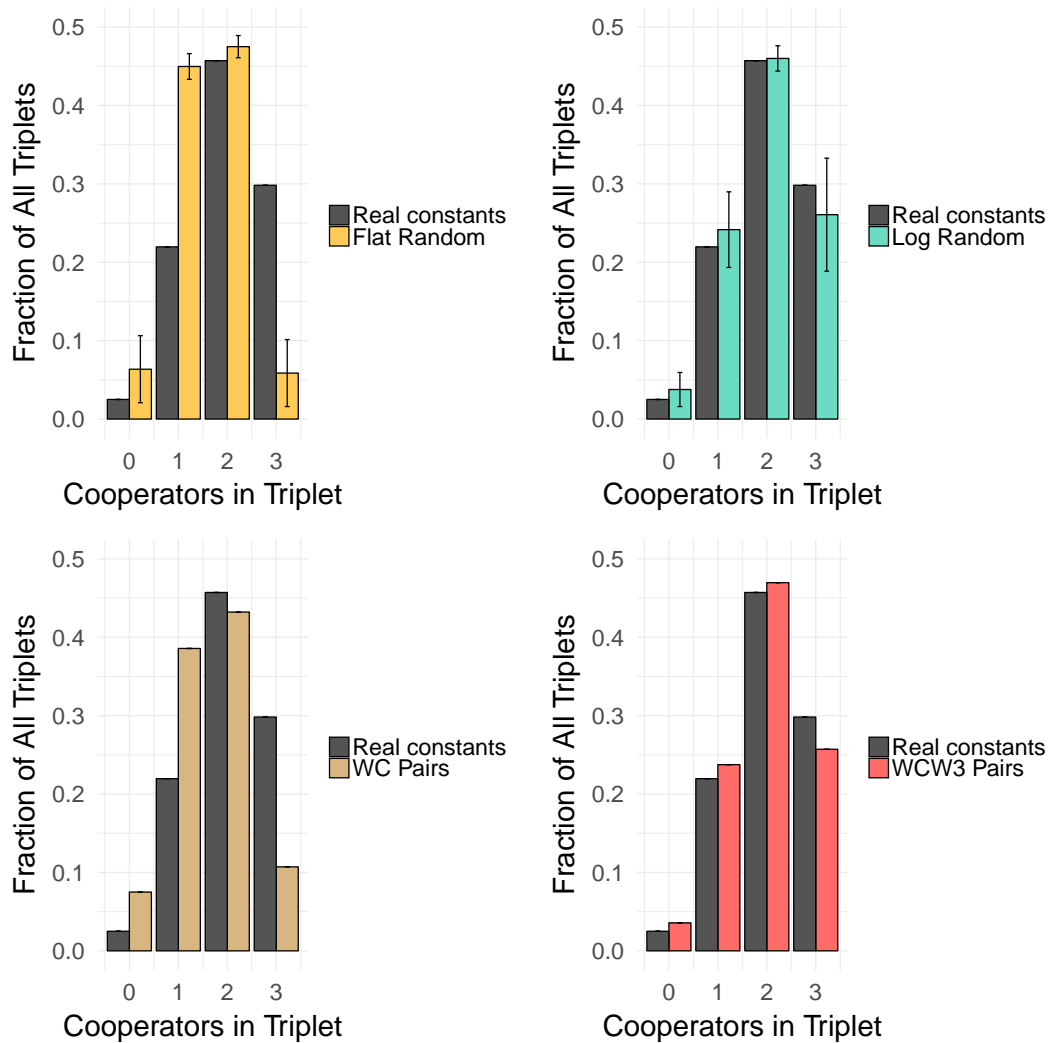


Figure 12. The effect of the distributions of ribozyme self-assembly rates on cooperation. Top, left: Flat Random *vs.* Real rate constants. The Empirically derived rate constants show an enhanced level of cooperative relative to a completely randomized control. Top, right: Log Random *vs.* Real rate constants. The empirically derived distribution of cooperation is well explained by a logarithmically distributed random distribution Bottom, left: WC pairs *vs.* Real. The empirically derived distribution of cooperation is not completely explained by including two values of rate constants. Bottom, right: WC with three wobble pairs *vs.* Real. The empirically derived rate distribution of cooperation is well explained by including wobble pairs in a simplified set of rate constants.

in a system increases. We further found that these enhanced rates of cooperation were well explained by the heterogeneous distribution of catalytic rate constants. These results have several consequences for understanding the evolution of cooperative catalytic networks.

While previous studies of collectively autocatalytic networks have provided topological constraints on such systems [57], results presented here provide new constraints on the dynamics of autocatalytic RNA networks. Prior investigations of autocatalytic chemical reaction networks find that they are rarely fixated in randomly assembled networks [127]. This rarity may be due (in part) to the uniform distribution of rate constants used in many of these studies. The results presented here suggest that the likelihood of observing such networks dynamically (either in the lab or *in silico*) would be enhanced by using a heterogeneous distribution of rate constants.

By classifying the degree of cooperation using these two scales of organization (improvement from individual rates, and improvement from pair-wise rates), we can begin to understand how cooperation in the *Azoarcus* RNA system changes under aggregation (e.g., addition of more unique genotypes). We find that there are some genotype triplets which are fully cooperative, and score a 6 on the second level classification, meaning that all the genotypes grow faster in the triplet than they do alone, and all possible pairs of genotypes grow faster in the triplet than they do in pairs. Triplets of this kind would be expected to have an enhanced stability in the face of external perturbations [128]. There are also cases when a semi-cooperative triplet scores a 5 or 6 on the second level classification. In this situation, there is at least one genotype which gains no net benefit from being in the triplet relative to being isolated but which gains an advantage relative to being in a pair with either other member of the triplet. Interestingly, we find that there are a few antagonistic

triplets, which score a 6 on the second level classification. This situation implies that each genotype in such a triplet does worse in combination with both of the other two, however if any two genotypes are present, they both gain an advantage by adding the third. These genotypes have the highest growth rates alone, but if they cannot be isolated completely they are better off in the triplet. These are just some examples of the types of evolutionary trade-offs present in this chemical system.

Understanding the evolution of cooperation in chemical systems will be essential to understanding the complexification of such systems as they make the transition to biological ones. In the current study, we tested the effect of the distribution of rate constants for ribozyme assembly. As can be seen in Figure 4, this distribution has a substantial effect on how fitness is partitioned between individual and group benefits. These rate constants in turn are determined by the chemical moieties in the base-pairing surface of a single nucleotide pair embedded in the middle of the IGS-tag triplet-triplet interactions. The steric and electronic configurations of this base-pairing surface in turn is determined by only a few atomic rearrangements, which in some cases involve as few as three atoms (compare a G-C pair with a G-U pair, for example). Thus here we are able to draw a line of causality from the atomic level to the group population level dynamics of complex macromolecules with molecular masses greater than 50,000 daltons. The dynamical behavior of ribozyme self-assembly then may be another manifestation of (or even determinant of) the chemical etiology as discussed by Eschenmoser [129].

One of the most salient features of the biosphere is its hierarchical organization [21, 23, 130–132]. The emergence of biological hierarchies can be viewed as a type of dynamic computation, where component parts are both partially constrained by and partially responsible for the higher level organization that emerges [130]. Notably, one

of the key features of the population level organization observed in the *Azoarcus* RNA system was due to the heterogeneous distribution of rate constants. In other aspects of biology, heterogeneous distributions of connections have been shown to enhance the information processing capabilities of biochemical networks [133, 134]. It is possible that the distribution of base-pairing energetics in nucleobases provides RNA systems with enhanced computational and informational abilities relative to other plausible prebiotic chemistries. That enhancement would have contributed to the ability of RNA-based systems to generate higher levels of hierarchical organization.

4.4 Materials and Methods

4.4.1 Kinetic Simulations

To explore the drivers of cooperation in this system, we developed a dynamic model *Azoarcus* ribozyme self-assembly. In contrast to previous models of this system, we considered a closed system with a fixed total number of **WXY** and **Z** fragments. We model the self-assembly process as a spontaneous reaction which can be catalyzed by assembled ribozymes. The spontaneous assembly propensity for a given genotype is determined by the total abundance of **WXY** fragments of that genotype and the number of **Z** fragments and the reaction rate constants k_s , such that the rate is proportional to $k_s[WXY_i][Z]$. The assembly of ribozymes can be catalyzed by other ribozymes. The degree of catalysis is determined by the IGS of the completed ribozyme and the tag of the ribozyme to be assembled. Catalysis increases the spontaneous propensity of assembly by a factor of $k_{ij}[X_i]$, where k_{ij} is the degree to which a ribozyme with IGS i catalyzes the assembly of ribozymes with tag j , and $[X_i]$ is the

number of ribozymes with IGS i . Thus the total assembly propensity of a ribozyme with tag j is proportional to

$$A_j \propto k_s [WXY_j][Z](1 + \sum_i k_{ij}[X_i]). \quad (4.1)$$

For the results presented here we have set the spontaneous rate constant to a small value ($k_s = 10^{-8}$), and we have approximated the values using the self-assembly autocatalysis rate constants provided in [116]. This model was implemented in a kinetic Monte Carlo algorithm [103].

4.4.2 Fitting Growth Constants

Understanding cooperation in this system requires comparing the growth rates of genotypes when they are isolated, and when they are coexisting with other genotypes. In order to make this comparison simulations were performed with genotypes individually, in pairs, and in triplet groups. For the individual simulations, the system was initialized with 1 completed ribozyme of the given genotype, 5000 **WXY** fragments of that genotype and 15000 **Z** fragments. For simulations with pairs of genotypes, the system was initialized with 2 complete ribozymes (one of each genotype), 10000 **WXY** fragments (5000 of each genotype) and 15000 **Z** fragments. Simulations with triplets were performed in a similar manner, with 3 completed ribozymes (one of each genotype), 15000 **WXY** fragments (5000 of each genotype) and 15000 **Z** fragments. For each simulation time series data for the abundance of each ribozyme genotype was recorded. To compare the growth rate of a given genotype under different conditions (e.g isolated, in triplet, or pair), an exponential curve of the form,

$$X_i^C(t) = X_0 \exp(R_i^C t), \quad (4.2)$$

was fit to the beginning of simulated time series data. Where X_i^C is the abundance of the completed ribozyme of genotype i , in condition C , and X_0 is the fitted growth rate for genotype i , in condition C . The possible conditions for a given ribozyme, are alone which we denote $C = i$ for ribozyme i , in pairs which we denote $C = ij$ for the combination of ribozymes and i and j , or in triplets which we denote $C = ijk$.

4.4.3 Randomization Experiments

In order to understand what drives the observed distribution of cooperation, we compared the distribution of cooperative ribozymes derived from the empirical catalytic rate constants to the distribution of cooperative ribozymes derived from other catalytic rate constants (e.g., different values for Table 1). Several different distributions of reaction rate constants were chosen to understand which features of the catalytic rate constants caused the observed distribution of rate constants. As a control, rate constants were drawn from a random uniform distribution with a range equal to the range of the empirical constants (0.0001, 0.0415), we refer to these as flat rate constants. Since the self-assembly process in the *Azoracus* ribozyme depends on the nucleobases in the IGS and the tag, the rate of catalytic reaction rate constants for any Watson Crick pair of IGS and tag nucleobases are much higher. Therefore the observed distribution of catalysis may be due to the spacing between the highest rate constants and the lowest. To capture these features we use three different synthetic distributions, to compare to the real distribution (Table 1). The first distribution, which we refer to as the Watson-Crick (WC) distribution, assigns the Watson-Crick IGS-tag pairs to have a uniform catalytic rate constant (0.0400), which is two orders of magnitude higher than all the other rate constants(0.0004). We also

used another distribution in which we again assign Watson-Crick IGS-tag pairs to have a high uniform rate constant (0.0400), which is one order of magnitude above the rate constant assigned to “wobble pair” IGS, tag pairs (0.0040), and all other rate constants were set to be one order of magnitude lower than the wobble pair constants (0.0004), we refer to this distributions as the Watson-Crick-Wobble (WCW3) distribution. The final distribution is one where rate constants are uniformly distributed on a logarithmic scale, which we refer to as the Log Random or heterogeneous rate constants.

4.5 Conclusions

Evolution of complexity is typically explained by understanding how individuals form cooperative groups. In the chemical events surrounding the origins of life, it appears that explaining the converse situation – how individuals arise from cooperative groups – will be an important task. The advent of cells or some type of semi-permanent encapsulation mechanism allowed for the physical linkage of genotypes with their own phenotypes, as well as energetic benefits [40, 56, 57, 135]. This “invention” of a barrier, such as proto-cells, likely led to an evolutionary advantage to individuals. However, reliably passing on the key molecular machinery to new containers likely would have required specialized chemical processes. Specialization to that degree requires a division of labor amongst prebiotic molecules, and it is well known in evolutionary theory that cooperation is a prerequisite of division of labor[131]. For these reasons we suggest that in a prebiotic chemical soup, the rules of cooperation such as those we investigated here, would have been the primary drivers of evolutionary change.

Chapter 5

NOISY CHANNELS, ERROR CORRECTION AND THE ORIGIN OF LIFE

5.1 Abstract

Life on Earth consists of evolutionary lineages that represent the stable propagation of information through billions of years of biological, ecological, and planetary evolution. A central problem in the origins of life lies in explaining the origin of this stable transmission of information. Here we present a model of molecular replication, where replicators are coupled to a noisy environment. We demonstrate that when mutation rates depend on the concentration of active monomers, information can be reliably propagated in the face of noise. We introduce a simple notion of function and demonstrate that replicating templates which catalyze changes in their environment are able to stabilize those environments which in turn further stabilizes the transmission of genetic information. We conclude with a discussion of how these result inform our understanding of the origin of life on Earth.

5.2 Introduction

Life on Earth is composed of evolutionary lineages which represent the stable transmission of information through billions of years of biological, ecological and planetary evolution [21]. In contemporary organisms this information is transmitted through the copying of genetic information stored in DNA [3]. In order for genetic information to be reliably transmitted between generations, the mutation rate per base

pair must be less than the inverse of the amount of information (in bits) contained in the genome. This constraint, known as *Eigen's error threshold*, implies that high fidelity replication would be required to enable the emergence of complex genetic lineages [19]. Experiments suggest that DNA polymerases used across the tree of life have error rates ranging from 1 mutation in every $10^8 - 10^{12}$ base pairs [136], allowing modern life to maintain very large genomes. The chemical origins of this high fidelity replication are currently unknown [7, 55]. Decades of research suggest that the spontaneous synthesis of a general purpose polymerase which could mediate this high fidelity replication, would have been an extremely unlikely, if not impossible event [7, 55, 137].

It has been proposed that non-enzymatic replication may have enabled proto-biological systems to store and transmit information prior to the emergence of a polymerase [137–139]. While non-enzymatic replication is chemically simple compared to enzymatic polymerization, it presents a number of different challenges [137]. Here we focus on the fidelity of non-enzymatic replication, which is much lower than the enzymatic replication seen in modern organisms [136, 137]. It is currently unclear if the fidelity of non-enzymatic replication is sufficient to enable stable transmission of genetic information [137, 138].

Empirical studies have estimated the error rate of non-enzymatic replication of RNA templates to be about 17% [140]. Importantly, this error rate is sensitive to the relative concentrations of activated monomers and 17% is the minimum expected given optimal monomer concentrations [140]. Some mechanisms to effectively lower this mutation rate have been identified [137, 141, 142]. Rajamani *et al* demonstrated that non-enzymatic replication proceeds much slower after a mutation has occurred [141]. This stalling phenomena implies that replication events with errors are completed at

a much slower rate than those without errors, enabling the correct copies to serve as templates sooner than incorrect copies. The net effect of this phenomena is an effective error rate which is at least 2 times less than expected [141]. Other authors have studied the effect of alternative biochemistry on error rates [142, 143]. Many errors in nucleic acids templates are caused by the mis-incorporation of uracil bases in the place of cytosine bases, which have similar thermodynamic stability when paired with guanine [142]. Testa *et al* demonstrated that changing one base, from uracil to 2-thio-uracil in RNA templates altered the thermodynamic stability of those mutations, reducing their likelihood [142].

Here we characterize a previously unexplored a *system level* mechanism which can reduce the effective error rate of non-enzymatic replication and enable Darwinian evolution. This mechanism manifests due to the feed-backs between replicating templates, and environmental concentrations of activated monomers, which in turn effect mutation rates. We demonstrate this mechanism using a computational model in which templates are replicated, degraded and recycled back into activated monomers. We assume the mutation rates of templates depend on the relative concentrations of free monomers. We also assume that some replicating templates are functional, and therefore cause changes in their environmental conditions. Within the context of this model, the simplest possible function a template can perform is to convert one type of monomer into another.

Using this model we find that coupling mutation rates to monomer availability mitigates the effects of the error threshold. Specifically, in this model, the effects of the error threshold manifest *stochastically*. Given the same initial conditions and parameters, some simulations enable the reliable transmission of information through selection, while others do not. We show the stochastic aspect of this phenomena is

driven by the coupling of mutation rates to resource availability. In this model we find selection can occur at mutation rates that are at least 1.5 times higher than those expected from traditional error threshold arguments. We demonstrate how this coupling, in combination with functional templates can enable the stable propagation of information in the face of even higher mutation rates. Finally, we discuss the implications of these results for the understanding the early evolutionary dynamics of information bearing systems.

5.3 Model Description

In order to investigate these dynamics we developed a kinetic model of non-enzymatic replication in which templates are explicitly coupled to their environment. We consider a *finite* pool of monomers composed of A and B monomer species. Templates are formed by strings of monomers of length ν . The particular sequence of monomers in the string determines the species i of the template. We refer to the abundance of a particular species as x_i . We use a_i and b_i to refer to the number of A and B monomers in an individual of species i , such that for any i , $a_i + b_i = \nu$. We refer to the number of free A and B monomers in the environment as \mathcal{A} and \mathcal{B} respectively. Templates spontaneously decay into monomers with a propensity, Q_d , given by equation 5.1.

Replication reactions occur with a propensity that depends on the composition of both the template and the free monomers available in the environment. Recent laboratory experiments suggest that non-enzymatic replication rates are approximately first order in the free monomers available [144]. The propensity of replication, Q_r is given by equation 5.2, where k_{ri} gives the replication rate constant for species i and

G_i encodes the resource dependence of the templates, ensuring that A rich templates are more likely to replicate in A rich environments. Since we are primarily concerned with the fidelity of replication we assume that all monomers are activated readily by the environment and that completed strings immediately dissociate from the template making them available to serve as templates or to degrade immediately. We designate a *master sequence* which is considered *more fit* than the other templates. Replication reactions involving this template occur at a rate which is ten times higher than all other sequences. We assume that mutations occurs during the replication process and the propensity of mutation is sensitive to the free monomers in the environment, such that if the number of free A monomers is higher, the likelihood of mis-incorporating an A monomer is higher, similarly for B monomers. The simplest way to incorporate this effect is to assume that during a replication reaction each B monomer has a chance to be replaced $M_{BA} = \mu \frac{A}{A+B}$. According to traditional arguments the error threshold in this model should manifest for $M_{BA} = M_{AB} \approx \frac{1}{\nu}$, when the number of free monomers of each type is equal this would reduce to $\mu \approx \frac{2}{\nu}$.

We further assume that A monomers can be converted to B monomers and visa versa. This might represent, for example, the deamination of cytosine into uracil. This process can occur spontaneously and it can be catalyzed. While we recognize it is in general not possible to convert any nucleobase into any other, we include this feature of the model as to explore the role functional feedback between templates and the relative concentrations of activated monomers. The propensity of transforming A monomers into B monomers, Q_t , is assumed to follow equation 5.3, where x_c is the abundance of catalyst template and f_c represents the effect of that catalyst on the reaction rate. The same functional form applies for the propensity of transforming B monomers into A monomers [145]. Particular species must be identified as catalysts,

and assigned a corresponding catalytic rate constant, f_c . We simulate these dynamics using a kinetic Monte Carlo algorithm [103, 104]. In what follows we set $\nu = 8$, $k_d = 1.0$, $k_t = 10^{-6}$, and $k_{ri} = k_r = 5 * 10^{-5}$. The total mass was set such that there were initially 20,000 of A and 20,000 B monomers, with 10 of each template species in the system, unless otherwise specified.

$$Q_d = \sum_i k_d x_i \quad (5.1)$$

$$Q_r = \sum_i k_{ri} x_i G_i \quad (5.2)$$

$$Q_t = \left(k_t + \sum_c f_c x_c \right) \mathcal{A} \quad (5.3)$$

$$G_i = (\mathcal{A} + \mathcal{B})(1 - \epsilon)^{a_i}(1 + \epsilon)^{b_i} \quad (5.4)$$

$$\epsilon = \frac{\mathcal{B} - \mathcal{A}}{\mathcal{B} + \mathcal{A}} \quad (5.5)$$

5.4 Results

In order to characterize the dynamics of this model, parameters were sampled using a Monte Carlo approach. We first investigated the effect of the replication rate k_r and the mutation rate μ , with a fixed master sequence ($BBBBBBBB$), no functional templates, and a very slow monomer conversion rate $k_t = 10^{-6}$. 9000 simulations were run with the values of k_r sampled from a log-uniform distribution over the range of $(10^{-5}, 10^{-2})$ ¹, and the mutation rates sampled from the uniform distribution between zero and unity, ($\mu \sim \mathcal{U}(0, 1)$). The time-averaged total number of replicators for these

¹This is done by sampling a value from the uniform distribution between $[-5, -2]$, exponentiating the result such that $k_r \sim 10^{\mathcal{U}(-2, -5)}$

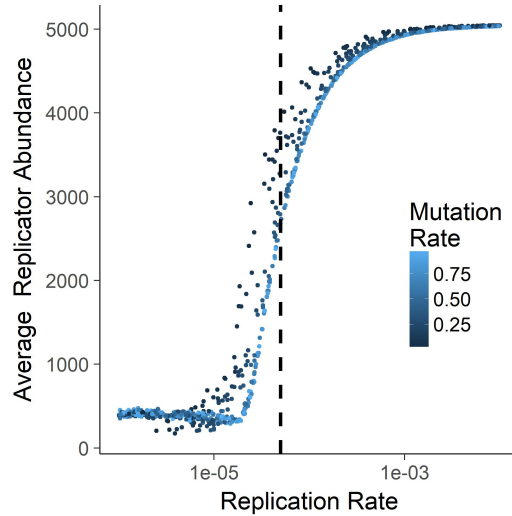


Figure 13. Average abundance of templates changes as a function of the replication rate k_r . The time-averaged, total template abundance is shown for different values of k_r and μ is shown. The value of $k_r = 5 * 10^{-5}$ was chosen for the following results, and is shown here with a vertical dashed line.

simulations is shown in figure 13. Based on these results we chose to use the value of $k_r = 5 * 10^{-5}$, for the remainder of the simulations because it enabled a wide variability in simulation outcomes.

The time series from a typical run with $k_r = 5 * 10^{-5}$ and $\mu \approx 0.230$ is shown in figure 14. The master sequence (shown in red) dominates the population, reaching a steady-state concentration of approximately 200. As predicted by quasi-species theory, a “cloud“ of mutants follows the population of the master sequence. Shown in grey is the averaged abundance of the all templates that are one mutation away from the master sequence, similarly the pink line shows the average abundance of all sequences which are two mutations away from the master.

In order to characterize the effect of the mutation rate μ on the dynamics, 1127 simulations were run with the replication rate fixed as $k_r = 5 * 10^{-5}$, with the mutation

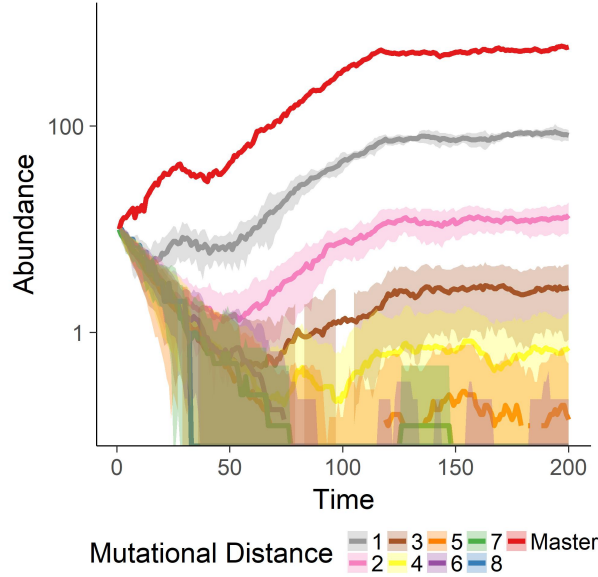


Figure 14. Time series of templates for a simulation with a single master sequence which has a factor of 10 greater replication rate than all other sequences. The colored line here correspond to the abundance different templates. The abundances are plotted on a logarithmic scale. The red line indicates the abundance of the master sequence at each time step. The grey line represents the *average* abundance of species which differ from the master by a single mutation, the pink line represents the *average* abundance of species which differ from the master by two mutations and so on.

rate sampled from the uniform distribution $\mu \sim \mathcal{U}(0, 1)$. We characterized these simulations based on whether or not *selection* for the master sequence occurred. We define selection to occur when the master sequence abundance both increases from its initial value and is greater than any other template in the system at the end of the simulation. In other models of molecular replication there is a mutation rate, known as the *error threshold*, above which the master sequence cannot be selected, despite an intrinsic fitness advantage [19]. Surprisingly we find that selection is possible above the predicted error threshold ($\mu_e \approx 0.25$) in this model. When the system is initialized with 10 of every species, we found that the master sequence was selected according to a probability that decreased with increasing mutation rates. Specifically in the range

$0.2 < \mu < 0.35$, there were many simulations which resulted in both possible outcomes, as shown in the top of figure 15. The bottom of figure 15, shows an estimate of this probability as a function of the mutation rate μ .

In order to understand the mechanisms responsible for this stochastic error threshold we compared our model with two other variants. First we compare it against a model in which "back mutations" are inhibited. Second we compare it against a model in which the mutation rate is constant. In the derivation of the error threshold it is assumed that no other sequence mutates into the master sequence such that mutations only serve to decrease the population of the master sequence [19]. Since the effect of back mutations is present in our model, we tested whether it could be responsible for the stochastic aspects of the error threshold. We ran simulations which were exactly the same except that any mutation which would've created a copy of the master sequence was inhibited. The results of that model were statistically indistinguishable from our results (not shown), implying that back mutations are not responsible for the stochastic error threshold observed in this model.

Another major distinction between the model presented here and others is the explicit coupling of the mutation rate to a finite and dynamic environment. In order to test whether this difference was driving the stochastic the error threshold in this model we ran simulations in which the mutation rate was fixed at the same value of $M = \frac{\mu}{2}$. This ensured the likelihood of mutations were constant with respect to the resource availability, with a value that was equal to the scenario when both monomers are equally abundant. We found that when the mutation rate is not coupled to a changing environment the master sequence could not be selected above $\mu \approx 0.31$. This is different than the original model, where many simulations showed the master sequence fixating small, but dominate populations at higher mutation rates.

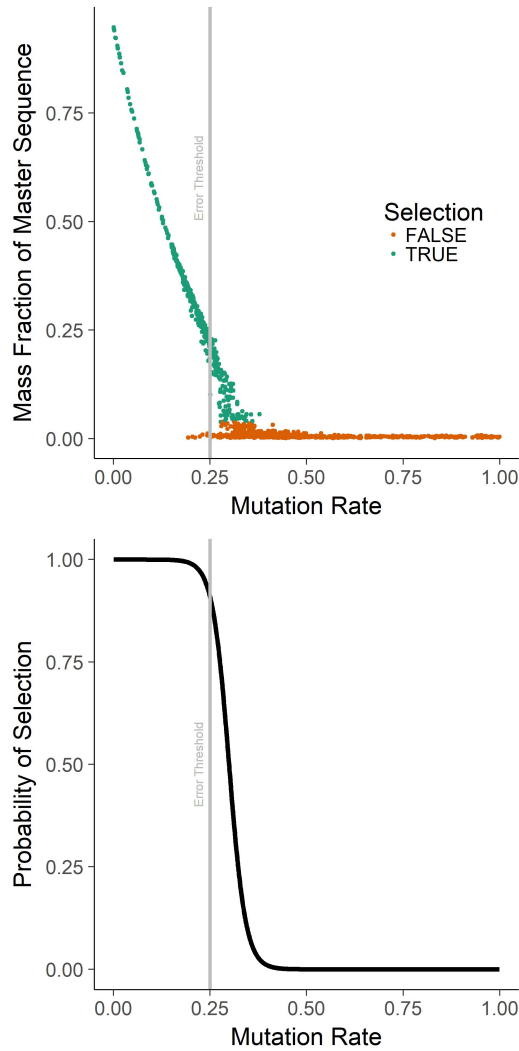


Figure 15. Selection of the master sequence occurs stochastically in this model. Top: The mass fraction is the time averaged ratio between the mass in the master sequence against the mass in all templates. The mass fraction of the master sequence for 1127 simulations is shown as a function of the mutation rate. Points are colored based on whether or not selection occurred in that simulation (see main text). Bottom: The estimated probability of selection occurring is shown as a function of the mutation rate μ (black line). We find that in contrast to previous models of molecular replication, the error threshold in this model does not represent a strict limitation. Selection of the master sequence occurs with a probability that decreases as the mutation rate increases. Selection can occur above the error threshold.

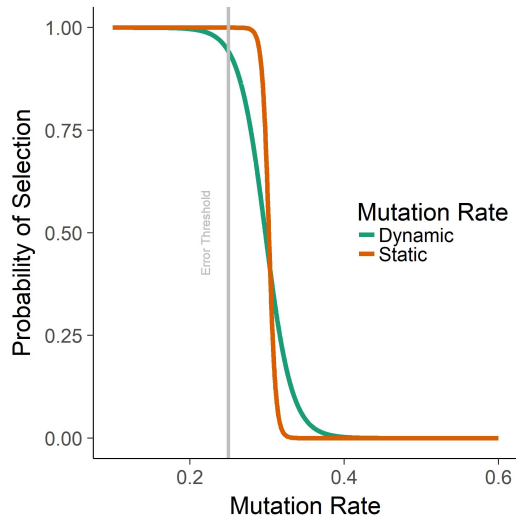


Figure 16. The probability of the master sequence fixating is shown as a function of the mutation rate μ , for two different models. In orange, the probability for a model with a static mutation rate is shown. In this case the error threshold is a sharp transition, as expected from quasi-species theory, when the mutation rate is above a critical value the master sequence cannot be selected. In green, the same probability is shown a model with a dynamic mutation rate. The model with a static mutation rate has a sharp error threshold, while the model with a dynamic mutation rate has a stochastic error threshold.

We next sought to determine the highest possible mutation rate at which the master sequence could be fixated in the original model. In order to do this we initialized the system with different initial concentrations. We used the previous simulation runs to help select new concentrations. We took the final concentrations from runs in which the master sequence had fixated. We added some small noise to those final concentrations by randomly increasing the concentration of some species and reducing the concentration of other species, while maintaining the total mass of the system. We then evolved these perturbed concentrations at higher mutation rates than the simulation which produced the initial conditions. This process was done iteratively. As the mutation rate increased fewer and fewer initial conditions led to the master sequence fixating. Further, the absolute number of templates also decreased. The

highest mutation rate with a stable population of the master sequence was observed to be $\mu = 0.41$.

We next investigated the effect of introducing functional templates. We ran 1000 simulations with the mutation rate fixed at $\mu = 0.4$, and the master sequence (*BBBBBBBB*) was designated as a catalyst for the transformation of *A* monomers into *B* monomers. The remaining parameters were the same as before. The catalytic rate constant was sampled from a uniform distribution $f_c \sim \mathcal{U}(0.0, 0.01)$. As the catalytic rate constant increased the probability of selection increased, as shown on the left side Figure 17, this analysis was repeated for higher mutation rates up to $\mu = 0.7$. As the mutation rate increases, the catalytic rate constant needed to be higher in order to ensure the same effect.

We next sought to determine whether the same effect could be achieved when the master sequence was not itself functional, but rather when it's mutational neighbors were functional. Once again we fixed the master sequence to be *BBBBBBBB*, but instead assigned the a functional template to be either *ABBBBBBB*, *AABBBBBB* or *AAABBBBB*, with the value of f_c drawn from the same distribution. We find that even when the master sequence is not functional itself selection can occur as the functional template is populated in the “cloud” of the master sequence and effectively biases the environmental conditions. Figure 17 right shows the result of this analysis. As the mutational distance between the master sequence and the functional template increases, greater catalytic rate constants are required to ensure selection.

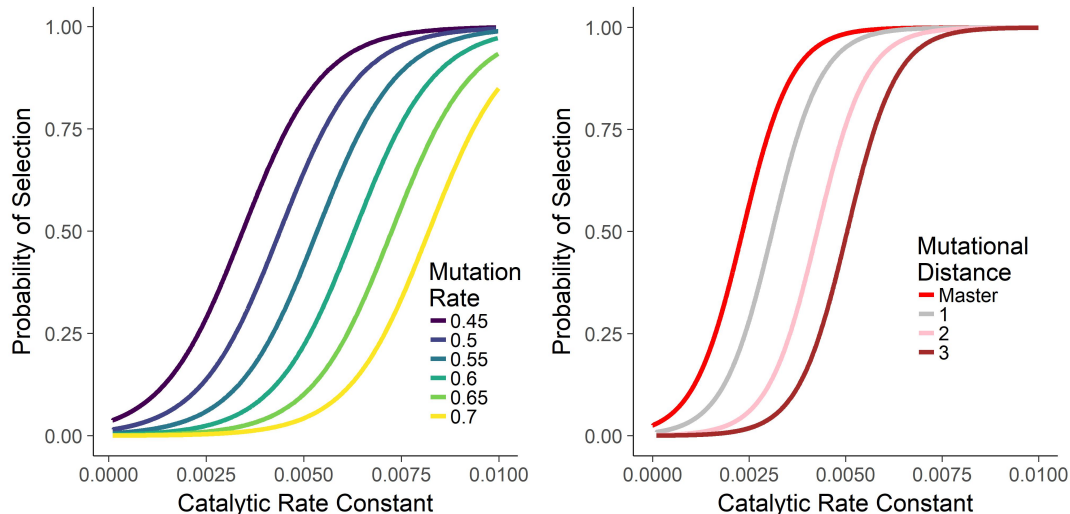


Figure 17. Functional templates which catalyze changes in the environment can enable selection in the face of high mutation rates. The probability of selection is shown as a function of the catalytic rate constant associated with functional templates. Left: The master sequence ($BBBBBBB$) is also a catalyst which converts A monomers into B monomers. As the catalytic rate constant f_c , is increased the probability of selection increases. For higher mutation rates, greater catalytic function is needed to achieve selection. Right: The master sequence can be different. Here the probability of selection is shown the cases where the catalyst is the same as the master sequence, as well as when the catalyst is a mutant of the master sequence (see text for particular sequences).

5.5 Discussion

We used a chemical kinetic model to explore the fidelity non-enzymatic replication. This model has two key differences when compared with previous models. First, it directly couples mutation rates to resource abundances. This modeling decision is motivated by empirical observations [140] and it has dynamical consequences. Specifically it mitigates the effect of the error threshold, enabling the reliable transmission of information at higher mutation rates than found without such a coupling. In closed systems such as those explored here this coupling induces a feedback between

replicating templates and the environment. The feedback means that the dynamics of templates and activated monomers mutually constrain each other.

This model also incorporates one of the essential features of life on Earth, its capacity to modify its environment [88, 146]. By including catalytic feed-back that cause templates to bias the resources available in the environments the templates modify their own dynamics. When the master sequence catalyzes the environmental conditions that favor itself, it can dynamically reduce its own error rate and enable the stable transmission of information in the face of high mutation rates. This effect can even be realized when the catalyst is not the master sequence itself but rather a different template in the mutational neighborhood of the master sequence.

The system level mechanisms explored here reduce the effect of the error threshold. Importantly, they do not preclude previously characterized mechanisms to improve the fidelity of non-enzymatic replication [141]. When taken together with previously explored mechanisms it is reasonable to assume that non-enzymatic replication may achieve the fidelity required to propagate genetic information [137]. However, due to the coupling between mutation rates and monomer concentrations, the fidelity required to enable selection may have emerged dynamically. Early in prebiotic evolution, low fidelity replication would have enabled faster exploration of sequence space. The emergence of functional templates would have stabilized environment conditions, reducing the mutation rates and thereby allowing selection on that function.

These processes suggest a multilevel view of evolutionary dynamics similar to those in [147]. We demonstrated how selection between templates can occur when templates are coupled to a dynamic environment. However that same coupling implies that the evolutionary fate of templates depends in part on the other templates in its mutational neighborhood. In a real system its possible that there could be competition not just

between individual templates, but also between mutational neighborhoods, where some particularly fit templates are at a disadvantage due to being further away from functional sequences. We leave the dynamics of such competition as a topic for future study.

Finally, we point out that finite and closed systems, such as those studied here, have different properties compared to infinite/open systems, which are more typically modelled using differential equations. In this case, this distinction has dynamical consequences. We believe that the finite aspect of this model enables it to more closely approximate laboratory systems, such as [139], as well as being more interesting theoretically [99]. To make progress towards understanding the origin of life on Earth, a deeper integration between computational modeling and laboratory experiments will be required.

Chapter 6

CONCLUSION

The title of this dissertation is *On the Origin of the Living State*. As you may have guessed, this dissertation does not, in fact, explain the origin of life on Earth. So what was it all for? First and foremost this work demonstrates how computational modeling can improve our understanding of chemical evolution. In chapters 3 and 5 I used computational models of molecular replication to explore dynamic chemical systems. In those studies, simple rules interacted to generate complex dynamics. In chapter 3, I demonstrated how information measures can be used to track those complex dynamics through a phase transition [88]. The model used in chapter 5 illustrated important features that have been completely missed in models based on ordinary differential equations [17, 19]. Those dynamics suggest that the origin of genetic information may be easier to explain than previously expected [137, 138]. As we saw in chapter 4, computational models can be particularly powerful when combined with empirical information [111]. Using that model I was able to test alternative chemical scenarios to determine what features of the real microscopic system are responsible for the macroscopic patterns observed in the lab. This is a direct example of the theoretical framework discussed in chapter 2, it demonstrates how quantitatively different microscopic dynamics can manifest in qualitatively different macroscopic properties. These computational techniques enable us to ask new questions which would be impossible in the lab and intractable on the whiteboard. Developing new scientific questions represents significant progress in its own right.

As I've discussed before the origin of life is perhaps the greatest unsolved mystery

in the history of science. Philosophers and theologians pondered this question for thousands of years, even before modern scientific traditions emerged. Questions about the nature of life, and the origin of biological processes are, and always have been related to the fundamental questions of existence. These questions do not fall into the disciplinary distinctions that define university departments and scientific journals, rather they manifest in different ways through many disciplines, including geology, chemistry, biology, physics and astronomy. Unfortunately, for years the question of life's origin was assumed to be a strictly chemical question. The discovery of DNA by Watson and Crick convinced scientists the essential features of living systems were molecular in nature. For decades that notion dominated studies of the origin of life. Science based on that idea has ultimately provided detailed knowledge of life's building blocks, but no explanation for how and why those building blocks came to form life on Earth.

In the past two decades new discoveries have shown us that many properties of living systems are not chemical in nature [21, 59]. For example, the molecular basis of genetics led many evolutionary biologists to posit that “the genes are in charge[148].” This implies that all of biological phenomena can be understood as a consequence of ~ 4 billion years of selection on DNA sequences, and that genes are the ultimate cause of biological processes. The discovery of genetic regulatory networks up-ended this idea. If some genes can regulate others, why assume that any one gene is the cause of a particular phenomena? This discovery and many other advances have forced biologists adopt a *systems* view [59].

This systems perspective is driving a radical reshaping the discourse related to the origins of life [21, 23, 59, 147]. One the of the key consequences of this perspective is that having the chemical building blocks of a living system is not equivalent to

having a living system, *e.g.* having the stuff is not enough. This is one of the motivating ideas behind viewing life as a state of matter, as explained in chapter 2 and exemplified by the dynamics shown in chapter 3. Researchers are just now beginning to investigate the organizational and informational features of living systems. In chapter 3 I demonstrated how information processing distinguishes the living state from the non-living state, and in chapter 4 I demonstrated how some organizational features of populations emerge as a consequence of the properties of molecules. Future research will be needed to understand what features of the living state “come for free,” based on the molecular building blocks, which features are highly evolved due ~ 4 billion years of evolutionary optimization, and which features are due to currently unknown organizational laws [5, 23]. In that effort computational models, like the kind used here, will be invaluable.

Advances in computation have reshaped the very nature of the scientific method [149]. Computational approaches to the natural sciences have emerged in the space between experimentation and theory. They are simultaneously used to deduce the consequences of theories and analyze the outcomes of experiments. Computation has even disrupted the fundamental notion, dating back to the scientific revolution and Galileo, that the language of the universe is mathematics. In the place of mathematical equations, science today is based on discovering, and mimicking the algorithms of nature. As others have remarked, the rise of computation will almost certainly be viewed as a revolution in the history of science [149]. As I’ve shown here, this revolution will undoubtedly contribute to understanding the origin of life, but another revolution may be required before the problem can be solved.

New technologies are fueling radical new approaches to scientific questions. Information technology, machine learning and 3D printing are putting more power

in the hands of every scientist, enabling individual researchers to pursue questions that would have required an entire lab just 10 years ago. Meanwhile, improvements in communication technology are facilitating deeper levels of collaboration between individuals as well as entire communities of scientists. However, in spite of these advances, scientific institutions and practices have remained largely unchanged. With few exceptions, universities are still organized around disciplines, rather than motivating questions, meaning they are organized according to the problems we've already solved, rather than the ones we haven't. Academic publications are still targeted to a relatively narrow set of experts, obscuring critical results behind a haze of specialist knowledge. This means their target audiences are people who already understand the problem, rather than people who don't. Taken together these institutional factors create barriers to productive scientific discourse. These barriers are stymieing progress on the origin of life, and overcoming them will be tantamount to a scientific revolution.

This dissertation is rather narrowly focused on computational models of prebiotic chemistry, with some brief interludes from the world of statistical physics. Nonetheless this dissertation may be considered highly interdisciplinary in its scope. The very core of this work is the bridging together of concepts, ideas, hypotheses and facts from many different disciplines into a coherent narrative. In that sense, this dissertation represents an important step towards solving the origin of life, in that it is the bringing together of many disciplinary ideas. The success of future origin of life research will depend on the ability of scientists to communicate and collaborate across institutional barriers.

One of the most dramatic divides in the origin of life community is between the world of experiment and theory. Most models of chemical evolution, such as the ones explored in chapters 3 and 5 have been developed by physicists and theoretical

biologists, and they have been poorly constrained or informed by empirical facts. Luckily, new experimental paradigms are emerging, which leverage advances in analytic chemistry, artificial intelligence and robotics. These technologies are enabling new experimental platforms which can generate vast amounts of data. Characterizing the outputs of these experiments will require a new class of models. The development and refinement of those models represents the next frontier in solving the mystery of the origin of life.

NOTES

REFERENCES

- [1] Kenzo Nakamura, Particle Data Group, et al. “Review of particle physics”. In: *Journal of Physics G: Nuclear and Particle Physics* 37.7A (2010), p. 075021.
- [2] Pisin Chen. “Recent Progress in Cosmology and Particle Astrophysics”. In: *arXiv preprint arXiv:1310.1107* (2013).
- [3] SE Bresler. *Introduction to molecular biology*. Elsevier, 2012.
- [4] Geoffrey West. “Scale: The Universal Laws of Growth, Innovation, Sustainability, and the Pace of Life”. In: *Organisms, Cities, Economies, and Companies* (2017).
- [5] Sara Imari Walker. “Origins of life: a problem for physics, a key issues review”. In: *Reports on Progress in Physics* 80.9 (2017), p. 092601.
- [6] Caleb Scharf et al. “A Strategy for Origins of Life Research”. In: *Astrobiology* 15.12 (2015), pp. 1031–1042.
- [7] Steven A Benner. “Paradoxes in the Origin of Life”. In: *Origins of Life and Evolution of Biospheres* 44.4 (2014), pp. 339–343.
- [8] Antonio Lazcano. “Historical development of origins research”. In: *Cold Spring Harbor perspectives in biology* 2.11 (2010), a002089.
- [9] Rudolf Stichweh. *Differentiation of scientific disciplines: causes and consequences*. 2003.
- [10] Thomas S Kuhn. *The structure of scientific revolutions*. University of Chicago press, 2012.
- [11] Nicolas Glansdorff, Ying Xu, and Bernard Labedan. “The last universal common ancestor: emergence, constitution and genetic legacy of an elusive forerunner”. In: *Biology direct* 3.1 (2008), p. 1.
- [12] Nick Lane, John F Allen, and William Martin. “How did LUCA make a living? Chemiosmosis in the origin of life”. In: *BioEssays* 32.4 (2010), pp. 271–280.
- [13] John F. Atkins, Raymond F. Gesteland, and Thomas R. Cech, eds. *The RNA World, Third Edition*. Cold Spring Harbor, 2005.

- [14] Orgel Leslie E. “Prebiotic chemistry and the origin of the RNA world”. In: *Critical reviews in biochemistry and molecular biology* 39.2 (2004), pp. 99–123. DOI: 10.1080/10409230490460765.
- [15] Everett L Shock and Eric S Boyd. “Principles of geobiochemistry”. In: *Elements* 11.6 (2015), pp. 395–401.
- [16] Addy Pross. “Toward a general theory of evolution: extending Darwinian theory to inanimate matter”. In: *Journal of Systems Chemistry* 2.1 (2011), p. 1. DOI: 10.1186/1759-2208-2-1.
- [17] Martin A Nowak and Hisashi Ohtsuki. “Prevolutionary dynamics and the origin of evolution”. In: *Proceedings of the National Academy of Sciences* 105.39 (2008), pp. 14924–14927. DOI: 10.1073/pnas.0806714105.
- [18] Sara Imari Walker, Martha A Grover, and Nicholas V Hud. “Universal sequence replication, reversible polymerization and early functional biopolymers: a model for the initiation of prebiotic sequence evolution”. In: *PLoS One* 7.4 (2012), e34166.
- [19] Manfred Eigen. “Natural selection: a phase transition?” In: *Biophysical chemistry* 85.2–3 (2000), pp. 101–123. DOI: 10.1016/S0301-4622(00)00122-8.
- [20] Jeremy L England. “Statistical physics of self-replication”. In: *The Journal of chemical physics* 139.12 (2013), p. 121923. DOI: 10.1063/1.4818538.
- [21] Eric Smith and Harold J Morowitz. *The Origin and Nature of Life on Earth: The Emergence of the Fourth Geosphere*. Cambridge University Press, 2016.
- [22] Sara Imari Walker. “Top-Down Causation and the Rise of Information in the Emergence of Life”. In: *Information* 5.3 (2014), pp. 424–439. DOI: 10.3390/info5030424.
- [23] Sara Imari Walker and Paul CW Davies. “The algorithmic origins of life”. In: *Journal of the Royal Society Interface* 10.79 (2013), p. 20120869.
- [24] Christoph Adami. “Information-theoretic considerations concerning the origin of life”. In: *Origins of Life and Evolution of Biospheres* 45.3 (2015), pp. 309–317.
- [25] Chiara Marletto. “Constructor theory of life”. In: *Journal of The Royal Society Interface* 12.104 (2015), p. 20141226.
- [26] Peter Godfrey-Smith. *Theory and reality: An introduction to the philosophy of science*. University of Chicago Press, 2009.

- [27] Paul Feyerabend. *Against method*. Verso, 1993.
- [28] Ansel Payne. “Why Physics Is Not a Discipline, Physics is not just what happens in the Department of Physics.” In: ().
- [29] Carol E Cleland. “Methodological and epistemic differences between historical science and experimental science”. In: *Philosophy of Science* 69.3 (2002), pp. 447–451.
- [30] David C Krakauer. “Darwinian demons, evolutionary complexity, and information maximization”. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 21.3 (2011), p. 037110.
- [31] Stephen J Freeland and Laurence D Hurst. “The genetic code is one in a million”. In: *Journal of molecular evolution* 47.3 (1998), pp. 238–248.
- [32] John P DeLong et al. “Shifts in metabolic scaling, production, and efficiency across major evolutionary transitions of life”. In: *Proceedings of the National Academy of Sciences* 107.29 (2010), pp. 12941–12945.
- [33] David C Krakauer et al. “The challenges and scope of theoretical biology”. In: *Journal of theoretical biology* 276.1 (2011), pp. 269–276.
- [34] Marc Tessera. “Origin of evolution versus origin of life: a shift of paradigm”. In: *International journal of molecular sciences* 12.6 (2011), pp. 3445–3458.
- [35] Lucas John Mix. “Defending definitions of life”. In: *Astrobiology* 15.1 (2015), pp. 15–19.
- [36] Serhiy A Tsokolov. “Why is the definition of life so elusive? Epistemological considerations”. In: *Astrobiology* 9.4 (2009), pp. 401–412.
- [37] Leroy Cronin et al. “The imitation game—a computational chemical approach to recognizing life”. In: *Nature biotechnology* 24.10 (2006), pp. 1203–1206.
- [38] Erwin Schrödinger. *What is life?: With mind and matter and autobiographical sketches*. Cambridge University Press, 1992.
- [39] Eörs Szathmáry and John Maynard Smith. “From replicators to reproducers: the first major transitions leading to life”. In: *Journal of Theoretical Biology* 187.4 (1997), pp. 555–571. DOI: 10.1006/jtbi.1996.0389.

- [40] Eörs Szathmáry. “The origin of replicators and reproducers”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 361.1474 (2006), pp. 1761–1776. DOI: 10.1098/rstb.2006.1912.
- [41] Martin A Nowak and Karl Sigmund. “Evolutionary dynamics of biological games”. In: *Science* 303.5659 (2004), pp. 793–799.
- [42] Hisashi Ohtsuki and Martin A Nowak. “Prelife catalysts and replicators”. In: *Proceedings of the Royal Society B: Biological Sciences* 276.1674 (2009), pp. 3783–3790. DOI: 10.1098/rspb.2009.1136.
- [43] Meng Wu and Paul G Higgs. “The origin of life is a spatially localized stochastic transition”. In: *Biology direct* 7.1 (2012), p. 42. DOI: 10.1186/1745-6150-7-42.
- [44] Alyssa M Adams et al. “Formal Definitions of Unbounded Evolution and Innovation Reveal Universal Mechanisms for Open-Ended Evolution in Dynamical Systems”. In: *arXiv preprint arXiv:1607.01750* (2016).
- [45] John von Neumann. *Theory of self-reproducing automata*. Ed. by Arthur W Burks. University of Illinois Press, 1966.
- [46] Edith M Sevick et al. “Fluctuation theorems”. In: *arXiv preprint arXiv:0709.3888* (2007).
- [47] Gavin E Crooks. “Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences”. In: *Physical Review E* 60.3 (1999), p. 2721.
- [48] Eric Smith. “Large-deviation principles, stochastic effective actions, path entropies, and the structure and meaning of thermodynamic descriptions”. In: *Reports on Progress in Physics* 74.4 (2011), p. 046601.
- [49] Hugo Touchette. “The large deviation approach to statistical mechanics”. In: *Physics Reports* 478.1 (2009), pp. 1–69.
- [50] Georgiy P Karev. “Replicator equations and the principle of minimal production of information”. In: *Bulletin of mathematical biology* 72.5 (2010), pp. 1124–1142.
- [51] Edwin T Jaynes. “Information theory and statistical mechanics”. In: *Physical review* 106.4 (1957), p. 620.
- [52] Roderick C Dewar. “Maximum entropy production as an inference algorithm that translates physical assumptions into macroscopic predictions: Don’t shoot the messenger”. In: *Entropy* 11.4 (2009), pp. 931–944.

- [53] LM Martyushev and VD Seleznev. “Maximum entropy production principle in physics, chemistry and biology”. In: *Physics reports* 426.1 (2006), pp. 1–45.
- [54] Edwin T Jaynes. “On the rationale of maximum-entropy methods”. In: *Proceedings of the IEEE* 70.9 (1982), pp. 939–952.
- [55] Paul G Higgs and Niles Lehman. “The RNA World: molecular cooperation at the origins of life”. In: *Nature Reviews Genetics* 16.1 (2015), pp. 7–17.
- [56] Vera Vasas et al. “Evolution before genes”. In: *Biology Direct* 7.1 (2012), p. 1.
- [57] Philippe Nghe et al. “Prebiotic network evolution: six key parameters”. In: *Molecular BioSystems* 11.12 (2015), pp. 3206–3217.
- [58] Eric Smith and Harold J Morowitz. “Universality in intermediary metabolism”. In: *Proceedings of the National Academy of Sciences of the United States of America* 101.36 (2004), pp. 13168–13173.
- [59] Stuart A Kauffman. *The origins of order: Self organization and selection in evolution*. Oxford University Press, USA, 1993.
- [60] Stuart A Kauffman. “Autocatalytic sets of proteins”. In: *Journal of theoretical biology* 119.1 (1986), pp. 1–24.
- [61] Wim Hordijk, Stuart A Kauffman, and Mike Steel. “Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems”. In: *International journal of molecular sciences* 12.5 (2011), pp. 3085–3101.
- [62] Wim Hordijk et al. “An investigation into irreducible autocatalytic sets and power law distributed catalysis”. In: *Natural Computing* 13.3 (2014), pp. 287–296.
- [63] Mike Steel. “The emergence of a self-catalysing structure in abstract origin-of-life models”. In: *Applied Mathematics Letters* 13.3 (2000), pp. 91–95.
- [64] Vera Vasas, Eörs Szathmáry, and Mauro Santos. “Lack of evolvability in self-sustaining autocatalytic networks constraints metabolism-first scenarios for the origin of life”. In: *Proceedings of the National Academy of Sciences* 107.4 (2010), pp. 1470–1475.
- [65] Omer Markovitch and Doron Lancet. “Excess mutual catalysis is required for effective evolvability”. In: *Artificial life* 18.3 (2012), pp. 243–266.

- [66] Wim Hordijk, Mike Steel, and Stuart Kauffman. “The structure of autocatalytic sets: Evolvability, enablement, and emergence”. In: *Acta biotheoretica* 60.4 (2012), pp. 379–392.
- [67] Jens Christian Claussen and Arne Traulsen. “Non-Gaussian fluctuations arising from finite populations: Exact results for the evolutionary Moran process”. In: *Physical review E* 71.2 (2005), p. 025101.
- [68] Nigel Goldenfeld and Carl Woese. “Life is physics: evolution as a collective phenomenon far from equilibrium”. In: *arXiv preprint arXiv:1011.4125* (2010).
- [69] Kalin Vetsigian, Carl Woese, and Nigel Goldenfeld. “Collective evolution and the genetic code”. In: *Proceedings of the National Academy of Sciences* 103.28 (2006), pp. 10696–10701.
- [70] Leroy Cronin and Sara Imari Walker. “Beyond prebiotic chemistry”. In: *Science* 352.6290 (2016), pp. 1174–1175.
- [71] Jan Spitzer, Gary J Pielak, and Bert Poolman. “Emergence of life: Physical chemistry changes the paradigm”. In: *Biology direct* 10.1 (2015), p. 1.
- [72] Sara I Walker et al. “Exoplanet Biosignatures: Future Directions”. In: *arXiv preprint arXiv:1705.08071* (2017).
- [73] Albert SZENT-GYORGYI. “THE STUDY OF ENERGY-LEVELS IN BIO-CHEMISTRY”. In: *Nature* 148.3745 (1941), p. 157.
- [74] Albert Szent-Gyorgyi. “Towards a new biochemistry?” In: *Science* 93.2426 (1941), pp. 609–611.
- [75] Lillian Hoddeson et al. *Out of the crystal maze: chapters from the history of solid state physics*. Oxford University Press, 1992.
- [76] Albert Szent-Gyorgyi. “The living state and cancer”. In: *Proceedings of the National Academy of Sciences* 74.7 (1977), pp. 2844–2847.
- [77] A Szent-Gyorgyi. “Introduction to a Submolecular Biology, Acad”. In: *Press, NY* (1960).
- [78] Gábor Vattay et al. “Quantum criticality at the origin of life”. In: *Journal of Physics: Conference Series*. Vol. 626. 1. IOP Publishing. 2015, p. 012023.

- [79] Daniel J Nicholson. “Biological atomism and cell theory”. In: *Studies in history and philosophy of science part c: Studies in history and philosophy of biological and biomedical sciences* 41.3 (2010), pp. 202–211.
- [80] Greer A Dolby et al. “Assessing the geological and climatic forcing of biodiversity and evolution surrounding the Gulf of California”. In: *Journal of the Southwest* 57.2 (2015), pp. 391–455.
- [81] Tran T Huynh and Christopher J Poulsen. “Rising atmospheric CO₂ as a possible trigger for the end-Triassic mass extinction”. In: *Palaeogeography, Palaeoclimatology, Palaeoecology* 217.3-4 (2005), pp. 223–242.
- [82] Charles H Lineweaver, Yeshe Fenner, and Brad K Gibson. “The galactic habitable zone and the age distribution of complex life in the Milky Way”. In: *Science* 303.5654 (2004), pp. 59–62.
- [83] Carlo Cercignani. *Ludwig Boltzmann: the man who trusted atoms*. Oxford University Press, 1998.
- [84] Nigel Goldenfeld. *Lectures on phase transitions and the renormalization group*. Addison-Wesley, Advanced Book Program, Reading, 1992.
- [85] Leo P Kadanoff. *Statistical physics: statics, dynamics and renormalization*. World Scientific Publishing Company, 2000.
- [86] Leo P Kadanoff. “More is the same; phase transitions and mean field theories”. In: *Journal of Statistical Physics* 137.5-6 (2009), p. 777.
- [87] Hyunju Kim et al. “Universal scaling across biochemical networks on Earth”. In: *bioRxiv* (2018), p. 212118.
- [88] Cole Mathis, Tanmoy Bhattacharya, and Sara Imari Walker. “The Emergence of Life as a First-Order Phase Transition”. In: *Astrobiology* 17.3 (2017), pp. 266–276.
- [89] Paul G Falkowski, Tom Fenchel, and Edward F Delong. “The microbial engines that drive Earth’s biogeochemical cycles”. In: *science* 320.5879 (2008), pp. 1034–1039.
- [90] Rogier Braakman, Michael J Follows, and Sallie W Chisholm. “Metabolic evolution and the self-organization of ecosystems”. In: *Proceedings of the National Academy of Sciences* 114.15 (2017), E3091–E3100.

- [91] Yinon M Bar-On, Rob Phillips, and Ron Milo. “The biomass distribution on Earth”. In: *Proceedings of the National Academy of Sciences* (2018), p. 201711842.
- [92] Peter Godfrey-Smith. *Philosophy of biology*. Princeton University Press, 2013.
- [93] Addy Pross. “On the emergence of biological complexity: life as a kinetic state of matter”. In: *Origins of Life and Evolution of Biospheres* 35.2 (2005), pp. 151–166. DOI: 10.1007/s11084-005-5272-1.
- [94] Michael Manapat et al. “Originator dynamics”. In: *Journal of theoretical biology* 256.4 (2009), pp. 586–595. DOI: 10.1016/j.jtbi.2008.10.006.
- [95] Meng Wu and Paul G Higgs. “Origin of Self-replicating Biopolymers: Auto-catalytic Feedback can Jump-start the RNA World”. In: *Journal of Molecular Evolution* 69.5 (2009), pp. 541–554. DOI: 10.1007/s00239-009-9276-8.
- [96] Bettina E Schirrmeister, Muriel Gugger, and Philip CJ Donoghue. “Cyanobacteria and the Great Oxidation Event: evidence from genes and fossils”. In: *Palaeontology* 58.5 (2015), pp. 769–785.
- [97] G.A. King. “Recycling, reproduction, and life’s origins”. In: *Biosystems* 15.2 (1982), pp. 89–97. DOI: 10.1016/0303-2647(82)90022-3.
- [98] G.A. King. “Was there a prebiotic soup?” In: *Journal of Theoretical Biology* 123.4 (1986), pp. 493–498. DOI: 10.1016/S0022-5193(86)80216-8.
- [99] David C Krakauer and Akira Sasaki. “Noisy clues to the origin of life”. In: *Proceedings of the Royal Society of London. Series B: Biological Sciences* 269.1508 (2002), pp. 2423–2428. DOI: 10.1098/rspb.2002.2127.
- [100] Jay T. Goodwin and David G. Lynn. “Template-directed synthesis: use of a reversible reaction”. In: *Journal of the American Chemical Society* 114.23 (1992), pp. 9197–9198. DOI: 10.1021/ja00049a067.
- [101] Nilesh Vaidya, Sara Imari Walker, and Niles Lehman. “Recycling of informational units leads to selection of replicators in a prebiotic soup”. In: *Chemistry & biology* 20.2 (2013), pp. 241–252. DOI: 10.1016/j.chembiol.2013.01.007.
- [102] Jay G. Forsythe et al. “Ester-Mediated Amide Bond Formation Driven by Wet–Dry Cycles: A Possible Path to Polypeptides on the Prebiotic Earth”. In: *Angewandte Chemie International Edition* (2015), n/a–n/a. DOI: 10.1002/anie.201503792. URL: <http://dx.doi.org/10.1002/anie.201503792>.

- [103] Daniel T Gillespie. “Exact stochastic simulation of coupled chemical reactions”. In: *The journal of physical chemistry* 81.25 (1977), pp. 2340–2361. DOI: 10.1021/j100540a008.
- [104] Daniel T Gillespie. “A general method for numerically simulating the stochastic time evolution of coupled chemical reactions”. In: *Journal of computational physics* 22.4 (1976), pp. 403–434. DOI: 10.1016/0021-9991(76)90041-3.
- [105] Peter Szabo et al. “In silico simulations reveal that replicators with limited dispersal evolve towards higher efficiency and fidelity”. In: *Nature* 420.6913 (2002), pp. 340–343. DOI: 10.1038/nature01187.
- [106] Jan P Amend et al. “The energetics of organic synthesis inside and outside the cell”. In: *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 368.1622 (2013), p. 20120255.
- [107] C. Adami. *Information-theoretic considerations concerning the origin of life*. 2014. arXiv: 1409.0590.
- [108] George FR Ellis. “Top-down causation and emergence: some comments on mechanisms”. In: *Interface Focus* 2.1 (2012), pp. 126–140.
- [109] Stephen Jay Gould and Niles Eldredge. “Punctuated equilibria: the tempo and mode of evolution reconsidered”. In: *Paleobiology* 3.2 (1977), pp. 115–151. URL: <http://www.jstor.org/stable/2400177>.
- [110] Imre Csiszar and János Körner. *Information theory: coding theorems for discrete memoryless systems*. Cambridge University Press, 2011.
- [111] Cole Mathis et al. “Prebiotic RNA Network Formation: A Taxonomy of Molecular Cooperation”. In: *Life* 7.4 (2017), p. 38.
- [112] Martin A Nowak. “Five rules for the evolution of cooperation”. In: *science* 314.5805 (2006), pp. 1560–1563.
- [113] Martin Nowak and Roger Highfield. *Supercooperators: Altruism, evolution, and why we need each other to succeed*. Simon and Schuster, 2011.
- [114] Irene A Chen and Martin A Nowak. “From prelife to life: How chemical kinetics become evolutionary dynamics”. In: *Accounts of chemical research* 45.12 (2012), pp. 2088–2096.
- [115] Katrin Bohl et al. “Evolutionary game theory: molecules as players”. In: *Molecular BioSystems* 10.12 (2014), pp. 3066–3074.

- [116] Jessica AM Yeates et al. “Dynamics of prebiotic RNA reproduction illuminated by chemical game theory”. In: *Proceedings of the National Academy of Sciences* 113.18 (2016), pp. 5030–5035.
- [117] Eric J Hayden and Niles Lehman. “Self-assembly of a group I intron from inactive oligonucleotide fragments”. In: *Chemistry & Biology* 13.8 (2006), pp. 909–918.
- [118] Jessica AM Yeates, Philippe Nghe, and Niles Lehman. “Topological and thermodynamic factors that influence the evolution of small networks of catalytic RNA species”. In: *RNA* 23.7 (2017), pp. 1088–1096.
- [119] Nilesh Vaidya et al. “Spontaneous network formation among cooperative RNA replicators”. In: *Nature* 491.7422 (2012), p. 72.
- [120] Barry Sinervo and Curt M Lively. “The rock-paper-scissors game and the evolution of alternative male strategies”. In: *Nature* 380.6571 (1996), p. 240.
- [121] Paul E Turner and Lin Chao. “Escape from prisoner’s dilemma in RNA phage $\Phi 6$ ”. In: *The American Naturalist* 161.3 (2003), pp. 497–505.
- [122] Laura Elizabeth Satterwhite, Jessica AM Yeates, and Niles Lehman. “Group I intron internal guide sequence binding strength as a component of ribozyme network formation”. In: *Molecules* 21.10 (2016), p. 1293.
- [123] Barbara Reinhold-Hurek and David A Shub. “Self-splicing introns in tRNA genes of widely divergent bacteria”. In: *Nature* 357.6374 (1992), p. 173.
- [124] Louis Y Kuo, Leslie A Davidson, and Stacy Pico. “Characterization of the Azoarcus ribozyme: tight binding to guanosine and substrate by an unusually small group I ribozyme”. In: *Biochimica et Biophysica Acta (BBA)-Gene Structure and Expression* 1489.2 (1999), pp. 281–292.
- [125] Feng Guo and Thomas R Cech. “In vivo selection of better self-splicing introns in *Escherichia coli*: the role of the P1 extension helix of the *Tetrahymena* intron”. In: *RNA* 8.5 (2002), pp. 647–658.
- [126] Will E Draper, Eric J Hayden, and Niles Lehman. “Mechanisms of covalent self-assembly of the Azoarcus ribozyme from four fragment oligonucleotides”. In: *Nucleic Acids Research* 36.2 (2007), pp. 520–531.
- [127] Alessandro Filisetti et al. “A stochastic model of the emergence of autocatalytic cycles”. In: *Journal of Systems Chemistry* 2.1 (2011), p. 2.

- [128] Wim Hordijk and Mike Steel. “A formal model of autocatalytic sets emerging in an RNA replicator system”. In: *Journal of Systems Chemistry* 4.1 (2013), p. 3.
- [129] Albert Eschenmoser. “Chemical etiology of nucleic acid structure”. In: *Science* 284.5423 (1999), pp. 2118–2124.
- [130] Jessica Flack. “12 Life’s Information Hierarchy”. In: *From Matter to Life: Information and Causality* (2017), p. 283.
- [131] Eörs Szathmáry. “Toward major evolutionary transitions theory 2.0”. In: *Proceedings of the National Academy of Sciences* 112.33 (2015), pp. 10104–10111.
- [132] Kunio Kawamura. “A Hypothesis: Life Initiated from Two Genes, as Deduced from the RNA World Hypothesis and the Characteristics of Life-Like Systems”. In: *Life* 6.3 (2016), p. 29.
- [133] Hyunju Kim, Paul Davies, and Sara Imari Walker. “New scaling relation for information transfer in biological networks”. In: *Journal of The Royal Society Interface* 12.113 (2015), p. 20150944.
- [134] Sara Imari Walker, Hyunju Kim, and Paul CW Davies. “The informational architecture of the cell”. In: *Phil. Trans. R. Soc. A* 374.2063 (2016), p. 20150057.
- [135] Nick Lane and William F Martin. “The origin of membrane bioenergetics”. In: *Cell* 151.7 (2012), pp. 1406–1416.
- [136] Lawrence A Loeb and Thomas A Kunkel. “Fidelity of DNA synthesis”. In: *Annual review of biochemistry* 51.1 (1982), pp. 429–457.
- [137] Jack W Szostak. “The eightfold path to non-enzymatic RNA replication”. In: *Journal of Systems Chemistry* 3.1 (2012), p. 2.
- [138] Leslie E Orgel. “Molecular replication”. In: *Nature* 358.6383 (1992), p. 203.
- [139] Katarzyna Adamala and Jack W Szostak. “Nonenzymatic template-directed RNA synthesis inside model protocells”. In: *Science* 342.6162 (2013), pp. 1098–1100.
- [140] Kevin Leu et al. “The prebiotic evolutionary advantage of transferring genetic information from RNA to DNA”. In: *Nucleic acids research* 39.18 (2011), pp. 8135–8147.

- [141] Sudha Rajamani et al. “Effect of stalling after mismatches on the error catastrophe in nonenzymatic nucleic acid replication”. In: *Journal of the American Chemical Society* 132.16 (2010), pp. 5880–5885.
- [142] Stephen M Testa et al. “Thermodynamics of RNA- RNA duplexes with 2-or 4-thiouridines: implications for antisense design and targeting a group I intron”. In: *Biochemistry* 38.50 (1999), pp. 16655–16662.
- [143] Abdalla EA Hassan et al. “High fidelity of base pairing by 2-selenothymidine in DNA”. In: *Journal of the American Chemical Society* 132.7 (2010), pp. 2120–2121.
- [144] Noam Prywes et al. “Nonenzymatic copying of RNA templates containing all four letters is catalyzed by activated oligonucleotides”. In: *Elife* 5 (2016).
- [145] Richard J Bagley and J Doyne Farmer. *Spontaneous emergence of a metabolism*. Tech. rep. Los Alamos National Lab., NM (USA), 1990.
- [146] F John Odling-Smee et al. *Niche construction: the neglected process in evolution*. 37. Princeton university press, 2003.
- [147] Paulien Hogeweg and Nobuto Takeuchi. “Multilevel selection in models of prebiotic evolution: compartments and spatial self-organization”. In: *Origins of Life and Evolution of the Biosphere* 33.4-5 (2003), pp. 375–403.
- [148] James Woodward. “Causation in biology: stability, specificity, and the choice of levels of explanation”. In: *Biology & Philosophy* 25.3 (2010), pp. 287–318.
- [149] Stephen Wolfram. *A new kind of science*. Vol. 5. Wolfram media Champaign, 2002.

APPENDIX A
STATEMENT OF CO-AUTHOR PERMISSIONS

All co-authors have granted their permissions to use articles Mathis et al. 2017 (Astrobiology) and Mathis et al. 2017 (Life) for Chapters 3 and 4 respectively.