



TIMELINE

The Santa Fe Institute Across the Decades

David C. Krakauer xxvii

1984–1999: **Mavericks**

1: Complex Adaptive Systems: A Primer
John H. Holland 1

2: Bounded Rationality and Other Departures
Kenneth Arrow and George Cowan 9

3: Can Physics Contribute to Economics?
Richard Palmer 21

4: Nature Conformable to Herself
Murray Gell-Mann 31

5: The Simply Complex: Trendy Buzzword
or Emerging New Science?
John Casti 41

6: Learning How to Control Complex Systems
Seth Lloyd 51

7: Beyond Extinction: Rethinking Biodiversity
Simon Levin 63

8: What Can Emergence Tell Us About Today’s Eastern Europe?
Cosma Shalizi 71

2000–2014: **Unifiers**

9: The Evolutionary Dynamics of Social Organization in
Insect Societies: From Behavior to Genes and Back
Joachim Erber and Robert E. Page, Jr. 81

10: Picasso and Perception: Attending to the Higher Order
Tom Kepler 93

11: Four Complications in Understanding the Evolutionary Process
Richard C. Lewontin 97

12:	Searching for the Laws of Life: Separating Chance From Necessity <i>D. Eric Smith and Harold J. Morowitz</i>	115
13:	Metaphors: Ladders of Innovation <i>David Gray and Michele Macready</i>	129
14:	The Numbers of Our Nature: Is There a Math of Style? <i>Dan Rockmore</i>	147
15:	On Time and Risk <i>Ole Peters</i>	161
16:	Transcience: Disciplines & the Advance of Plenary Knowledge <i>David C. Krakauer</i>	169
17:	What Biology Can Teach Us About Banking <i>Lord Robert May</i>	175
18:	Imagining Complex Societies <i>Scott G. Ortman</i>	181
19:	Complexity: A Different Way to Look at the Economy <i>W. Brian Arthur</i>	193
20:	Life's Information Hierarchy <i>Jessica C. Flack</i>	201

2015 AND BEYOND: Terraformers


21:	Complexity: Worlds hidden in plain sight <i>David C. Krakauer</i>	229
22:	A Planet of Cities <i>Luis M.A. Bettencourt and Geoffrey B. West</i>	235
23:	Predicting the Next Recession <i>Rob Axtell and J. Doyne Farmer</i>	243
24:	Are Humans Truly Unique? How Do We Know? <i>Jennifer A. Dunne and Marcus J. Hamilton</i>	249
25:	Engineered Societies <i>Jessica C. Flack and Manfred D. Laubichler</i>	257
26:	Why People Become Terrorists <i>Mirta Galesic</i>	265
27:	Beehives and Voting Booths <i>John H. Miller</i>	273

28:	The Source Code of Political Power <i>Simon DeDeo</i>	281
29:	The Complex Economics of Self-Interest <i>Samuel Bowles</i>	289
30:	Water Management Is a Wicked Problem, But Not an Unsolvable One <i>Christa Brelsford</i>	295
31:	What Can Mother Nature Teach Us About Managing Financial Systems? <i>Simon Levin and Andrew Lo</i>	303
32:	What Happens When the Systems We Rely on Go Haywire? <i>John H. Miller</i>	311
33:	When an Alliance Comes with Strings Attached <i>Paula L.W. Sabloff</i>	317
34:	Thanksgiving 2050: To Feed the World We Have to Stop Destroying Our Soil <i>Molly Jahm</i>	323
35:	How Complexity Science Can Help Keep the Lights On <i>Seth Blumsack</i>	329
36:	Why Predicting the Future Is More Than Just Horseplay <i>Daniel B. Larremore and Aaron Clauset</i>	339
37:	Emergent Engineering: Reframing the Grand Challenge for the 21st Century <i>David C. Krakauer</i>	349

C	Computation
O	Origins & Emergence
M	Machines & Culture
P	Populations & Collectives
L	Life & Evolution
E	Economies & Scaling
X	eXploration & Theory

HOW TO READ THIS VOLUME

The interconnected nature of complexity science enables you, the reader, to choose your own adventure, as it were. At left are the broad themes tying together the eclectic essays in this book. Follow the shaded tabs on the right edge of the page to pursue a topic.



LIFE'S INFORMATION HIERARCHY

Jessica C. Flack, University of Wisconsin–Madison
SFI Bulletin, April 2014

Biological systems—from cells to tissues to individuals to societies—are hierarchically organized (e.g., Feldman and Eschel 1982; Buss 1987; Smith and Szathmáry 1998; Valentine and May 1996; Michod 2000; Frank 2003). To many, *hierarchical organization* suggests the nesting of components or individuals into groups, with these groups aggregating into yet larger groups. But this view—at least superficially—privileges space and matter over time and information. Many types of neural coding, for example, require averaging or summing over neural firing rates. The neurons' spatial location—that they are in proximity—is, of course, important, but at least as important to the encoding is their behavior in time. Likewise, in some monkey societies, as I will discuss in detail later in this review, individuals estimate the future cost of social interaction by encoding the average outcome of past interactions in special signals and then summing over these signals.

In both examples, information from events distributed in time as well as space (fig. 1) is captured with encodings that are used to control some behavioral output. My collaborators and I in the Center for Complexity & Collective Computation (C4) are exploring the idea that hierarchical organization at its core is a nesting of these kinds of functional encodings. As I will explain, we think these functional encodings result from biological systems manipulating space and time (fig. 2) to facilitate information extraction, which in turn facilitates more efficient extraction of energy.

This information hierarchy appears to be a universal property of biological systems and may be the key to one of life's greatest

PHASE II: UNIFIERS

202

mysteries—the origins of biological complexity. In this essay, I review a body of work by David Krakauer, myself, and our research group that has been inspired by many years of work at the Santa Fe Institute (e.g., Crutchfield 1994; Gell-Mann 1996; Gell-Mann and Lloyd 1996; Fontana and Buss 1996; West, Brown, and Enquist 1997; Fontana and Schuster 1998; Ancel and Fontana 2000; Stadler, Stadler, Wagner, and Fontana 2001; Smith 2003; Crutchfield and Görnerup 2006; Smith 2008). Our work suggests that complexity and the multiscale structure of biological systems are the predictable outcome of evolutionary dynamics driven by uncertainty minimization (Krakauer 2011; Flack 2012; Flack, Erwin, Elliot, and Krakauer 2013).

This recasting of the evolutionary process as an inferential one¹ (Bergstrom and Rosvall 2009; Krakauer 2011) is based on the premise that organisms and other biological systems can be viewed as hypotheses about the present and future environments they or their offspring will encounter, induced from the history of past environmental states they or their ancestors have experienced. This premise, of course, only holds if the past is prologue—that is, has regularities, and the regularities can be estimated and even manipulated (as in niche construction) by biological systems or their components to produce adaptive behavior.

If these premises are correct, life at its core is computational, and a central question becomes: How do systems and their components estimate and control the regularity in their environments and use these estimates to tune their strategies? I suggest that the answer to this question, and the explanation for complexity, is that biological systems manipulate spatial and temporal structure to produce order—low variance—at local scales.

¹ This idea is related to work on Maxwell's demon (e.g., Krakauer 2011; Mandal, Quan, and Jarzynski 2013) and the Carnot cycle (e.g., Smith 2003), but we do not yet understand the mapping.

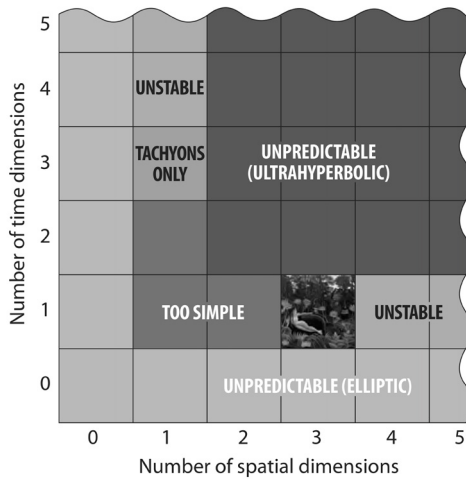


FIGURE 1. The dimensionality of the time–space continuum, with properties postulated when x does not equal 3 and y is larger than 1. Life on earth exists in three spatial dimensions and one temporal dimension. Biological systems effectively “discretize” time and space to reduce environmental uncertainty by coarse-graining and compressing environmental time series to find regularities. Components use the coarse-grained descriptions to predict the future, tuning their behavior to their predictions. [MODIFIED BY JAN UNDERWOOD FOR SFI FROM ORIGINAL BY MAX TEGMARK, WIKIMEDIA COMMONS]




FIGURE 2. Biological systems—from (left to right) Volvox colonies, to slime molds, to animal societies, to large-scale ecosystems such as reefs, to human cities—are hierarchically organized, with multiple functionally important time and space scales. All feature: components with only partially aligned interests exhibiting coherent behavior at the aggregate level; components that turn over and that co-exist in the system at varying stages of development; social structure that persists but component behavior that fluctuates; and macroscopic variation in temporal and spatial structure and coupling with microscopic behavior, which has functional implications when the components can perceive—in evolutionary, developmental, or ecological time—regularities at the macroscopic scale.

Uncertainty Reduction

204

The story I want to tell starts with the observation that with each new level of organization typically comes new functionality—a new feature with positive payoff consequences for the system as a whole, or for its components (Flack, Erwin, Elliot, and Krakauer 2013). Policing in a pigtailed macaque group is an example. Once a heavy tailed distribution of power—defined as the degree of consensus in the group that an individual can win fights (see Flack and Krakauer 2006; Boehm and Flack 2010; Brush, Krakauer, and Flack 2013)—becomes effectively institutionalized (meaning hard to change) policing (an intrinsically costly strategy) becomes affordable, at least to those animals that sit in the tail of the power distribution: those superpowerful monkeys who are rarely or never challenged when they break up fights (Flack, de Waal, and Krakauer 2005; Flack, Girvan, de Waal, and Krakauer 2006).

My collaborators and I propose that a primary driver of the emergence of new functionality such as policing is the reduction of environmental uncertainty through the construction of nested dynamical processes with a range of characteristic time constants (Flack, Erwin, Elliot, and Krakauer 2013). These nested dynamical processes arise as components extract regularities from fast, microscopic behavior by coarse-graining (or compressing) the history of events to which they have been exposed.

Proteins, for example, can have a long half-life relative to RNA transcripts and can be thought of as the summed output of translation. Cells have a long half-life relative to proteins, and are a function of the summed output of arrays of spatially structured proteins. Both proteins and cells represent some average measure of the noisier activity of their constituents. Similarly, a pigtailed macaque's estimate of its power is a kind of average measure of the collective perception in the group that the macaque is capable of winning fights, and this is a better predictor of the cost the macaque will pay during fights than the outcome of any single melee, as

these outcomes can fluctuate for contextual reasons. These coarse-grainings, or averages, are encoded as slow variables (Flack and de Waal 2007; Flack 2012; Flack, Erwin, Elliot, and Krakauer 2013; see also Feret, Danos, Krivine, Harner, and Fontana 2009, for a similar idea). Slow variables may have a spatial component as well as a temporal component, as in the protein and cell examples (fig. 6), or, minimally, only a temporal component, as in the monkey example.

As a consequence of integrating overabundant microscopic processes, slow variables provide better predictors of the local future configuration of a system than the states of the fluctuating microscopic components. In doing so, they promote accelerated rates of microscopic adaptation. Slow variables facilitate adaptation in two ways: they allow components to fine-tune their behavior, and they free components to search, at low cost, a larger space of strategies for extracting resources from the environment (Flack 2012; Flack, Erwin, Elliot, and Krakauer 2013). This phenomenon is illustrated by the power-in-support-of-policing example and also by work on the role of neutral networks in RNA folding. In the RNA case, many different sequences can fold into the same secondary structure. This implies that over evolutionary time, structure changes more slowly than sequence, thereby permitting sequences to explore many configurations under normalizing selection.

New Levels of Organization

As an interaction history builds up at the microscopic level, the coarse-grained representations of the microscopic behavior consolidate, becoming for the components increasingly robust predictors of the system's future state.

We speak of a new organizational level when the system's components rely to a greater extent on these coarse-grained or compressed descriptions of the system's dynamics for adaptive decision-making than on local fluctuations in the microscopic behavior

205

C

P

L

X

ENDOMESODERM SPECIFICATION UP TO 30 HOURS

Maternal Inputs

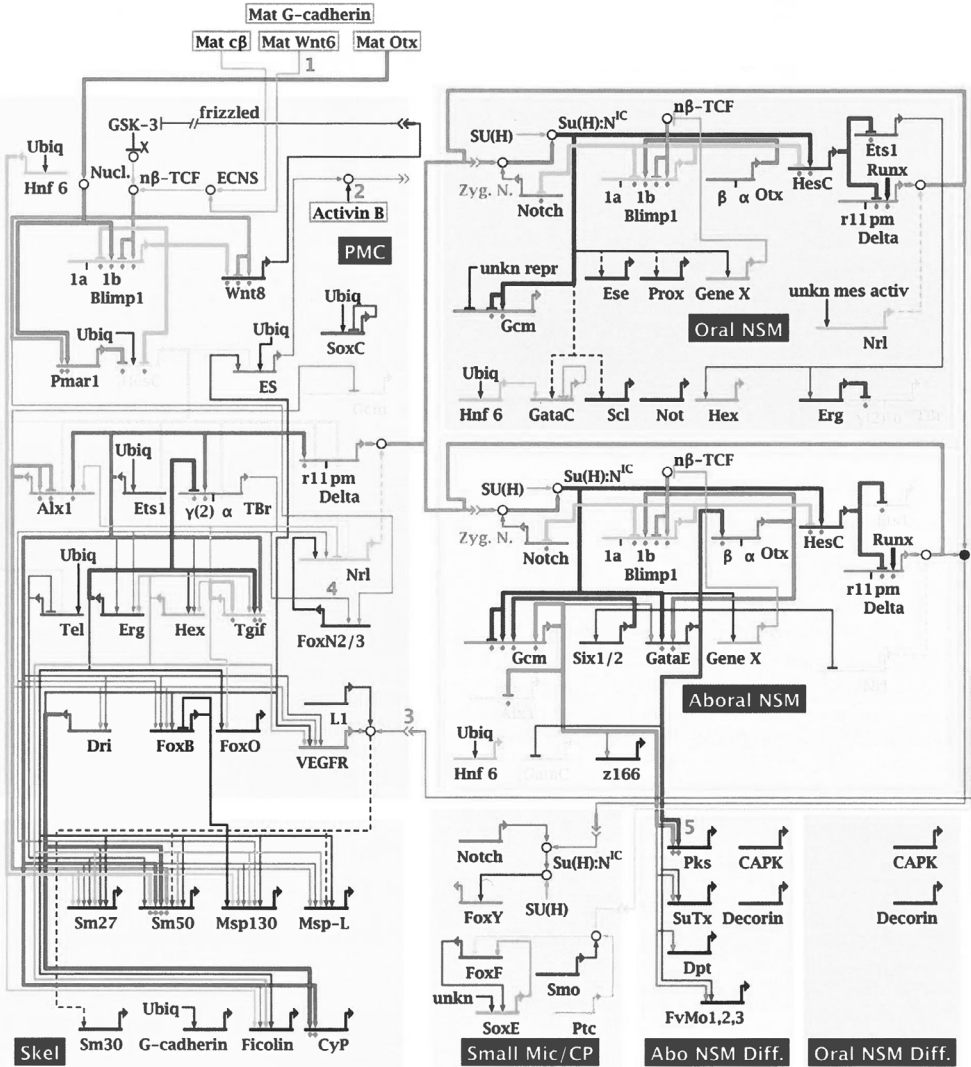
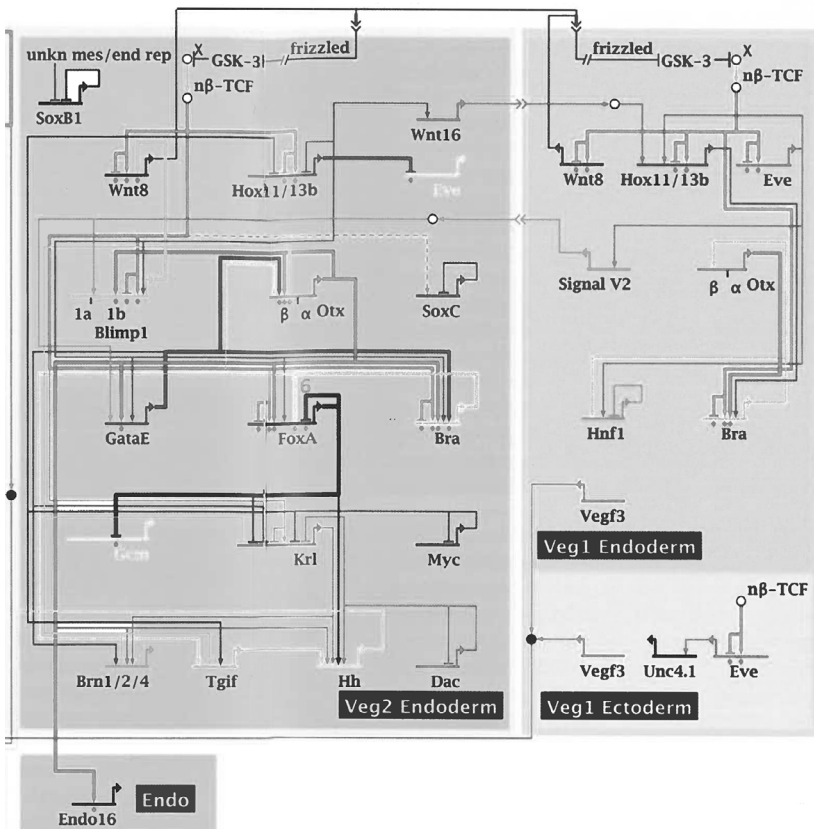


FIGURE 3. A sea urchin gene regulatory circuit. The empirically derived circuit describes the Boolean rules for coordinating genes and proteins to produce aspects of the sea urchin's phenotype—in this case, the position of cells in the endomesoderm at 30 hours since fertilization. Edges indicate whether a node induces a state change in another node, here genes and proteins. The circuit is a rigorous starting point for addressing questions about the logic of development and its evolution. In computational terms, the input is the set of relevant genes and proteins and the output is the target phenotypic feature.



Ubiqu = ubiquitous; Mat = maternal; activ = activator; rep = repressor; unkn = unknown; Nucl. = nuclearization; X = β -catenin source; n β -TCF = nuclearized b- β -catenin-Tcf1; ES = early signal; ECNS = early cytoplasmic nuclearization system; Zyg. N. = zygotic Notch

Additional data sources for selected notes: 1: McClay lab; 2: Angerer lab; 3, 4: McClay lab; 5: Rogers and Calestani, 2010; 6: Croce and McClay

This model is frequently revised. It is based on the latest laboratory data [as of April 2014], some of which was not then published.

PHASE II: UNIFIERS

208

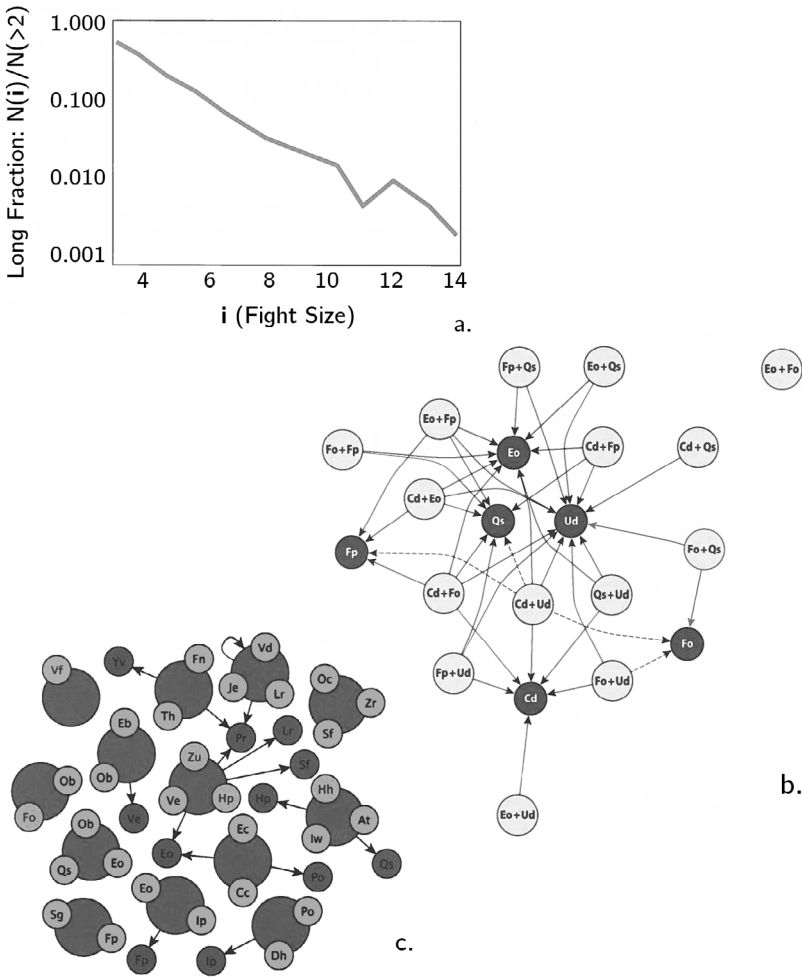


FIGURE 4. Cognitive effective theories for one macroscopic property of a macaque society: the distribution of fight sizes (a). To reduce circuit complexity we return to the raw time series data and remove as much noise as possible by compressing the data. In the case of our macaque dataset, this reveals which individuals and subgroups are regular and predictable conflict participants. We then search for possible strategies in response to these regular and predictable individuals and groups. This approach returns a family of circuits (b is an example), each of which has fewer nodes and edges than the full circuit (c). These circuits are simpler and more cognitively parsimonious. We then test the reduced circuits against each other in simulation to determine how well they recover the target macroscopic properties.

and when the coarse-grained estimates made by components are largely in agreement. The idea is that convergence on these “good enough” estimates underlies nonspurious correlated behavior among the components. This in turn leads to an increase in local predictability (e.g., Flack and de Waal 2007; Brush, Krakauer, and Flack 2013) and drives the construction of the information hierarchy. (Note that increased predictability can seem the product of downward causation in the absence of careful analysis of the bottom-up mechanisms that actually produced it.)

209

The Statistical Mechanics & Thermodynamics of Biology

Another way of thinking about slow variables is as a functionally important subset of the system's potentially many macroscopic properties. An advantage of this recasting is that it builds a bridge to physics, which over the course of its maturation as a field grappled with precisely the challenge now before biology: understanding the relationship between behavior at the individual or component level and behavior at the aggregate level.

C

In physics

As discussed in Krakauer and Flack (2010), the debate in physics began with thermodynamics—an equilibrium theory treating aggregate variables—and came to a close with the maturation of statistical mechanics—a dynamical theory treating microscopic variables.

P

Thermodynamics is the study of the macroscopic behavior of systems exchanging work and heat with connected systems or their environment. The four laws of thermodynamics all operate on average quantities defined at equilibrium—temperature, pressure, entropy, volume, and energy. These macroscopic variables exist in fundamental relationships with each other, as expressed, for example, in the ideal gas law. Thermodynamics is an extremely powerful framework as it provides experimentalists with explicit, principled recommendations

L

X

PHASE II: UNIFIERS

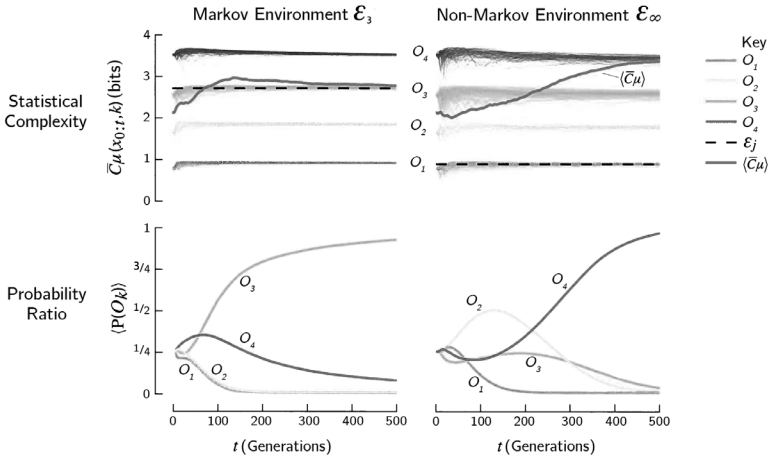


FIGURE 5. A comparison of Markov organisms in two environments: a Markov environment (left) and a non-Markov environment (right). In the top two plots, organismal complexity is plotted against time for each organism (organisms are represented by varying colors) and for many different sequences of 500 environmental observations; the bold red line shows the average organismal complexity, which in the Markov environment tends toward the environmental complexity and in the non-Markov environment exceeds it. In the bottom plots, the probability that a random organism has order k is plotted against time.




FIGURE 6. The cell can be thought of as a slow variable to the extent it is a function of the summed output of arrays of spatially structured proteins and has a long half-life compared to its proteins. Features that serve as slow variables provide better predictors of the local future configuration of a system than the states of the fluctuating microscopic components. We propose that when detectable by the system or its components, slow variables can reduce environmental uncertainty and, by increasing predictability, promote accelerated rates of microscopic adaptation.

about what variables should be measured and how they are expected to change relative to each other, but it is not a dynamical theory and offers no explanation for the mechanistic origins of the macroscopic variables it privileges. This is the job of statistical mechanics. By providing the microscopic basis for the macroscopic variables in thermodynamics, statistical mechanics establishes the conditions under which the equilibrium relations are no longer valid or expected to apply. The essential intellectual technologies behind much of statistical mechanics are powerful tools for counting possible microscopic configurations of a system and connecting these to macroscopic averages.

211

In biology

This brief summary of the relation between thermodynamics and statistical mechanics in physics is illuminating for two reasons. On the one hand it raises the possibility of a potentially deep division between physical and biological systems: so far—and admittedly biology is young—biology has had only limited success in empirically identifying important macroscopic properties and deriving these from first principles rooted in physical laws or deep evolved constraints.² This may be the case because many of the more interesting macroscopic properties are slow variables that result from the collective behavior of adaptive components, and their functional value comes from how components use them, making them fundamentally subjective (see Gell-Mann and Lloyd 1996 for more on subjectivity) and perhaps even nonstationary.³

On the other hand, the role of statistical mechanics in physics

² The work on scaling in biological systems shows a fundamental relationship between mass and metabolic rate, and this relationship can be derived from the biophysics (e.g., West, Brown, and Enquist 1997). Bettencourt and West are now investigating whether similar fundamental relationships can be established for macroscopic properties of human social systems, like cities (e.g., Bettencourt, Lobo, Helbing, Kuhnert, and West 2007; Bettencourt 2013).

³ With the important caveat that in biology the utility of a macroscopic property as a predictor will likely increase as consensus among the components about the estimate increases, effectively reducing the subjectivity and increasing stationarity (see also Gell-Mann and Lloyd 1996).

C

P

L

X

suggests a way forward. If we have intuition about which macroscopic properties are important—that is, which macroscopic properties are slow variables—and we can get good data on the relevant microscopic behavior, we can proceed by working upward from dynamical many-body formalisms to equilibrium descriptions with a few favored macroscopic degrees of freedom (Levin, Grenfell, Hastings, and Perelson 1997; Krakauer and Flack 2010; Krakauer et al. 2011; Gintis, Doebeli, and Flack 2012).

212

A Statistical Mechanics–Computer Science– Information Theoretic Hybrid Approach

The most common approach to studying the relationship between micro and macro in biological systems is perhaps dynamical systems and, more specifically, pattern formation (for examples, see Sumpter 2006; Ball 2009; Couzin 2009; Payne et al. 2013). However, if, as we believe, the information hierarchy results from biological components collectively estimating regularities in their environments by coarse-graining or compressing time series data, a natural (and complementary) approach is to treat the micro and macro mapping explicitly as a computation.

If we have intuition about which
macroscopic properties are important—
that is, which macroscopic properties are
slow variables—and we can get good
data on the relevant microscopic behavior,
we can proceed by working upward
from dynamical many-body formalisms
to equilibrium descriptions with a few
favored macroscopic degrees of freedom.

Elements of computation in biological systems

Describing a biological process as a computation minimally requires that we are able to specify the output, the input, and the algorithm or circuit connecting the input to the output (Flack and Krakauer 2011; see also Mitchell 2010; Valiant 2013). A secondary concern is how to determine when the desired output has been generated. In computer science this is called the termination criterion or halting problem. In biology it potentially can be achieved by constructing nested dynamical processes with a range of timescales, with the slower timescale processes providing the “background” against which a strategy is evaluated (Flack and Krakauer 2011), as discussed later in this paper in the section “Couplings.”

A macroscopic property can be said to be an output of a computation if it can take on values that have functional consequences at the group or component level, if it is the result of a distributed and coordinated sequence of component interactions under the operation of a strategy set, and if it is a stable output of input values that converges (terminates) in biologically relevant time (Flack and Krakauer 2011). Examples studied in biology include aspects of vision such as edge detection (e.g., Olshausen and Field 2004), phenotypic traits such as the average position of cells in the developing endomesoderm of the sea urchin (e.g., Davidson 2010; Peter and Davidson 2011), switching in biomolecular signal-transduction cascades (e.g., Smith, Krishnamurthy, Fontana, and Krakauer 2011), chromatin regulation (e.g., Prohaska, Stadler, and Krakauer 2010), and social structures such as the distribution of fight sizes and the distribution of power in monkey societies (e.g., Flack 2012; Flack, Erwin, Elliot, and Krakauer 2013).

The input to the computation is the set of elements implementing the rules or strategies. As with the output, we do not typically know a priori which of many possible inputs is relevant, and so we must make an informed guess based on the properties of the output. In the case of the sea urchin's endomesoderm, we might start with a list of genes that have been implicated in the regulation

213

C

P

L

X

PHASE II: UNIFIERS

of cell position. In the case of the distribution of fight sizes in a monkey group, we might start with a list of individuals participating in fights.

Reconstructing the microscopic behavior

214

In biological systems the input plus the strategies constitute the system's microscopic behavior. There are many approaches to reconstructing the system's microscopic behavior. The most powerful is an experiment in which upstream inputs to a target component are clamped off and the output of the target component is held constant. This allows the experimentalist to measure the target component's specific contribution to the behavior of a downstream component (Pearl 2010). This type of approach is used to construct gene regulatory circuits mapping gene–gene and gene–protein interactions to phenotypic traits (fig. 3).

When such experiments are not possible, causal relationships can be established using time series analysis in which clamping is approximated statistically (Ay 2009; Pearl 2010). My collaborators and I have developed a novel computational technique, called Inductive Game Theory, that uses a variant of this statistical clamping principle to extract strategic decision-making rules, game structure, and (potentially) strategy cost from correlations observed in the time series data.

Collective computation through stochastic circuits

In all biological systems, of course, there are multiple components interacting and simultaneously coarse-graining to make predictions about the future. Hence the computation is inherently collective. A consequence of this is that it is not sufficient to simply extract from the time series the list of the strategies in play. We must also examine how different configurations of strategies affect the macroscopic output. One way these configurations can be captured is

by constructing Boolean circuits describing activation rules as illustrated by the gene regulatory circuit shown in figure 3, which controls cell position (the output) at thirty hours from fertilization in the sea urchin (Peter and Davidson 2011). In the case of our work on micro to macro mappings in animal societies, we describe the space of microscopic configurations using stochastic “social” circuits (fig. 4).

Nodes in these circuits are the input to the computation. As discussed above, the input can be individuals or subgroups, or they can be defined in terms of component properties like age or neurophysiological state. A directed edge between two nodes indicates that the “receiving node” has a strategy for the “sending node”—and the edge weight can be interpreted as the above-null probability that the sending node plays the strategy in response to some behavior by the receiving node in a previous time step. Hence, an edge in these circuits quantifies the strength of a causal relationship between the behaviors of a sending and receiving node.

Sometimes components have multiple strategies in their repertoires. Which strategy is being played at time t may vary with context. These metastrategies can be captured in the circuit using different types of gates specifying how a component's myriad strategies combine (see also Feret, Davis, Krivine, Harmer, and Fontana 2009). By varying the types of gates and/or the strength of causal relationships, we end up with multiple alternative circuits—a family of circuits—all of which are consistent with the microscopic behavior, albeit with different degrees of precision. Each circuit in the family is essentially a model of the micro–macro relationship and so serves as a hypothesis for how strategies combine over nodes (inputs) to produce the target output. We test the circuits against each other in simulation to determine which can best recover the actual measured macroscopic behavior of our system.

215

C

P

L

X

Cognitive effective theories for collective computation

The circuits describing the microscopic behavior can be complicated, with many “small” causes detailed, as illustrated by the gene regulatory circuit shown in figure 3. The challenge—once we have rigorous circuits—is to figure out the circuit logic (Flack and Krakauer 2011; see also Feret, Davis, Krivine, Harmer, and Fontana 2009).

216 There are many ways to approach this problem. Our approach is to build what’s called in physics an effective theory: a compact description of the causes of a macroscopic property. Effective theories for adaptive systems composed of adaptive components require an additional criterion beyond compactness. As discussed earlier in this essay, components in these systems are tuning their behaviors based on their own effective theories—coarse-grained rules (see also Feret, Davis, Krivine, Harmer, and Fontana 2009)—that capture the regularities (Daniels, Krakauer, and Flack 2012). If we are to build an effective theory that explains the origins of functional space and timescales—new levels of organization—and ultimately the information hierarchy, the effective theory must be consistent with component models of macroscopic behavior, as these models, through their effects on strategy choice, drive that process. In other words, our effective theory should explain how the system itself is computing.

We begin the search for cognitively principled effective theories using what we know about component cognition to inform how we coarse-grain and compress the circuits. This means taking into account, given the available data, the kinds of computations components can perform and the error associated with these computations at the individual and collective levels, given component memory capacity and the quality of the “data sets” components use to estimate regularities (Krakauer, Flack, DeDeo, and Farmer 2010; Flack and Krakauer 2011; Daniels, Krakauer, and Flack 2012; all building on Gell-Mann 1996).

As we refine our understanding of the micro–macro mapping

through construction of cognitive effective theories, we also refine our understanding of what time series data constitute the “right” input—and hence the building blocks of our system. And, by investigating whether our best-performing empirically justified circuits can also account for other potentially important macroscopic properties, we can begin to establish which macroscopic properties might be fundamental and what their relation is to one another—the thermodynamics of biological collectives.

217

The question we must answer is,
What is the optimal coupling between
macroscopic and microscopic change,
and can systems, by manipulating how
components are organized in space and
time, get close to this optimal coupling?

C

Couplings, information flow, and macroscopic tuning

Throughout this essay I have stressed the importance of slowness (effective stationarity) for prediction. Slowness also has costs, however. Consider our power example. The power structure must change slowly if individuals are to make worthwhile investments in strategies that work well given the structure, but it cannot change too slowly or it may cease to reflect the underlying distribution of fighting abilities on which it is based and, hence, cease to be a good predictor of interaction cost (Flack 2012; Flack, Erwin, Elliot, and Krakauer 2013). The question we must answer is, What is the optimal coupling between macroscopic and microscopic change, and can systems, by manipulating how components are organized in space and time, get close to this optimal coupling?

P

L

X

PHASE II: UNIFIERS

One approach to this problem is to quantify the degeneracy of the target macroscopic property and then perturb the circuits by either removing nodes, up- or down-regulating node behavior, or restructuring higher-order relationships (subcircuits) to determine how many changes at the microscopic level need to occur to induce a state change at the macroscopic level.

218 Another approach is to ask how close the system is to a critical point—that is, how sensitive the target macroscopic property is to small changes in parameters describing the microscopic behavior. Many studies suggest that biological systems of all types sit near the critical point (Mora and Bialek 2011). A hypothesis we are exploring is that sitting near the critical point means that important changes at the microscopic scale will be visible at the macroscopic scale. Of course this also has disadvantages as it means small changes can potentially cause big institutional shifts, undermining the utility of coarse-graining and slow variables for prediction (Flack, Erwin, Elliot, and Krakauer 2013).

If balancing trade-offs between robustness and prediction on the one hand, and adaptability to changing environments on the other, can be achieved by modulating the coupling between scales (Flack, Hammerstein, and Krakauer 2012; Flack, Erwin, Elliot, and Krakauer 2013), we should be able to make predictions about whether a system is far from, near, or at the critical point based on whether the data suggest that robustness or adaptability is more important given the environment and its characteristic timescale. This presupposes that the system can optimize where it sits with respect to the critical point, implying active mechanisms for modulating the coupling. We are working to identify plausible mechanisms using a series of toy models to study how the type of feedback from the macroscopic or institutional level to the microscopic behavior influences the possibility of rapid institutional switches.

Complexity

This essay covers a lot of work, so allow me to summarize. I suggested that the origins of the information hierarchy lie in the manipulation of space and time to reduce environmental uncertainty. I further suggested that uncertainty reduction is maximized if the coarse-grained representations of the data the components compute are in agreement (because this increases the probability that everyone is making the same predictions and so tuning the same way). As this happens, the coarse-grained representations consolidate into robust, slow variables at the aggregate level, creating new levels of organization and giving the appearance of downward causation.

I proposed that a central challenge lies in understanding what the mapping is between the microscopic behavior and these new levels of organization. (How exactly do everyone's coarse-grainings converge?) I argued that, in biology, a hybrid statistical mechanics–computer science–information theoretic approach (see also Krakauer et al. 2011) is required to establish such mappings. Once we have cognitively principled effective theories for mappings, we will have an understanding of how biological systems, by discretizing space and time, produce information hierarchies.

Where are we, though, with respect to explaining the origins of biological complexity?

The answer we are moving toward lies at the intersection of the central concepts in this essay. If evolution is an inferential process with complex life being the result of biological systems extracting regularities from their environments to reduce uncertainty, a natural recasting of evolutionary dynamics is in Bayesian terms. Under this view, organism and environment can be interpreted as k -order Markov processes and modeled using finite-state hidden Markov models (fig. 5). Organisms update prior models of the environment with posterior models of observed regularities. We are exploring how the Markov order (a proxy for memory) of organisms changes as organisms evolve to match their environment, quantifying fit

219

C

P

L

X

PHASE II: UNIFIERS

to the environment with model selection. We use information-theoretic measures to quantify structure. Our approach allows us to evaluate the memory requirements of adapting to the environment given its Markov order, quantify the complexity of the models organisms build to represent their environments, and quantitatively compare organismal and environmental complexity as our Markov organisms evolve. We hypothesize that high degrees of complexity result when there is regularity in the environment, but it takes a long history to perceive it and an elaborate model to encode it. 🐙

220

ACKNOWLEDGMENTS

This essay summarizes my view of the past, present, and predicted future of the core research program at the Center for Complexity & Collective Computation (C₄). In addition to our current collaborators—Nihat Ay, Dani Bassett, Karen Page, Chris Boehm, and Mike Gazzaniga—and the super-smart students and postdoctoral fellows Eleanor Brush, Bryan Daniels, Simon DeDeo, Karl Doron, Chris Ellison, the late Tanya Elliot, Evandro Ferrada, Eddie Lee, and Philip Poon, who have carried out much of this work on a daily basis, I am deeply grateful to the Santa Fe Institute for its support over the years and to the Santa Fe folks whose ideas have provided inspiration. First and foremost this includes David Krakauer, my main collaborator. Other significant SFI influences include Jim Crutchfield, Doug Erwin, Walter Fontana, Lauren Ancel Meyers, Geoffrey West, Eric Smith, Murray Gell-Mann, Bill Miller, David Padwa, and Cormac McCarthy. I am indebted to Ellen Goldberg for making possible my first postdoctoral position at SFI. Finally, much of this research would not be possible without the generous financial support provided by the John Templeton Foundation through a grant to SFI to study complexity and a grant to C₄ to study the mind–brain problem, a National Science Foundation grant (0904863), and a grant from the US Army Research Laboratory and the US Army Research Office under contract number W911NF-13-1-0340.

REFERENCES CITED

- Ancel, L. W., and W. Fontana. 2000. "Plasticity, Evolvability, and Modularity in RNA." *Journal of Experimental Zoology. Part B, Molecular and Developmental Evolution* 288: 242–83.
- Ay, N. 2009. "A Refinement of the Common Cause Principle." *Discrete Applied Mathematics* 157: 2439–57.
- Ball, P. 2009. *Nature's Patterns: A Tapestry in Three Parts*. Oxford, UK: Oxford University Press.
- Bergstrom, C. T., and M. Rosvall. 2009. "The Transmission Sense of Information." *Biology and Philosophy* 26: 159–76.
- Bettencourt, L. M. A. 2013. "The Origins of Scaling in Cities." *Science* 340: 1438–41.
- Bettencourt, L. M. A., J. Lobo, D. Helbing, C. Kuhnert, and G. B. West. 2007. "Growth, Innovation, Scaling, and the Pace of Life in Cities." *PNAS* 104: 7301–6.
- Boehm, C., and J. C. Flack. 2010. "The Emergence of Simple and Complex Power Structures through Niche Construction." In *The Social Psychology of Power*, edited by A. Guinote and T. K. Vescio, 46–86. New York, NY: Guilford Press.
- Brush, E. R., D. C. Krakauer, and J. C. Flack. 2013. "A Family of Algorithms for Computing Consensus About Node State from Network Data." *PLOS Computational Biology* 9: e1003109.
- Buss, L. W. 1987. *The Evolution of Individuality*. Princeton, NJ: Princeton University Press.
- Couzin, I. D. 2009. "Collective Cognition in Animal Groups." *Trends in Cognitive Science* 13: 36–43.
- Crutchfield, J. P. 1994. "The Calculi of Emergence: Computation, Dynamics, and Induction." *Physica D* 75: 11–54.
- Crutchfield, J. P., and D. P. Feldman. 2001. "Synchronizing to the Environment: Information-Theoretic Constraints on Agent Learning." *Advances in Complex Systems* 4: 251–64.
- Crutchfield, J. P., and O. Görnerup. 2006. "Objects That Make Objects: The Population Dynamics of Structural Complexity." *Journal of the Royal Society Interface* 22: 345–49.
- Daniels, B., D. C. Krakauer, and J. C. Flack. 2012. "Sparse Code of Conflict in a Primate Society." *PNAS* 109: 14259–64.
- Daniels, B., D. C. Krakauer, and J. C. Flack. N.d. *Conflict Tuned to Maximum Information Flow*. In preparation.

221

C

P

L

X

PHASE II: UNIFIERS

- Davidson, E. H. 2010. "Emerging Properties of Animal Gene Regulatory Networks." *Nature* 468: 911–20.
- Dedeo, S., D. C. Krakauer, and J. C. Flack. 2010. "Inductive Game Theory and the Dynamics of Animal Conflict." *PLOS Computational Biology* 6: E1000782.
- Ellison, C., J. C. Flack, and D. C. Krakauer. n.d. *On Inferential Evolution and the Complexity of Life*. In preparation.
- Feldman, M., and I. Eschel. 1982. "On the Theory of Parent–Offspring Conflict: A Two-Locus Genetic Model." *American Naturalist* 119: 285–92.
- 222
- Feret, J., V. Danos, J. Krivine, R. Harmer, and W. Fontana. 2009. "Internal Coarse-Graining of Molecular Systems." *PNAS* 106: 6453–58.
- Ferrada, E., and D. C. Krakauer. n.d. *The Simon Modularity Principle*. In preparation.
- Flack, J. C. 2012. "Multiple Time-Scales and the Developmental Dynamics of Social Systems." *Philosophical Transactions of the Royal Society B: Biological Sciences* 367: 1802–10.
- Flack, J. C., and D. C. Krakauer. 2006. "Encoding Power in Communication Networks." *American Naturalist* 168: E87–102.
- Flack, J. C., and F. B. M. De Waal. 2007. "Context Modulates Signal Meaning in Primate Communication." *PNAS* 104: 1581–86.
- Flack, J. C., and D. C. Krakauer. 2011. "Challenges for Complexity Measures: A Perspective from Social Dynamics and Collective Social Computation." *Chaos* 21: 037108.
- Flack, J. C., F. B. M. De Waal, and D. C. Krakauer. 2005. "Social Structure, Robustness, and Policing Cost in a Cognitively Sophisticated Species." *American Naturalist* 165: E126–39.
- Flack, J. C., P. Hammerstein, and D. C. Krakauer. 2012. "Robustness in Biological and Social Systems." In *Evolution and the Mechanisms of Decision-Making*, edited by P. Hammerstein, and J. Stevens, 129–51. Cambridge, MA: MIT Press.
- Flack, J. C., D. Erwin, T. Elliot, and D. C. Krakauer. 2013. "Timescales, Symmetry, and Uncertainty Reduction in the Origins of Hierarchy in Biological Systems." In *Cooperation and Its Evolution*, edited by K. Sterelny, R. Joyce, B. Calcott, and B. Fraser, 45–74. Cambridge, MA: MIT Press.
- Flack, J. C., M. Girvan, F. B. M. De Waal, and D. C. Krakauer. 2006. "Policing Stabilizes Construction of Social Niches in Primates." *Nature* 439: 426–29.

- Fontana, W., and L. W. Buss. 1996. "The Barrier of Objects: From Dynamical Systems to Bounded Organizations." In *Boundaries and Barriers*, edited by J. Casti and A. Karlqvist, 56–116. Reading, MA: Addison–Wesley.
- Fontana, W., and P. Schuster. 1998. "Continuity in Evolution: On the Nature of Transitions." *Science* 280: 1451–55.
- Frank, S. A. 2003. "Repression of Competition and the Evolution of Cooperation." *Evolution* 57: 693–705.
- Gell-Mann, M. 1996. "Fundamental Sources of Unpredictability." Talk presented at conference of the same name. http://www-physics.mps.ohio-state.edu/~perry/p633_spo7/articles/fundamental-sources-of-unpredictability.pdf. Accessed 01/10/2014.
- Gell-Mann, M., and S. Lloyd. 1996. "Information Measures, Effective Complexity, and Total Information." *Complexity* 2: 44–53.
- Gintis, H., M. Doebeli, and J. C. Flack. 2012. "The Evolution of Human Cooperation." *Cliodynamics: Journal of Theoretical and Mathematical History* 3.
- Krakauer, D. C. 2011. "Darwinian Demons, Evolutionary Complexity, and Information Maximization." *Chaos* 21: 037110.
- Krakauer, D. C., N. Bertschinger, N. Ay, E. Olbrich, and J. C. Flack. n.d. "The Information Theory of Individuality." In *What Is an Individual?*, edited by L. Nyhart and S. Lidgard. Chicago, IL: University of Chicago Press. In review.
- Krakauer, D. C., and J. C. Flack. 2010. "Better Living Through Physics." *Nature* 467: 661.
- Krakauer, D. C., and P. Zanotto. 2009. "Viral Individuality and Limitations of the Life Concept." In *Protocells: Bridging Non-Living and Living Matter*, edited by M. A. Rasmussen et al., 313–36. Cambridge, MA: MIT Press.
- Krakauer, D. C., J. C. Flack, S. DeDeo, and D. Farmer. 2010. "Intelligent Data Analysis of Intelligent Systems." *IDA 2010 LNCS* 6065: 8–17.
- Krakauer, D. C., J. P. Collins, D. Erwin, J. C. Flack, W. Fontana, M. D. Laubichler, S. Prohaska, G. B. West, and P. Stadler. 2011. "The Challenges and Scope of Theoretical Biology." *Journal of Theoretical Biology* 276: 269–76.
- Lee, E., B. Daniels, D. C. Krakauer, and J. C. Flack. n.d. *Cognitive Effective Theories for Probabilistic Social Circuits Mapping Strategy to Social Structure*. In preparation.

223

C

P

L

X

PHASE II: UNIFIERS

- Levin, S. A., B. Grenfell, A. Hastings, and A. S. Perelson. 1997. "Mathematical and Computational Challenges in Population Biology and Ecosystems Science." *Science* 275: 334–43.
- Mandal, D., H. T. Quan, and C. Jarzynski. 2013. "Maxwell's Refrigerator: An Exactly Solvable Model." *Physical Review Letters* 111: 030602.
- Michod, R. E. 2000. *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*. Princeton, NJ: Princeton University Press.
- Mitchell, M. 2010. "Biological Computation." Working Paper 2010-09-021, Santa Fe Institute, Santa Fe, NM.
- 224 Mora, T., and W. Bialek. 2011. "Are Biological Systems Poised at Criticality?" *Journal of Statistical Physics* 144: 268–302.
- Olshausen, B., and D. Field. 2004. "Sparse Coding of Sensory Inputs." *Current Opinion in Neurobiology* 14: 481–87.
- Payne, S., L. Bochong, Y. Cao, D. Schaeffer, M. D. Ryser, and L. You. 2013. "Temporal Control of Self-Organized Pattern Formation Without Morphogen Gradients in Bacteria." *Molecular Systems Biology* 9: 697.
- Pearl, J. 2010. *Causality*, 2nd ed. Cambridge, MA: Cambridge University Press.
- Peter, I. S., and E. H. Davidson. 2011. "A Gene Regulatory Network Controlling the Embryonic Specification of Endoderm." *Nature* 474: 635–39.
- Poon, P., J. C. Flack, and D. C. Krakauer. n.d. *Niche Construction and Institutional Switching Via Adaptive Learning Rules*. In preparation.
- Prohaska, S. J., P. F. Stadler, and D. C. Krakauer. 2010. "Innovation in Gene Regulation: The Case of Chromatin Computation." *Journal of Theoretical Biology* 265: 27–44.
- Sabloff, J. A., ed. n.d. *The Rise of Archaic States: New Perspectives on the Development of Complex Societies*. Santa Fe Institute. In preparation.
- Schuster, P., and W. Fontana. 1999. "Chance and Necessity in Evolution: Lessons from RNA." *Physica D: Nonlinear Phenomena* 133: 427–52.
- Smith, E. 2003. "Self-Organization from Structural Refrigeration." *Physical Review E* 68: 046114.
- . 2008. "Thermodynamics of Natural Selection I: Energy Flow and the Limits on Organization." *Journal of Theoretical Biology*, <http://dx.doi.org/10.1016/j.jtbi.2008.02.010>, PDF.
- Smith, E., S. Krishnamurthy, W. Fontana, and D. C. Krakauer. 2011. "Nonequilibrium Phase Transitions in Biomolecular Signal Transduction." *Physical Review E* 84: 051917.

Chapter 20: Life's Information Hierarchy

- Smith, J. M., and E. Szathmáry. 1998. *The Major Transitions in Evolution*. Oxford: Oxford University Press.
- Stadler, B. M. R., P. F. Stadler, G. Wagner, and W. Fontana. 2001. "The Topology of the Possible: Formal Spaces Underlying Patterns of Evolutionary Change." *Journal of Theoretical Biology* 213: 241–74.
- Sumpter, D. J. T. 2006. "The Principles of Collective Animal Behaviour." *Philosophical Transactions of the Royal Society B* 361: 5–22.
- Valentine, J., and R. May. 1996. "Hierarchies in Biology and Paleontology." *Paleobiology* 22: 23–33.
- Valiant, L. 2013. *Probably Approximately Correct*. New York, NY: Basic Books. 225
- West, G. B., J. H. Brown, and B. J. Enquist. 1997. "A General Model for the Origin of Allometric Scaling Laws in Biology." *Science* 276: 122–26.

C

P

L

X