

A prolegomenon to nonlinear empiricism in the human behavioral sciences

CHARLES EFFERSON^{1,3,*} and PETER J. RICHERSON^{1,2,3}

¹Graduate Group in Ecology, University of California, Davis, USA; ²Graduate Group in Animal Behavior, University of California, Davis, USA; ³Department of Environmental Science and Policy, University of California, Davis, USA; *Author for correspondence (e-mail: cmefferson@ucdavis.edu)

Received 14 December 2005; accepted in revised form 15 December 2005

Key words: Cultural evolution, Human evolutionary ecology, Nonlinear dynamics

Abstract. We propose a general framework for integrating theory and empiricism in human evolutionary ecology. We specifically emphasize the joint use of stochastic nonlinear dynamics and information theory. To illustrate critical ideas associated with historical contingency and complex dynamics, we review recent research on social preferences and social learning from behavioral economics. We additionally examine recent work on ecological approaches in history, the modeling of chaotic populations, and statistical application of information theory.

Introduction

“Models do not explain history, but they may tell us where to look.”

Bowles (2004)

“All reasonable ecological models of closed systems are non-linear.”

Murdoch et al. (2003)

“Think Non-Linearity”

A theoretical ecologist’s t-shirt¹

Complexity marks human behavior. We live in huge societies with an intricate division of labor. We employ a vast array of technologies in virtually every ecological system in practically every corner of the planet. The biological foundation of our behavior is a metabolically expensive brain whose full complexity remains largely beyond the ken of science. Moreover, human culture stands as the world’s most intricate system for transmitting behavior extra-genetically. It leads to rapid evolution and prodigious behavioral diversity across societies.

The social sciences all take complex human behavior as their object of study. In spite of this commonality, the various disciplines go about their respective

¹ Our experience with this t-shirt is through Alan Hastings, a theoretical ecologist at UC Davis who from time to time proceeds through the day with the above admonition emblazoned across his torso.

studies with little communication across disciplinary boundaries. The chasm between the social sciences and biology has been even harder to bridge than the chasms among the social sciences. In addition, in spite of the fact that mathematics has proven its value countless times as a language for confronting complexity, the social sciences, with the notable exception of economics, have developed as the least mathematical of the sciences. The quantitative skills of many social scientists do not extend beyond static linear statistical models, and the proliferation of postmodernism has further eroded the role of mathematics in the social sciences.

Happily, we do not have to choose, nor could we, between words and quantities. We can have both poetry and equations. We can have both rhetoric and deduction. In an applied setting both typically have their uses. The most mathematical treatment of a problem requires some discussion, even persuasion, and a little math in turn can often provide a degree of precision simply beyond the reach of words. Together, words and mathematics allow us to travel freely from the intuitive and the compelling to the precise and the counterintuitive, and then back again to boot. More practically, decades of experience in other sciences suggest that math is just what we need as scientists to discipline our thinking when phenomena get complex. Nothing about human behavior recommends a crucial difference or even the legitimacy of a wholly non-mathematical approach for the social sciences.

Here we summarize several of our views on the integration of theory and empiricism in the study of human behavioral complexity. We review recent research that both illustrates and played a role in the development of our thinking. Several themes figure prominently. First, human behavior is complex, and mathematically this complexity translates into the need for an increased emphasis on nonlinear dynamical models. We can derive a simple taxonomy of mathematical models by focusing on two orthogonal dimensions. Models can be linear or nonlinear, and they can be static or dynamical. Static linear models are the simplest models to understand, but also have the most limited potential to summarize complex phenomena. This limitation, in conjunction with the widely recognized complexity of human behavior, further emphasizes the likely problems associated with the fact that many social scientists do not have technical skills extending beyond linear statistical models. At the other end of the spectrum, nonlinear dynamical models have the greatest potential to summarize complex phenomena. The combination of nonlinearity and dynamical feedbacks has an extraordinary potential to yield complex behavior with few variables. Indeed, as May (1974, 1976) first showed, one-dimensional nonlinear difference equations with a single parameter can exhibit a staggering array of behaviors from steady states to chaos. If we have an interest in complexity, the mathematical reality is that nonlinear dynamics can take us far beyond the offerings of static linearity.

Second, empiricism is essential. Mathematics invites a deductive approach, but math without empiricism, especially in a high-dimensional nonlinear environment, represents simply an abundance of logical possibilities. As ecologists

Burnham and Anderson (2002, p. 20) put it, “we believe that ‘truth’ (full reality) in the biological sciences has essentially infinite dimension, and hence full reality cannot be revealed with only finite samples of data and a ‘model’ of those data.... We can only hope to identify a model that gives a good approximation of the data available.” Although this humility in the face of complexity is a lesson that would also serve social scientists well, we will draw additional lessons from Burnham and Anderson. Our primary lessons pertain to their methods for fitting models to data and choosing the best model among a set of candidate models. Model fitting and model selection are essential because, even though behavioral phenomena may be high-dimensional, we cannot necessarily justify thinking about them this way. The best model, given finite data, will generally be much simpler than the exceedingly complicated reality we know is “out there.” The phenomena in question, the data at hand, and the models under consideration all matter in our selection of the best model.

The cognitive-module hypothesis of Tooby and Cosmides (1992) illustrates the problem. No doubt the human brain and the problems it has to solve have “essentially infinite dimension.” Tooby and Cosmides, plausibly enough, posit hundreds or thousands of cognitive modules that, via their interaction with the environment, govern human behavior and its variation within and across societies. Unfortunately, we currently have no way to assess the validity of this idea. Although the hypothesis that we have many cognitive modules rests on a typically verbal theory, thinking about it as a possible basis for integrating theory and empiricism suggests some of the problems likely to arise. Consider a single behavioral response variable measured in many different societies. Initially we have no basis for excluding some modules and including others with respect to an effect on the response. Consequently, an exceedingly simple version of the cognitive-module model would be a linear regression with a dummy variable for each module. Such a regression is no doubt only one of many possible approaches left unspecified by a verbal theory. Nonetheless, we take it as our benchmark. Putting aside the significant problem of how to identify the modules and code which ones are active versus inactive, to say nothing of variable levels of activation, the resulting model would involve estimating a parameter for each module a coefficient. Under thousands of modules, the dimensionality of the system would be unruly, to say the least. Any such model would severely *overfit* the data relative to other candidate models. Indeed, leaving aside the possibility of a model with hundreds or thousands of variables, a model with only tens of terms representing only a relative handful of modules may fit the data on hand extremely well, yet utterly fail to predict any new data we collect. Specifically, assuming we even had enough data to estimate the parameters, the estimates would probably be so imprecise as to lack all meaning. Alternatively, and as a different kind of problem, a model with too many terms can spuriously fit the error in a necessarily finite data set and potentially give apparent statistical support to a highly misleading model. Anyone can overfit data using linear regression and similar techniques. Many biologists and social scientists are in the habit of

collecting large datasets in the hopes of finding statistically significant relationships between independent and dependent variables. This is called *data dredging*. If you collect a large number of independent variables, you will not be disappointed by the fits you get. Even so, you are likely to discover that a colleague who collects a parallel set of data will not replicate your fit, but will report a quite different structure. First chocolate causes cancer, then it protects.

The great irony is that in practice many social scientists, especially outside of economics, rely on static, linear, and relatively low-dimensional models to analyze data. Often they merely set up a test that potentially rejects a null model in favor of a proposed alternative. If the dimensionality of the alternative happens to be too low, but the alternate still fits better than the null, such an approach will lead to *underfitting*. Any number of additional models might fit the data as well or better. Tooby and Cosmides (1992), for example, have studied a well known logical task in which subject performance improves, relative to an abstract version of the problem, when the task is framed in terms of identifying the violation of a social rule. The rejected null in this case is that subjects perform equally well on any task with the same logical structure regardless of how the problem is framed. Tooby and Cosmides take rejection of the null as support for a cognitive module that evolved among ancient hunter-gatherers and is dedicated to detecting the violation of social norms. Alternative models might be that social life creates the module by reinforcement or that culturally transmitted strategies play a role in superior performance. We certainly know in general that innate characteristics, individual learning, and cultural transmission interact to improve performance in various types of decision making. We expect that the best-fitting, low-dimensional model for a given dataset will typically include terms that represent more than one of these three causal elements. Thus the verbal theories Tooby and Cosmides propose and the mathematical methods they employ in data analysis are fundamentally at odds. The grand strategy sounds like overfitting; the actual practice sounds like underfitting.

Evolutionary psychologists no doubt realize that one cannot fit a model, for example, that has 100 variables when one only has 99 data points. A model with 15 variables can also overfit, however, or even a model with 1 or 2 variables. As Forster and Sober (1994) argue, philosophers of science long contended that a pure empiricist would derive a model that fit every data point perfectly, but practicing scientists do not do so because they value simplicity, an extra-empirical concern. Modern statisticians have aimed for a *principled* form of parsimony that steers between the intimidating complexity of the phenomena we study and the distressing limitations of the data we can bring to bear on them. The key insight is that data include both signal and noise. Thus simplicity is not an extra-empirical desideratum; we definitely do not want to fit every data point exactly for, if we do, we shall fit the noise as well as the signal. We need to discriminate between the signal and the noise as we walk a line between underfitting by ignoring important components of the signal

inherent in the data and overfitting by paying attention to the stochastic features unique to the particular sample of data available (Forster and Sober 1994).

The information theoretic criterion originally proposed by Akaike (1973) and summarized in Burnham and Anderson provides a formal approach to this problem. The central concept is information loss. Models lose information intrinsic to data because they are not true; they are models. Given a set of candidate models, however, we would like to identify the model estimated to lose the least amount of information about the processes that generated the data. We cannot know *absolutely* how much information any given model loses because we do not know the truth, but estimating the *relative* information loss turns out to be straightforward. In addition, a deep relationship exists between minimizing information loss and optimizing the trade-off between underfitting and overfitting. This relationship is captured beautifully by Akaike's (1973) Information Criterion (AIC), the centerpiece of Burnham and Anderson's (2002) thoughts on optimal model fitting. We discuss the details below.

Importantly, we justifiably fuss a great deal about the benefits of both Akaike and nonlinear dynamics, perhaps even to the extent that they may seem related. AIC and nonlinear dynamics, however, are conceptually distinct. One can easily calculate AIC values for linear models or nonlinear static models from the standard output of common statistical software. Burnham and Anderson give the recipes. A shift toward AIC would be an improvement over today's standard practice in most of the social sciences. Nonetheless we suspect that AIC and nonlinear dynamics should often go hand and hand in the social sciences. A move toward nonlinear dynamics would improve the set of models under consideration, while a move toward Akaike and information theory would improve the tendency for practicing scientists to select useful models.

As our third theme, we promote disrespecting disciplinary boundaries. The study of human behavior was compromised throughout the 20th Century by the divorce of the social sciences from biology and from each other. Each of the social sciences has its strengths (economics, deductive theory; psychology, experimental methods; anthropology, a grip on human diversity; sociology and political science, an appreciation of social institutions), but each would benefit from borrowing methods, theories, and data from the others. We are after all talking about the behavior of only a single species.

As unfruitful as the divorces among the social sciences have been, the divorce from biology has been at least as disastrous. In principle all the social sciences should be sub-disciplines of biology with porous boundaries both vertically and horizontally. Although we are happy to hail the complexity of human behavior from the front row, nothing about complexity implies that ideas and explanations drawing on biology are inappropriate. Indeed, ecology and evolution have long been at the vanguard with respect to the study of complexity. In this vein we provide a taste below of how evolutionary ecologists are doing some of the most exciting current research integrating theory and empiricism in

complex environments. A growing number of social science researchers are crossing the various disciplinary divides and using ecology, evolutionary biology, game theory, and economic concepts, broadly conceived, with no shame and to good effect. Their numbers ought to increase.

We organize the remainder of the paper in the following way. We begin by reviewing three examples of how nonlinear dynamics figure in the study of human behavioral diversity. The first example is an innovative study using experimental economics to chart cultural variation across societies (Henrich et al. 2004a). An ostinato throughout this paper centers on the tendency of nonlinear dynamics to yield different systems that evolve in different, historically contingent ways under the *same set of causal mechanisms*. We find this a compelling line of inquiry with regard to formalizing the study of history and the macroevolutionary consequences of cultural evolution. To pursue such a line of inquiry, however, we need to document cultural diversity. Anthropologists, of course, have made considerable strides toward this objective. Most of this work has been qualitative. The experimental economic study we review represents a significant methodological development in what we refer to as quantitative ethnography, a useful and even necessary supplement to more conventional ethnographic genres. The second example pertains to the rise and fall of agrarian societies as a problem in nonlinear dynamics (Turchin 2003). This study takes us directly to the heart of the matter; it shows how the mathematics of nonlinear dynamical systems can produce precise insights simply unavailable via strictly rhetorical methods. The final example initiates a discussion of some of the distinctive challenges likely to arise in experiments designed to study learning. We focus in particular on nonlinear social learning and an associated incipient research program that merges the methods of experimental economics with gene-culture coevolutionary theory (Baum et al. 2004; McElreath et al. 2005; Efferson et al. submitted). We continue with two sections on the methodological tools we find especially promising with regard to integrating theory, nonlinear dynamics, and empiricism in the social sciences. We focus on the work of Cushing et al. (2003), a group of ecologists whose research on the population dynamics of flour beetles represents the state of the art in experimental nonlinear dynamics. We round out our survey by discussing the use of information theory (Burnham and Anderson 2002) in data analysis.

Experimental comparative ethnography

“The obvious question at this point was: were the Machiguenga results anomalous in some fashion? Or, were they an indication of some substantial amount of cultural variation lurking outside the student-subject pools of the world’s universities?”

Henrich et al. (2004a)

“After these games it was clearly evident from the comments of participants that many made the association between this game and the institution of *harambee*.”

Ensminger (2004)

The use of experimental methods in economics has exploded in the last three decades. In this short period of time, experiments have demonstrated that many of the key assumptions typical of economic theory in the Walrasian tradition are wrong. Loosely put, we can characterize this tradition as one of decision-making based on constrained optimization in a world with stable, exogenous, self-regarding preferences and non-strategic interactions. Moreover, although resources are scarce, ideal social institutions and complete, fully enforceable contracts are somehow freely available under Walrasian assumptions (Bowles and Gintis 2000; Bowles 2004). Among other findings at odds with the Walrasian tradition, experiments have shown that people exhibit contingent, endogenous, and other-regarding preferences. Kagel and Roth (1995) and Camerer (2003) provide extensive reviews of this research. One particularly consistent finding is that people exhibit preferences over the social characteristics of different outcomes. Under these social preferences, people do not simply care about their own payoffs. They also care about the distribution of payoffs among individuals in some relevant reference group. Correspondingly, players consistently sacrifice absolute payoffs to influence the distribution of payoffs in the group.

In addition to many experiments designed strictly to measure social preferences, public goods games have also received substantial attention. Public goods are goods that are rival in the sense that more for one necessitates less for another. But unlike private goods, public goods have some characteristic that makes excluding those who did not contribute to the production of the good impractical or even impossible. Thus an incentive to free-ride is present: enjoy the public good provisioned by others. In the limit, however, no one contributes in any way, and the public good is not produced. The orthodox prediction from economic theory is that in all periods all players contribute nothing to the provision of the public good. The basic reasoning is straightforward. One would rather contribute nothing and possibly free-ride on others in lieu of contributing only to be exploited by free-riders. In public goods experiments, however, players typically contribute about half of their endowment to the public good in the first period, but contributions decline steadily in subsequent periods. The experimenter can often prevent this tendency for contributions to unravel by introducing an institution like the punishment of non-contributors. Such an institution tends to stabilize average contributions at nearly all of a player's endowment (Fehr and Gächter 2002; Fehr and Fischbacher 2003, 2004), a fruitful situation with respect to both individual payoffs and social efficiency.

As is the case in psychology, the pool of subjects in experimental economics exhibits a heavy bias toward university undergraduates in Europe, North

America, and Japan. Thus no matter how many times researchers confirm a particular result, one is left wondering to what extent experiments have tapped human universals as opposed to the comparatively homogeneous culture of university undergraduates in the world's OECD countries. To address this question, a group of anthropologists, psychologists, and economists recently undertook a remarkable study (Henrich et al. 2004a) to conduct some classic economic experiments in small-scale societies around the world. The most complete data in the study come from the ultimatum game, and these data serve to illustrate some central points.

The ultimatum game is a simple game that taps social preferences. It involves two players, a proposer and a respondent. The experimenter provides a quantity of money, and the proposer has to propose a division of the money between herself and the respondent. The experimenter then presents the offer to the respondent, who can either accept or reject. If the respondent accepts, both players receive the specified amounts. If the respondent rejects, neither player gets anything. To control for interpersonal effects, the experimenter typically mediates the interaction so that neither the proposer nor the respondent ever sees, talks to, or knows the person with whom she plays. The Walrasian prediction is that, if offered any money at all, the respondent's choice is between some positive amount of money and nothing. Because some money is better than no money, the respondent will always accept a positive offer. The proposer, anticipating this behavior, will offer the smallest positive amount² with no fear of rejection. Experimentalists have conducted variations of the ultimatum game thousands of times. The basic results among university undergraduates are highly regular. The most common offer is generally 40–50% of the total. Respondents rarely reject such offers, but they frequently reject offers less than approximately 20% (Camerer 2003). Researchers have carefully varied experimental protocol to unpack the motivations behind these behaviors. The general consensus seems to be that respondents, seeing a low offer as unfair, respond by rejecting because rejection is the only available means of punishing proposers. Proposers, in turn, make offers based on some mix of a concern for fairness and a strategic desire to avoid a rejected offer. To develop an economic theory consistent with such findings, a handful of researchers have developed novel utility functions for use in economic analysis (Fehr And Schmidt 1999; Bolton and Ockenfels 2000; Charness and Rabin 2002; Bowles 2004).

The researchers in the Henrich et al. study all had considerable field experience working among at least one small-scale society somewhere in the world. In all the study encompassed 15 societies, 12 countries, and four continents plus New Guinea. Their key result, in stark contrast to the highly regular play

² The conventional theory does not make a strong prediction at the boundary when the proposer offers zero to the respondent. Technically, the respondent is indifferent to accepting or rejecting. One can assume that respondents always accept when indifferent with any probability in the interval [0,1]. Any specific assumption is arbitrary.

of university undergraduates, documented a previously uncharted degree of variation in participant behavior. The extremes from the ultimatum game data illustrate. The Machiguenga of Peru are at one end of the spectrum. Social interactions among the Machiguenga are primarily at the family level. Until recently, the Machiguenga lived as highly mobile familial bands of subsistence foragers and horticulturalists. Although their lives are increasingly settled, social interactions continue to be primarily with close kin, and cooperation beyond the kin group is extremely rare. They also exhibit little concern for social disapproval, and until recently did not have personal names (Henrich et al. 2004b, Henrich and Smith 2004). Congruent with this ethnographic sketch, Machiguenga proposers offered little in the ultimatum game. The most common offer was 15% of the total. Moreover, even though offers were low, respondents rarely rejected low offers. Only 1 out of 10 offers equal to or less than 20% of the total was rejected. At the other end of the spectrum are the Lamalera whalers of Indonesia. Whaling among the Lamalera is by necessity a highly cooperative activity. A clear division of labor and a highly elaborate system of social norms characterize the capture of whales and the careful distribution of whale meat (Alvard 2004). Among the Lamalera, the modal offer was 50%, the mean offer was 57%, and virtually the entire distribution of offers was at 50% and above. Alvard (2004) had to fabricate low offers, an atypical procedure in experimental economics, to get any data on rejection rates. The Lamalera rejected his sham low offers.

Two other groups that produced some fascinating results are the Au and Gnau of Papua New Guinea. In both of these societies, people value generosity. To refuse a feasible request brings social sanctions, especially if one repeatedly does so. Unsolicited generosity is also possible, but acceptance brings a tacit and nebulous obligation to reciprocate in the future (Tracer 2004). In keeping with these ethnographic details, Au and Gnau proposers, like the Lamalera, exhibited a notable tendency to offer more than 50% of the total. Unlike Lamalera respondents, however, Au and Gnau respondents often rejected such offers. Conventional Walrasian assumptions simply provide no basis for interpreting such results. If Ultimatum play among university undergraduates was initially surprising, these results from the Au and Gnau are absolutely stunning.

A net result of the Henrich et al. study is that culture matters. Although most people might agree with this idea in the abstract, showing that culture matters and demonstrating how it matters are quite a different undertakings. Henrich et al. (2004b) provide a careful statistical analysis that quantifies the explanatory importance of cultural variation in the full cross-cultural data set, but an anecdote from the study illustrates the same conclusion. The Orma, a group of Kenyan pastoralists, have an institution called *harambee*. It is a village-level institution for collecting funds to provide public goods like a new school building. When Ensminger (2004) played the public goods game among the Orma, they immediately dubbed it the “*harambee*” game. Moreover, the institution of *harambee* is progressive in the sense that wealthier

members of the community are expected to contribute more than poorer members. In the public goods game, Ensminger found the same pattern. Wealthier participants contributed more to the public good in the experiment than poor participants.

In essence, the experimental environment is not neutral with respect to the norms, interpretations, and cultures of participants. Although some experimentalists will no doubt see this finding as problematic because it implies a lack of experimental control, a more optimistic interpretation would see culture as, in effect, a treatment. Comparing a cultural treatment to an acultural control group might be impossible, but researchers can of course compare play in one culture to play in another culture. To view the matter another way, if researchers can show that play in experimental games consistently corresponds to salient cultural variables, then experimental economics has the potential to become invaluable as a methodological tool in comparative ethnography, a technique for measuring the macro-evolutionary consequences of cultural transmission. Experimental economic games in essence offer the possibility of a theory-driven, quantitative ethnographic map of human cultural variation. Such a map would bring anthropology and economics, perhaps the two most distant of the social sciences, closer together and probably bring psychology, political science, and sociology into the fold along the way. After all, 17 researchers participated in the Henrich et al. project. They included a thorough mix of anthropologists, economists, psychologists, and evolutionary ecologists. When was the last time you saw that in the social sciences? We take the Henrich et al. project as exemplary in this regard. Many important questions await attention along the boundaries between disciplines. A few years ago, one might have asked what such a collection of 17 researchers would have found to do together. Remarkably, tellingly, and perhaps unsurprisingly, the answer now seems obvious.

Historical nonlinear dynamics

“Mathematical modeling is a key ingredient in this research program because quantitative dynamical phenomena, often affected by complex feedbacks, cannot be fully understood at a purely verbal level.”

Turchin (2003)

“Thus, our task ... is to translate various postulated ... mechanisms into models, determine whether these models are in principle capable of generating second-order dynamics, and, if so, derive testable predictions from them.”

Turchin (2003)

“I attempted mathematics, and even went during the summer of 1828 with a private tutor (a very dull man) to Barmouth, but I got on very slowly. The work was repugnant to me, chiefly from my not being able to see any meaning in the early steps of algebra. This impatience was very foolish, and

in after years I have deeply regretted that I did not proceed far enough at least to understand something of the great leading principles of mathematics, for men thus endowed seem to have an extra sense.”

Darwin's *Autobiography*

Evolutionary ecology has recently played a prominent role with regard to some of the big questions in the study of human history. Diamond (1997), for example, examined how seemingly irrelevant details such as the orientation of continental axes and the global distribution of wild big-seeded grasses could have contributed to the eventual rise of the European colonial powers. He recently followed with an analogous approach to the study of societal collapse (Diamond 2005). Additional research has employed human evolutionary ecology to focus on questions related to the rise of agriculture (Richerson et al. 2001), the path-dependent nature of history (Boyd and Richerson 1992), the evolution of complex social institutions (Bowles 2004; Richerson and Boyd 2005), and the dynamics of technological innovation and diffusion (Henrich 2001, 2004).

The theoretical ecologist and history buff Peter Turchin recently joined this tradition with his book on the rise and fall of agrarian empires (Turchin 2003). Many historians have asked why historical agrarian societies often grew into vast empires only to collapse eventually. These single oscillations represent the empirical regularity historians seek to explain. Many of their theories posit endogenous mechanisms that drive the imperial rise and fall. In short, effects related to expansion itself eventually yield collapse; collapse does not come from completely outside the system. Virtually all of this traditional historical theorizing is nonmathematical. In this setting an exclusive reliance on rhetoric carries a huge methodological risk. Simply put, Turchin's overarching point is that, if the task is to explain single oscillations via endogenous mechanisms, the mathematics of nonlinear dynamics must play a role. Without exogenous variables driving the system, a linear model only produces monotonic dynamics. Variables either rise or fall through time. A system that rises and then falls is not possible. Thus the endogenous expansion and decline of agrarian empires is inherently a problem in nonlinear dynamics. Further, strictly verbal arguments are dubious for two reasons. First, verbal assumptions that seem precise often support multiple mathematical formulations, and the differences matter. If an assumption happens to be sufficiently precise to support only one mathematical expression, our hypothetical historian is probably engaging in the difficult, bizarre, and completely unnecessary task of verbal mathematics. Even Darwin, who is especially notable for his lucid intuition unaided by mathematics, knew keenly the limitations of strictly verbal inquiry, as the quote at the top of this section reveals. Second, because we cannot all be Darwin, human intuition is in general spectacularly unreliable at inferring the aggregate dynamics of variables in a nonlinear system. One could thus spend considerable time with a verbal theory that appears plausible but when translated into the language of nonlinear dynamics simply cannot

produce the hypothesized result. If the verbal theory supports an alternative formulation that does produce the hypothesized result, one can ask no more. Math has done its job, at least for the moment, and we're off and running.

As Turchin demonstrates, a shift to a nonlinear dynamical theory in history can yield immediate payoffs. For example, we know from other disciplines such as theoretical ecology that oscillations require delayed feedbacks. As a consequence, before we derive our first set of equations for the rise and fall of agrarian empires, we already know our system must have some kind of delay. We simply cannot obtain the required boom and bust otherwise. To demonstrate, let $x(t)$ be a variable that denotes imperial scope as a function of time, and let $\dot{x} = f(x(t))$ be some dynamical model. The simplest approach to endogenous expansion and decline implies $\exists \bar{x} > 0$ such that $x < \bar{x} \Rightarrow f(x) > 0$ and $x > \bar{x} \Rightarrow f(x(t)) < 0$. In other words, when scope is large it gets smaller, and when scope is small it gets larger. The boundary between large and small is \bar{x} , the unique non-trivial steady state. Although such a model can produce both expansion and decline endogenously, it cannot produce even a single oscillation. When scope is above its steady-state value, the negative feedback is immediate. An oscillation requires overshoot, and overshoot requires a delay in feedback. Essentially, some mechanism has to be in place such that, even though \bar{x} continues to be the steady-state value, scope declines only after $x > \bar{x} + \Delta x > \bar{x}$ for some sufficiently large Δx . Because of the delayed feedback, imperial scope continues to expand for a while even though the empire is larger than the sustainable steady state. Consequently, when the negative feedback finally does take hold, it does so with a vengeance because the society is overextended. The society does not just decline; it crashes suddenly while at the height of its power. Without such a mechanism, theory can provide no basis for understanding the boom and bust of an agrarian empire. One way to delay a feedback is to introduce an additional variable that mediates the effect. Expanding our system to include two variables, assume

$$\begin{aligned}\dot{x} &= f(x, y) \\ \dot{y} &= g(x, y),\end{aligned}\tag{1}$$

where $g_x > 0$ and $f_y < 0$. In this system x has a positive effect on y , which has a negative effect on x , thus further indicating that x has a negative effect on itself via y . If this process of mediated effects produces a sufficiently long delay, oscillations result. The delay has to be long enough, however, and this fact accentuates the importance of relative temporal rates of change in nonlinear dynamics.

Turchin applies exactly this kind of reasoning to geopolitical theories of imperial collapse. Geopolitical theories postulate several relationships between variables, but one negative feedback loop in particular has the potential to yield endogenous decline. An empire's success in war increases its territory. Territorial increase further increases the difficulties of governing the far reaches of the empire, which via the resulting logistical burden decreases the empire's success in war. Turchin argues, however, that this mediated negative feedback

loop is not sufficient by itself to generate the boom and bust cycle he's looking for because the relative time scales are not right. Consequently, his mathematical model of geopolitics does not produce oscillations. In fact it reduces to a single ordinary differential equation in the end. Once again experience with nonlinear dynamics from other disciplines immediately tells us that such a system cannot oscillate. A one-dimensional system must have a delay to oscillate (Murdoch et al. 2003). Turchin (2003, p. 19) concludes, "The feedback loop involving logistic loads ... should operate without an appreciable time lag. For example, the effect of distance ... is instantaneous: as soon as territory is conquered, the boundary moves away from the center, and the state immediately incurs the higher costs of projecting its power to the new boundary."

Turchin essentially follows with his own pursuit of the appropriate lag for the rise and fall of agrarian empires, and he places a particular emphasis on the effectiveness of collective social action as structured by ethnicity. As he acknowledges, the durability of his specific models remains a question for the future. Any model with the right kind of lag can produce a single oscillation. Turchin's argument that logistic loads operate without a significant lag, for example, may one day prove wide of the mark. Nonetheless, by insisting on the language of nonlinear dynamics he raises the precision of the debate and demonstrates the usefulness of formal techniques in even traditionally non-mathematical territory like history. In the specific matter of agrarian empires, Turchin's use of nonlinear dynamics takes us to the center of the problem with an efficiency that verbal arguments typically do not achieve. Namely, in a closed system the boom and bust of a characteristic agrarian empire requires a negative feedback that operates with a sufficiently long lag. Thus, in spite of the complexity that certainly characterizes the historical rise and fall of empires, the conclusion still follows. If a society booms and busts, either the system is not closed in some critical way, or a negative feedback operates with a sufficiently long delay.

More generally Turchin exemplifies the pervasive significance of relative time scales in nonlinear dynamic systems. The dimensional analysis of such systems (Nisbet and Gurney 1982; Gurney and Nisbet 1998) is an extremely useful but underexploited analytical method that, among other advantages, formalizes the importance of relative time scales. If we slightly modify the notation of (1) and explicitly note the presence of two parameters, a and b , that are proportional to $[\text{time}]^{-1}$ and have the same units, the system is

$$\begin{aligned}\frac{dx}{dt} &= f(x, y; a) \\ \frac{dy}{dt} &= g(x, y; b).\end{aligned}\tag{2}$$

Dimensional analysis stems from Buckingham's (1914) π Theorem, and the central idea comes naturally. Empirical phenomena do not depend on the units of measurement. A given woman is a particular height regardless of whether we

measure her in centimeters or inches. Thus quantities with units, as in absolute temporal rates, cannot control the dynamics of system (2). If, however, we define $x = \hat{x}\chi$, $y = \hat{y}\psi$, and $t = \hat{t}\tau$, where χ , ψ , and τ are pure numbers without units, simple algebra shows that (2) is equivalent to the non-dimensional system

$$\begin{aligned}\frac{d\chi}{d\tau} &= \frac{\hat{t}}{\hat{x}}f(x, y; a) \\ \frac{d\chi}{d\tau} &= \frac{\hat{t}}{\hat{y}}g(x, y; b).\end{aligned}\tag{3}$$

Such an exercise typically results in the conclusion that the non-dimensional parameter a/b controls (2) and by extension (3). This conclusion is a direct consequence of the fact that relative time scales matter. Technically, a/b is a quantity with a unitless component [time]/[time].

Experimental nonlinear learning

“Most of the research on learning and evolution is theoretical, but it is unlikely that theorizing alone will explain how people learn without the input of empirical regularity and careful testing.”

Camerer (2003)

“[I]mitation is most compelling when players are symmetric, and not at all sensible when players are asymmetric...Imitation of successful players who are similar is probably a heuristic way of moving toward strategies with high forgone payoffs.”

Camerer (2003)

Linear social learning, as we now define it, occurs when individuals adopt a particular behavior via learning from conspecifics with a probability equal to that behavior’s frequency in the population. For example, assume two alternative behaviors, 0 and 1, are present, and denote the behavior of individual b in period t by $c_{b,t}$. Call the frequency of behavior 1 in a period q_t . Linearity means that for each b one can on average summarize the social learning process with the following proposition:

$$P(c_{b,t+1} = 1) = q_t.$$

In essence the probability a given individual adopts behavior 1 in the next period is equal to the proportion of the population exhibiting that behavior in the current period. Linearity is important because social learning, if linear, simply reproduces the existing structure of the population from one period to the next. The aggregate dynamics of behavior are neutral with respect to linear social learning. In this case changes in q_t cannot arise from social learning, but only from other processes like chance or individual learning.

The importance of linearity also points toward the importance of nonlinearity. If one cannot summarize social learning with the above proposition, the mechanics of social learning generate endogenous evolutionary dynamics at the population level and can do so quite apart from properties intrinsic to the behaviors in question. The distinction between linear and nonlinear social learning is central to gene-culture coevolutionary theory and its hypothesis that humans use cognitive biases that structure social learning (Boyd and Richerson 1985, 2005; Richerson and Boyd 2005). Specific biases are essentially hypotheses about the form of nonlinear social learning.

A basic difficulty underlies the empirical study of social learning. We would like to know if social learning is linear or nonlinear, and we would like to know how social learning, if nonlinear, deviates from linearity. The difficulty arises from a common property of nonlinear dynamical systems. Small changes can have big effects. The consequences of this fact for the empiricist are substantial. To illustrate with a whimsical example, imagine a population of nonconformists who deviate from linear social learning by exhibiting a tendency to adopt the least common behavior in the population. We can model nonconformity as

$$P(c_{b,t+1} = 1) = \lambda q_t + (1 - \lambda)\{q_t + q_t(1 - q_t)(1 - 2q_t)\}, \quad (4)$$

where λ is a parameter that weights the relative importance of linearity and nonconformity. We can make the nonconformist tendency arbitrarily small by letting λ approach 1. As long as $\lambda < 1$, however, no matter how close to 1 it is, the long-term aggregate behavior of the system is completely different from the linear case in which $\lambda = 1$. In the linear case, the system never moves. For any value of q_t the system stays at that value forever.³ In the nonlinear case, in contrast, no matter how weak the nonconformity, the system has three steady states at $\hat{q} = 0$, $\hat{q} = 1$, and $\hat{q} = 1/2$. The steady state at $\hat{q} = 1/2$ is globally stable. In other words, given enough time and any mixture of behaviors, our theoretical population of nonconformists must end up with half the population exhibiting behavior 0 and half exhibiting behavior 1. Once this state obtains, it persists forever. More to the point, a seemingly trivial deviation from linearity at the individual level does not necessarily have trivial consequences at the aggregate level. Our example is similar to a classic result from population genetics. Even when selection is weak at the individual level, populations respond rather rapidly on times scales of a relatively few generations.

If an evolutionary social scientist wanted to study nonconformity, what should she do? She could try to measure nonconformist deviations from linearity at the individual level with carefully designed behavioral experiments, but nonconformity can be arbitrarily small and still have the same aggregate consequences given enough time. Reaching the steady state takes longer as

³ We ignore the other forces that could cause changes in order to focus on the critical difference between linear and nonlinear social learning.

nonconformity gets weaker, but the equilibrium prediction does not change. Consequently, nonconformity can be too weak to measure at the individual level but nonetheless present and thus important at the aggregate level. The alternative empirical approach would involve a study of aggregate dynamics. In this case theory predicts the system should tend toward $\hat{q} = 1/2$. Unfortunately, the model

$$P(c_{b,t} = 1) = 1/2 \tag{5}$$

makes exactly the same prediction, as do numerous other models we could imagine. Transitional dynamics can vary. Model (5), for example, takes one period to reach its steady state from any initial condition, whereas model (4) can take anywhere from few to several periods. Nevertheless the fact remains that overlapping predictions among models at the aggregate level weaken the empiricist's ability to understand what she observes.

Fitting models derived from first principles to empirical data provides a partial solution. This method is firmly established in the experimental economics literature on learning (e.g. Camerer and Ho 1999), and Camerer (2003) reviews much of the literature. The objective is to identify the model among a set of models that fits the data at hand best. The data in this case take the form of a sequence of behaviors or choices in an economic experiment. Questions of particular interest concern how people learn individually and socially. In other words, imagine participants in a setting where they know some behavior is optimal, but they do not know which behavior it is. Perhaps payoffs involve a random component, and thus a short sequence of choices does not reliably reveal the optimal choice. Alternatively, the problem may be difficult because participants face a large number of options but only a limited opportunity to choose. Such a situation is kind of like subsistence farming. Choices matter with respect to crop yields, but at one or two crops a year any given farmer simply cannot try all the alternatives available to him. Should he plant his seeds on March 15 or March 25? Should he plant his traditionally favored landrace of corn in his most distant field, another variety in his second most distant field, and beans in his near field, or should he consider some other combination? When considering all orthogonal behavioral dimensions, the resulting space is impossibly large for one farmer to explore. In essence, all experiments of learning are like our farmer. When any given player does not know what to do, how does she explore the relationship between the behavioral and payoff spaces facing her?

Experiments and model fitting offer a controlled approach to this question that holds considerable promise. Experimenters can control the information available to players through time and, in conjunction with subsequent model fitting, examine how people use private and culturally transmitted information. Moreover, this method can shed light on the crucial matter of whether information use is linear or nonlinear and what the effects are with respect to aggregate behavioral dynamics and cultural evolution. One can also

incorporate contingent biases or biases that change through time. All such principles can inform the derivation of a model to fit to a dataset.

In the early days of cultural evolutionary theory (e.g. Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985), the question of how to develop an associated empirical program remained largely unaddressed. Recently, however, an incipient research program combining cultural evolutionary theory and experimental economics (Baum et al. 2004; McElreath et al. 2005; Efferson et al. submitted) has begun to address some of the basic questions in cultural evolution. We contend that not just some but every general theoretical point in cultural evolutionary theory can be addressed, at least in part, with an appropriate economic experiment that controls the flow of individual and social information through time as each player makes a sequence of decisions.

Model fitting should play a critical although not exclusive role in such an enterprise. Model fitting, in essence, allows the researcher to take theoretical models summarizing the consequences of dynamically updated private and social information and fit them to the data. Once models have been fit to a given dataset, one can identify the model that fits best. The epistemological assumption is that the best fitting model best summarizes the behavioral consequences of dynamically updated information. We say much more about this method below, but here we point out that in complex behavioral situations such a model is never right. It can, however, be better than other models under consideration. Our suspicion is that the better models will carry critical nonlinearities pertaining to the use of social information. Critical nonlinearities, in turn, imply that the structure of social learning almost certainly matters, that success today can yield endogenous collapse tomorrow, or that success today might yield disproportionately greater success tomorrow, and that behavioral, economic, and institutional dynamics involve a certain erratic flair that presents the scientist with unique but surmountable difficulties.

Experimental nonlinear dynamics

“[I]n nonlinear systems one should not necessarily expect data to be simply ‘fuzzy’ versions of some attractor. Instead, one is likely to see dynamics that are mixtures of attractors, transients, and even unstable entities.”

Cushing et al. (2003)

“One of the features of nonlinearity is that responses are not (necessarily) proportional to disturbances.”

Cushing et al. (2003)

Now assume our population of social learners from above is infinitely large, and thus we can ignore individual choice probabilities and model changes in q ,

directly. In particular, posit dynamics that change deterministically according to some weighted combination of linearity and conformity,⁴

$$q_{t+1} = \lambda q_t + (1 - \lambda)\{q_t + q_t(1 - q_t)(2q_t - 1)\}. \quad (6)$$

In short, with probability λ individuals maintain their current probability of adopting behavior 1 in the next period, while with probability $1 - \lambda$ individuals exhibit a tendency to select the most common current behavior in the next period. If $\lambda < 1$, in the long run this system will persist in one of only three states. If the initial frequency of behavior 1 in the population is less than one half, the long-run frequency is $\hat{q} = 0$. If the initial frequency is greater than one half, the long-run frequency is $\hat{q} = 1$. Lastly, if the initial frequency is exactly one half, then it will remain so indefinitely.

This enumeration of steady states, however, glosses over their stability properties. Models (5) and (6) have the same steady states, but their stability properties are exactly the opposite of each other. Under the conformity model, (6), the steady states $\hat{q} = 0$ and $\hat{q} = 1$ are locally stable. They are attractors in the sense that, if the system starts sufficiently close to the attractor, the system moves toward the attractor. For example, if we start with an initial value at $t = 0$ of $q_0 = 0.2$, the system tends toward the $\hat{q} = 0$ steady state through time. An initial value of $q_0 = 0.8$, in contrast, means the system tends toward the $\hat{q} = 1$ steady state. The steady state $\hat{q} = 1/2$, however, is unstable. For any value of $q_t = 1/2 \pm \varepsilon$, where $\varepsilon > 0$, the system moves away from the steady state $\hat{q} = 1/2$ no matter how small ε .

In a deterministic world of infinite populations and perfect predictions, this distinction causes no practical problems. One need only characterize the steady states and initial conditions, and all falls into place. Alternatively and more realistically, imagine a setting where model (6) yields highly accurate but not perfect predictions. In this case the residual variation is small, and for many purposes the deterministic model (6) might serve admirably. In effect the deterministic signal is so strong relative to any remaining effects, which we may choose to model stochastically, that few interpretive problems arise. If given several replicate data sets with initial conditions $q_0 = 1/2 - \varepsilon$ for some small $\varepsilon > 0$, we recognize some probability that a given time series will tend toward $\hat{q} = 1$ rather than $\hat{q} = 0$. But as long as ε is not too small, and as long as determinism is strong enough, we nonetheless expect the time series to tend toward $\hat{q} = 0$ in general. In other words, in the long run we expect dynamical data to be “fuzzy” versions of some appropriately modeled deterministic attractor.

⁴ A model such as (6) cannot be fully deterministic in a finite population. Feasible values of q_t will form a lattice in the sense that, in a population of size N , the variable q_t can only take $N + 1$ different values. Model (6) places no such restrictions of the possible values of q_t , and thus it can only be approximately deterministic if N is finite. In a nonlinear experimental environment, one should not trivialize this distinction as it can prove critical to understanding the predictive power of a model (see Cushing et al. 2003, ch. 5). For present expository purposes, our assumption of an infinitely large population, among other advantages, allows us to ignore lattice effects.

But what if residual variation is large? This difficulty is intrinsic to high-dimensional systems. Many phenomena involve dozens, hundreds, or even thousands of relevant causal variables. As scientists, however, we need to restrict the set of variables under consideration in some way. Given a particular dataset, restricting the set of explanatory variables increases the residual variation. In a static linear environment, this increased residual variation may not hamper insight. A linear trend can remain a linear trend even in the face of substantial increases in residual dispersion. In a dynamical nonlinear environment, however, small changes can have big effects. In model (6), for instance, $q_t=0.49$ yields a prediction about the future state of the system that is wildly different from the prediction under $q_t=0.51$. As a consequence, if residual variation is large enough, one essentially loses the ability to predict if one only focuses on deterministic attractors. This difficulty is intrinsic to high-dimensional nonlinear dynamical systems. High dimensionality ensures that a tractable and thus typically low-dimensional model leaves a noteworthy amount of residual variation. Nonlinear dynamics, in turn, ensure that small differences, like those due to residual variation, can yield disproportionately large differences in future dynamics. Together, these two forces suggest that empirical time series need not be “fuzzy” versions of a deterministic attractor even if this attractor derives from a model that summarizes important deterministic forces in the empirical system.

Consider, for example, the hypothetical time series in Figure 1. An unconsidered focus on the attractors of (6) would lead us to conclude this model is fundamentally incompatible with the hypothetical data in the figure. As a strictly deterministic model, conformity (i.e. model (6)) offers no basis for making sense of such a time series. As discussed earlier, conformity implies that when aggregate dynamics are near $q_t=1$, the system moves toward 1, not away from it. In Figure 1, however, aggregate dynamics reach $q_t=1$ only then to

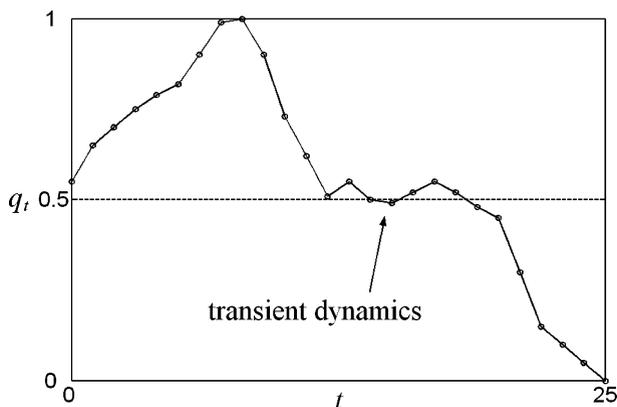


Figure 1. A hypothetical time series fundamentally incompatible with a deterministic conformity model but perfectly compatible with a stochastic conformity model based on conditional predictions.

move away and end up, of all places, at $q_t=0$. Whatever Figure 1 is, it's not simple deterministic conformity. We could, however, alternatively treat (6) as a conditional expectation for a nonlinear stochastic model. For example, assume we model the probability that in the next period A_{t+1} out of N individuals will choose behavior 1, conditioned on the frequency of 1 in the current period, q_t . Then our stochastic nonlinear dynamical model with conditional predictions takes the form

$$P(A_{t+1}|q_t) \sim \text{binomial}(q_t + (1 - \lambda)q_t(1 - q_t)(2q_t - 1), N). \quad (7)$$

This model, although stochastic, preserves the structure of (6) under conditional expectations in the following sense,

$$E[q_{t+1}|q_t] = \frac{1}{N}E[A_{t+1}] = q_t + (1 - \lambda)q_t(1 - q_t)(2q_t - 1).$$

The difference in terms of confronting data (e.g. Figure 1) with models, however, is fundamental. Unlike the deterministic model (6), under model (7) stochasticity can move the system away from one of the two attractors. Moreover, under the right circumstances the system can end up near the unstable steady state at $\hat{q} = 1/2$. If so, this unstable entity can affect the dynamics of the system transiently as in Figure 1. The issue here is that, although the unstable steady state is unstable, it is still a steady state. In effect, the forces pulling the system toward either of the two attractors are weaker here. Consequently, transients can result. Strictly speaking, transients are any dynamical behavior that is not final (Hastings 2004). Here we focus on temporary dynamics that occur when the system lingers near an unstable steady state because the forces pulling it away are weak. In this case, one can show that as λ increases the deterministic forces drawing the system away from $\hat{q} = 1/2$ grow weaker for any value of q_t . Near the unstable steady state, these forces can indeed be arbitrarily small, and this fact will protract the transient dynamics in the vicinity of the unstable steady state. Equipped with these ideas, the time series in Figure 1 is at least interpretable under the deterministic forces summarized by model (6). When we embed these nonlinear deterministic forces in a stochastic model as in (7), we can tentatively conclude conformity effects are compatible with the data. Initially the system tends toward the attractor at $\hat{q} = 1$, but subsequent stochastic effects place the system near the unstable entity, $\hat{q} = 1/2$, which in turn affects dynamics transiently as q_t lingers near $1/2$ under the potentially weak forces attracting the system to other steady states. At some point the system enters the basin of attraction for $\hat{q} = 0$ and escapes the transiently negligible effects of conformity near the unstable steady state. Deterministic conformity forces come to dominate again, and the system tends toward the attractor at $\hat{q} = 0$. To see this possibility, however, we have to consider determinism and stochasticity together. Dynamical nonlinear determinism produces a complex landscape of attractors and unstable entities. Stochasticity adds uncertainty on top of this landscape and so destabilizes

trajectories but correspondingly gives us a means of interpreting otherwise evasive patterns.

In practical terms, these kinds of nuances are intuitive. We recognize that the deterministic model (6) yields a sharp distinction between $q_t=0.49$ and $q_t=0.51$, but we can also imagine the response of an actual group of 100 conformists with two alternative behaviors. If, for some reason, they end up with 49 out of 100 choosing behavior 1 at a given point in time, and if they know this fact, we still would not be surprised if the system moved to 100 out of 100 choosing behavior 1 instead of 0 out of 100. A large, probably extremely large, number of unconsidered forces affecting decision making could produce this outcome even if conformity is important. Under some degree of conformity the probability the system will move toward 0 is greater than the probability it will move toward 100. No matter how small this tendency, if we could replicate the precise situation above enough times, we could detect the effects of conformity simply by counting replicates. Nonetheless, as stochasticity becomes more important as we substitute low-dimensional systems for high-dimensional reality, and as threshold effects enter the fray due to nonlinear feedbacks, the number of necessary replicates becomes impractically large. The crux lies here; the dynamical interaction of nonlinearity, determinism, and strong residual stochasticity presents unique problems to those who would integrate theory and empiricism in the evolutionary social sciences.

Interestingly and not coincidentally, population biologists face the same difficulties. Recently a group of ecologists and applied mathematicians, Jim Cushing, Robert Costantino, Brian Dennis, Robert Desharnais, Shandelle Henson, and Aaron King (collectively known as “The Beetles”), have developed a provocative set of tools for integrating theory and empiricism in experimental nonlinear stochastic environments. They study the population dynamics of the flour beetle, *Tribolium castaneum*, and their task is a demanding one. Specifically, beetles in the genus *Tribolium* exhibit a taste for eating conspecifics, and this cannibalistic tendency creates strong nonlinear interactions between the different stages in the life cycle. The resulting range of dynamical behaviors in controlled laboratory environments runs from simple steady states to chaos and in between includes an exotic and surprising collection of intermediate dynamical regimes. Cushing et al. (2003) summarize a decade’s worth of research to develop and use a single low-dimensional model to predict, qualitatively and quantitatively, the myriad dynamical regimes they observe in their experimental *Tribolium* populations. Their achievement is stunning. Figures 2 and 3 reprint two of the figures in their book. The figures summarize some of their results, offer a comparison between the predictions of their model and the observed population dynamics, and finally stand as two of the most remarkable figures we have encountered in the ecological literature. This kind of predictive accuracy rarely, if ever, happens in ecology, and it is a testimony to how far a sustained research program in experimental nonlinear dynamics can go.

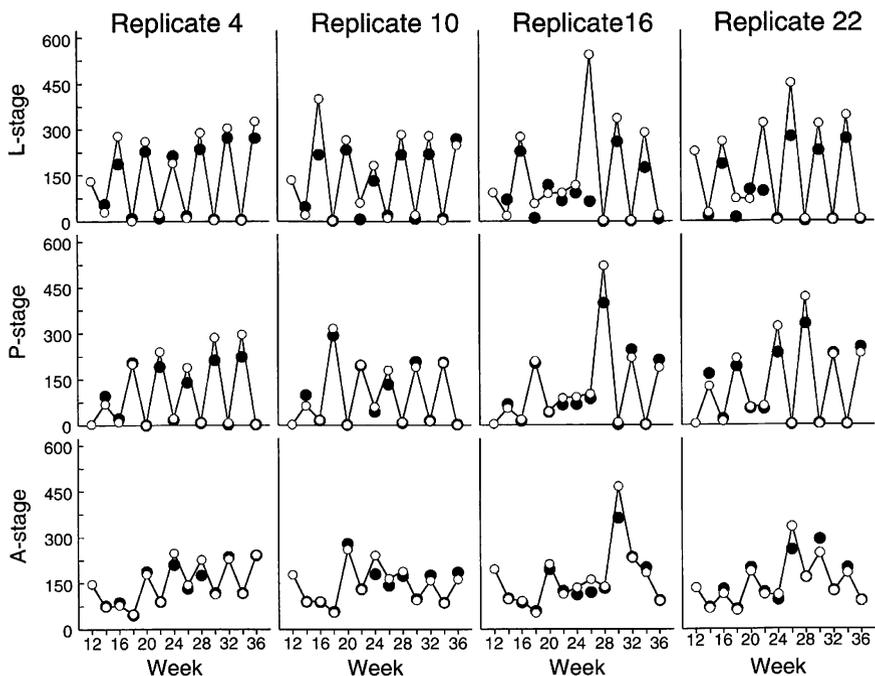


Figure 2. The census data (open circles with lines) for the SS strain of *Tribolium castaneum* and the one-step conditional predictions (solid circles) derived from the three-dimensional model of Cushing et al. (2003). The three stages of the life cycle shown are the larval (L) stage, the pupal (P) stage, and the adult (A) stage. (From J. M. Cushing, R. F. Costantino, B. Dennis, R. A. Desharnais, and S. M. Henson. 2003. *Chaos in Ecology: Experimental Nonlinear Dynamics*. Academic Press. Reprinted with permission from Harcourt, Inc.)

With regard to the study of human evolutionary ecology, and in particular a nascent experimentalism in cultural evolution (Baum et al. 2004; McElreath et al. 2005, Efferson et al. submitted), the importance of Cushing et al. (2003) is primarily methodological. Many of the methodological implications also apply more broadly. History, for instance, may indeed be one damn thing after another, but experimental *Tribolium* populations substantiate the importance of theory in such a world. The achievement of Cushing et al. is impressive largely because they have become so adept at analyzing and even predicting the behavior of a system with such elusive patterns complicated by history. Although they limit themselves to beetle dynamics in the laboratory, they proceed with equal facility whether facing transient dynamics, simple dynamics like a stable steady state, or chaos, the hallmark of dynamical complexity. Far from dismissing the study of flour beetles as irrelevant in the social sciences, we take their work as a challenge to behavioral scientists of all kinds. Additionally, we take the work of Cushing et al. as pointing toward the methodological potential of stochastic nonlinear dynamics. The study of

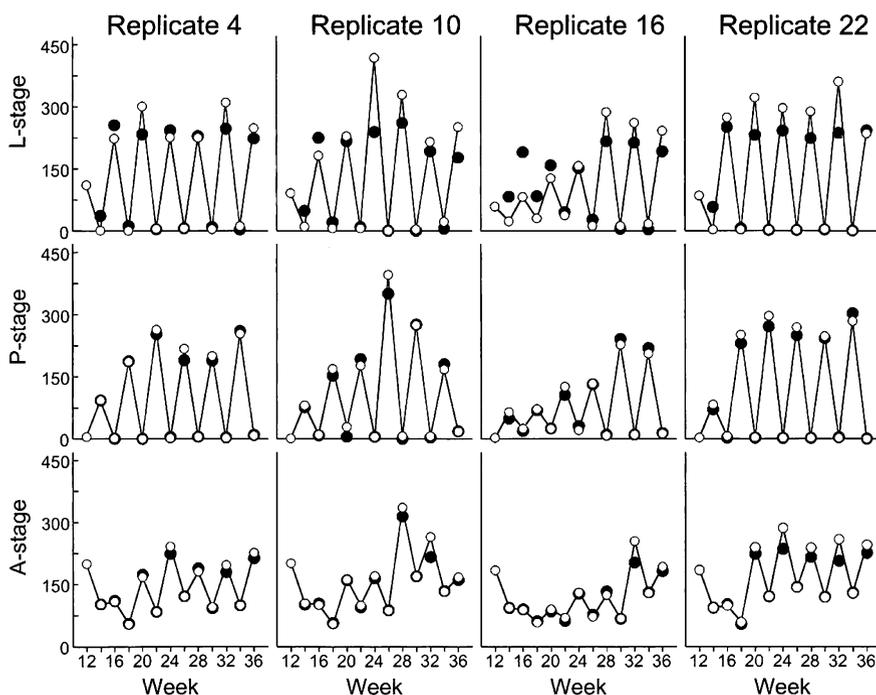


Figure 3. The census data (open circles with lines) for the RR strain of *Tribolium castaneum* and the one-step conditional predictions (solid circles) derived from the three-dimensional model of Cushing et al. (2003). The three stages of the life cycle shown are the larval (L) stage, the pupal (P) stage, and the adult (A) stage. (From J. M. Cushing, R. F. Costantino, B. Dennis, R. A. Desharnais, and S. M. Henson. 2003. *Chaos in Ecology: Experimental Nonlinear Dynamics*. Academic Press. Reprinted with permission from Harcourt, Inc.)

human behavior has all the hallmarks of a discipline in need of stochastic nonlinear dynamics as a central part of its methodological repertoire. Many variables interact through time at different rates, and the result is often large residual variation, evanescent patterns, thresholds, and without the discipline of applied mathematics a seemingly inexplicable procession of one damn thing after another. We describe some of the key methods of Cushing et al. below and discuss their implications specifically with regard to the experimental study of social learning and cultural evolution.

Low-dimensional models from first principles

Many phenomena in biology and the social sciences are high-dimensional in the sense that even simple systems involve more causal variables than we would like to address. As a consequence, a key question, and one that will figure prominently in our discussion of model selection below, is the question of how to find a suitable low-dimensional surrogate to examine high-dimensional

phenomena. All else being equal, in nonlinear stochastic settings a mechanistic model can typically do more with less than a corresponding phenomenological model. Loosely put, a mechanistic model is one derived from specific assumptions about how the system works. It is based on ideas about the mechanisms that govern the phenomena in question. A phenomenological model, in contrast, is not based on specific assumptions about the system but rather on the fact that one can always use certain kinds of mathematical functions to summarize data without any real regard for the processes that generated the data. Standard statistical approaches such as linear regression provide examples of phenomenological models that should be familiar to most social scientists. Phenomenological models, though they bring the advantage of flexibility across various applications, also bring important costs. In particular, as observed time series become more complex, phenomenological models typically require an increasingly opaque array of nonlinear terms to capture the variation. In the limit, one can end up with a senseless model that does an admirable job of summarizing observed variation. Mechanistic models, in contrast, derive from specific assumptions about causation. They can be far less flexible and possibly even unique to the application at hand. But if the phenomenon under study includes distinctive causal mechanisms, an analogously distinctive approach to models can take one to the heart of the mechanisms in a way that phenomenology cannot.

Consider a person who rarely cooks but needs at least one knife in the kitchen. An 8" cook's knife would be an obvious choice and would probably serve our reluctant cook admirably, even though it's terrible for filleting fish and even worse for making rosettes out of tomato skins. The reason it would serve admirably is because, in spite of its limitations, it is a flexible all-around design for a kitchen knife. A professional chef, however, would find the same knife unacceptable in many situations. A professional chef would want a large selection of knives so she can choose one suited to the present task. Pure phenomenology in science is like a professional chef who only has an 8" cook's knife.

Conditional predictions

Because small differences in parameter values and the values of state variables can yield large differences in dynamical behavior in stochastic nonlinear models, one needs some way to prevent even small errors in prediction from accumulating through time. Nonlinearity increases the chances that, even if a model predicts well, compounding errors of this sort will lead to predictions that diverge wildly from the observed time series. As a consequence, Cushing et al. (2003) rely exclusively on conditional predictions. They take the state of the system in one time period as given by the actual observed state and predict one period forward only. In essence, if $X_t \in R^N$ is a vector of random state variables with expected values determined by $f: R^N \rightarrow R^N$, predictions take the following form,

$$E[X_{t+1}|x_t] = f(x_t),$$

where x_t are the observed values of the state variables at t . In essence, one uses the observed values at $t-1$ to predict the values at t , and then one uses the observed values at t to predict the values at $t+1$. This resetting process continues recursively over the entire dataset and allows one to estimate parameters in the model and make predictions without the potentially serious consequences of stochastic errors accumulating in a nonlinear environment through time. The approach is similar to weather prediction. Each day the weather services predict weather for the next seven days or so. But the forecasting models are not run every seven days. Rather every day fresh observations are fed into the model and a new forecast based on up-to-date data is computed. In principle, one would predict forward in time as far as one reasonably could without accumulating errors due to nonlinear feedbacks. *A priori*, however, this question can be difficult or impossible to address. Moreover, the answer will often depend specifically on the current state of the system. Thus single-period predictions represent a useful starting point, and Cushing et al. demonstrate just how far one can go with such predictions.

Different deterministic entities and their effects

As our exercise with Figure 1 demonstrates, when stochasticity is important in nonlinear empirical settings one cannot focus solely on the attractors of a given model. Deterministic entities with different stability properties can interact with stochastic effects to yield time series with a historically contingent character. In Figure 1, for example, we imagined a population of conformists that ended up at $q_t = 1$. Subsequently, stochastic effects placed the population near the unstable steady state at $\hat{q} = 0.5$. Additional small stochastic fluctuations accumulated and the system eventually entered the basin of attraction for $\hat{q} = 0$. This steady state is exactly where the population ended up, a steady state with no one exhibiting the behavior in question (i.e. $\hat{q} = 0$). Nonetheless, when a population is near the unstable steady state, the probability it will enter the basin of attraction for $\hat{q} = 1$ is equivalent to the probability it will enter the basin of attraction for $\hat{q} = 0$. An otherwise identical population of conformists could thus end up with everyone exhibiting the behavior in question ($\hat{q} = 1$). The two populations would be identical in terms of how individuals make decisions, but the historical interaction of stochasticity and determinism would produce two populations as far apart behaviorally as possible. Cushing et al. (2003) repeatedly find evidence for this kind of historical contingency, and Boyd and Richerson (1992) discuss it specifically with respect to cultural evolution. The solution for Cushing et al. is to model deterministic and stochastic forces together in ways rooted firmly in the biology of their study organism. Moreover, once a given model is in place, they develop various

model-specific approaches to examining the effects of different deterministic entities when stochasticity is important. More generally, Cushing et al. (2003) do not average the time series from replicate experiments in search of average trends. Although such a procedure may be useful in some situations, one must recognize the potential importance of history in a dynamical world, as our own, with a thorough mix of nonlinearity, determinism, and stochasticity. Averaging loses the information inherent in history.

Bifurcation experiments

A bifurcation is a change in the deterministic attractor of a dynamical system as some parameter of interest changes. Although Cushing et al. (2003) repeatedly emphasize the need to expand the study of nonlinear dynamical systems beyond the study of attractors, the fact remains that their core model is amazingly accurate at predicting the dynamics of experimental *Tribolium*

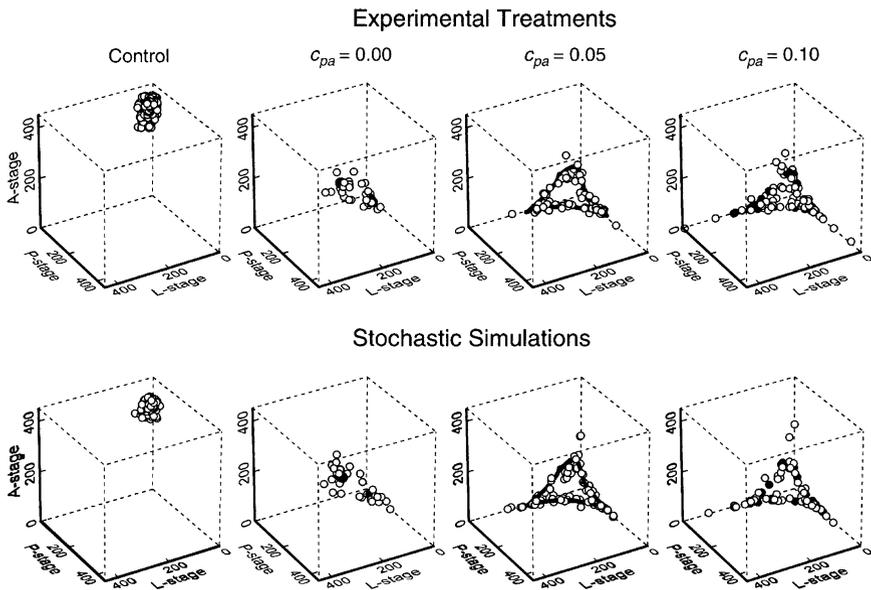


Figure 4. Results from a bifurcation experiments in which c_{pa} , the rate at which adults cannibalize pupae, is being controlled. As this rate increases, the deterministic attractor (solid circles/lines) of the three-dimensional model of Cushing et al. (2003) changes. Additionally, as the parameter is experimentally controlled, laboratory populations respond, and census data are also shown (open circles). Beneath each plot is a corresponding plot (i.e. same parameter values) showing the deterministic attractor and predictions (open circles) based on a stochastic version of the model of Cushing et al. The three stages of the life cycle shown are the larval (L) stage, the pupal (P) stage, and the adult (A) stage. (From B. Dennis, R. A. Desharnais, J. M. Cushing, S. M. Henson, and R. F. Costantino. 2001. Estimating chaos and complex dynamics in an insect population. *Ecological Monographs*. 2: 277–303. Reprinted with permission from the Ecological Society of America.)

populations. Thus an analysis of attractors constitutes a crucial part of their general research program. Moreover, although removing data certainly requires some justification, one can remove transients from a time series if the particular task at hand requires a focus on attractors. In this regard Cushing et al. (2003) make effective use of what they call “bifurcation experiments.” Essentially, given a model of an experimental situation, at least some of the parameters in the model will be under experimental control. One can vary one of these parameters theoretically, and the result is a bifurcation diagram. This diagram is a description of how the set that attracts the dynamical system changes as the parameter of interest changes. One can then select appropriate values of this parameter from the bifurcation diagram to implement as different experimental treatments. If the observed sequence of bifurcations matches the predicted sequence, the bifurcation experiment represents a severe and successful test of the theoretical model. Figures 4 and 5 show the results from one of the bifurcation experiments described in Cushing et al. (2003) and are some of the few figures from the ecological literature we find more impressive than Figures 2 and 3. As the figures show, their core three-dimensional model has a remarkable ability to predict changes in the behavior of experimental *Tribolium* dynamics as some parameter, in this case the rate at which adults cannibalize pupae, changes.

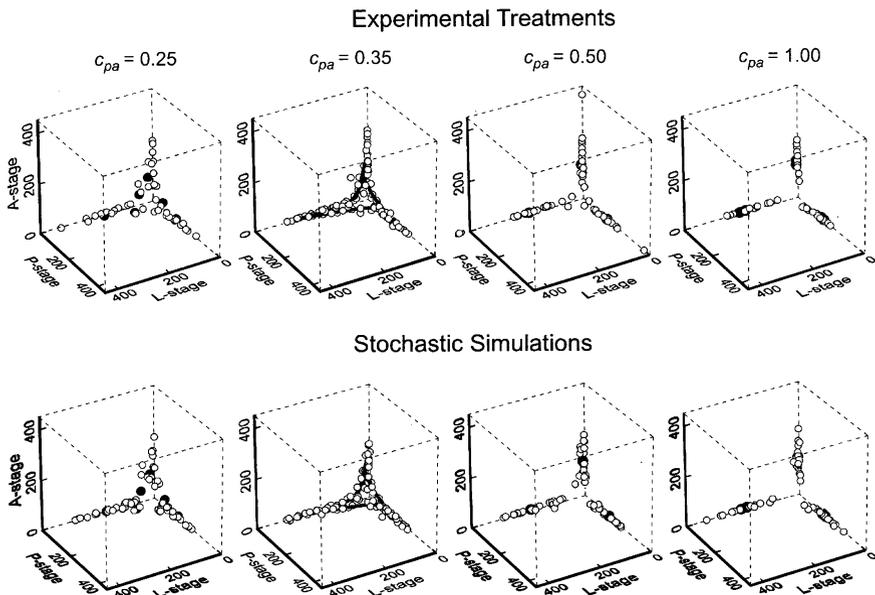


Figure 5. Results from the same bifurcation experiment shown in Figure 4. (From B. Dennis, R. A. Desharnais, J. M. Cushing, S. M. Henson, and R. F. Costantino. 2001. Estimating chaos and complex dynamics in an insect population. *Ecological Monographs*. 2: 277–303. Reprinted with permission from the Ecological Society of America.)

We suspect that the experimental study of cultural evolution is a long way from the successful use of such a technique, but we would like to put it forward as an eventual goal. Before such a goal can be realized, however, bifurcation analysis must become a central part of the theoretical study of cultural evolutionary dynamics.

Model selection and minimized expected relative information loss

“While there are no models that exactly represent full reality ..., full truth can be denoted as f .”

Burnham and Anderson (2002)

“*Truth, f , Drops Out as a Constant*”

Burnham and Anderson (2002)

Without empiricism, of course, models are not even models. They’re just mathematics. Once we accept the need for a link between models and empirical phenomena, we have the remaining question of how to select the best model. This question has three parts. First, how does one choose a set of models to consider? One cannot select an unconsidered model as the best model; it is not under consideration. The art of scientific research resides here, and science proceeds in part as an ongoing modification of the set of models under consideration. This process is “art” in the sense that no formal criteria exist for specifying a set of models to use in data analysis. Indeed, given that the set of potential models is probably infinite for problems of any complexity, the nature of such criteria is far from obvious. The second part of finding the best model is specifying what “best” means. Much of statistical theory implicitly or explicitly addresses this question, and consequently many formal criteria exist for identifying a best model given a set of models under consideration. The final part of the question centers on the technical details. Once we have defined what we mean by a “best” model, we still have to settle the practical matter of how one can apply the criterion to actual data.

Here we focus on the definition of a best model. Take a dataset, x , that depends on some process. Designate the process truth, and label it f . The task is to make inferences about f given the nexus, x . The procedure is to posit a set of R alternative models, g_i , where i is an index in the set $\{1, \dots, R\}$ and specifies one of the R models. Each model has a structure (e.g. functional form), $g_i(\cdot)$, which includes a vector of parameters, θ , and for each set of parameter values the model predicts the values of x . In essence, x (the data) contains information about f , and we use the g_i to model this information. A given model also contains information, and this information depends on both the structure of the model and the specific parameter values used. For example, the models $g_1(x) = 1 + 2x$ and $g_2(x) = -157,309 + 9001x$ have the same structure, but they differ profoundly in terms of informational content because the parameter values are so different. Because the models are models and not f itself, where f

need not even be effable in mathematical or verbal terms, the ability of any given g_i to capture the information in x is invariably fragmentary. Information loss is the inevitable result. This information loss is readily quantified as the Kullback–Leibler (K–L) distance from g_i to f (Burnham and Anderson 2002). The distance is directed in the sense that the K–L distance from a given g_i to f is not the same as the distance from f to the same g_i , a fact that merely indicates we cannot exchange truth and models. A straightforward and intuitive criterion for model selection is to identify the model with the minimum K–L distance to f . Equivalently, we seek the model with the minimum directed information loss.

Two problems present themselves. The first is that we do not know truth, f , and thus we cannot calculate the Kullback–Leibler distance. Although seemingly insurmountable, this problem turns out to be less of an obstacle than one might think. Consider the following heuristic. Call the K–L distance from g_i to f the following,

$$D(f, g_i(x|\theta)) = I(f) - I(g_i(x|\theta)),$$

where $I(f)$ is the information intrinsic to truth, and $I(g_i(x|\theta))$ is the information intrinsic to the i th model under a dataset, x , and a specific set of parameter values, θ . Now consider the K–L distance from another model, g_j , with another specific set of parameter values, Θ ,

$$D(f, g_j(x|\Theta)) = I(f) - I(g_j(x|\Theta)).$$

Although we cannot calculate either of these quantities because we do not know $I(f)$, we can subtract one distance from the other to obtain

$$D(f, g_i(x|\theta)) - D(f, g_j(x|\Theta)) = -I(g_i(x|\theta)) + I(g_j(x|\Theta)).$$

Truth and the information intrinsic to truth drop out of the equation. We are left only with quantities at hand, namely quantities related to the information intrinsic to the models given the data available and the specific sets of parameter values. The compromise we make is that we say nothing about how far we are from truth in an absolute sense. What we gain is that we can say which model is closest to truth, and we can say how much further away the other models are. It's as if Charles and Pete were standing at the tram stop at Bellevue in Zürich with Charles three feet to the south of Pete. Neither of us would know how far we were from Cologne, but we would know that Charles was three feet farther than Pete. To push the analogy, we would not even know if we wanted to go to Cologne. Maybe the destination would be the more distant Copenhagen or perhaps only as far as Stuttgart. Regardless, we would know the objective was to the north, and Pete would be three feet closer than Charles.

The second problem is that the data analyst has to estimate the parameters, θ , in any given g_i using the data, x . Thus the distances above must be estimated distances because the informational content of any g_i depends on both the

structure of the model and the specific parameter values used. The net result of addressing these two problems is Akaike's Information Criterion (AIC). Given a set of models under consideration and appropriate data, AIC allows one to identify the model in the set estimated to have the least information loss that must exist when one substitutes a model for unknown truth, f . In practical terms, AIC amounts to finding the model in the set most likely to fit the signal in the data as it penalizes models with more parameters in an appropriate fashion to prevent fitting the noise as well. In particular AIC optimizes the tradeoff between the competing objectives of reducing bias and reducing variance, the problems of underfitting and overfitting described in the introduction. As one increases the number of parameters one estimates from a given amount of data, the estimates become less biased. In effect, the simpler one's approach to the data (i.e. the fewer parameters estimated), the more likely one is to leave out an important consideration and arrive at a biased conclusion. The solution is to estimate more parameters and thereby increase complexity by asking more questions, which reduces the tendency of the model to underfit the data. As one increases the number of parameters, however, the estimates become more uncertain. Asking a fixed amount of data to answer a larger number of questions means the answers are less conclusive, a common indication of overfitting. Empiricists consequently face a basic dilemma. When should they stop asking questions of the data in hand? AIC provides a remarkable answer. The model that optimizes the tradeoff between reducing variance and reducing bias also has the least estimated information loss. This conclusion is one of the most profound propositions of modern statistics. It is, however, intuitive. With regard to the magnitude of estimated effects, variance (i.e. uncertainty) and bias (i.e. systematic errors) represent the two basic forms of information loss in data analysis. Unfortunately, one cannot reduce both variance and bias. Reducing one increases the other. Thus minimizing information loss should amount to optimizing the trade-off between reducing variance and reducing bias. This conclusion is exactly what Akaike's criterion tells us. Burnham and Anderson (2002) provide an excellent survey of Akaike's criterion, its derivative criteria, and their many advantages over traditional methods of data analysis rooted in hypothesis testing.

Conclusion

The key to this prolegomenon is our conviction that, to study human behavior as evolutionary social scientists, we should generally be prepared to give behavior a nonlinear dynamical representation. Societies rise and fall, expand and contract (Diamond 1997, 2005; Turchin 2003). Technologies persist at low levels in a population only to spread suddenly and quickly (Henrich 2001). Some countries grow rich, while others remain poor or get even poorer (Barro and Sala-i-Martin 2004; Bowles 2004). Residential neighborhoods segregate racially even as residents profess a desire for integration (Bowles 2004). Success

breeds success (Frank and Cook 1995), and people do not form social networks randomly (Henrich and Gil-White 2001; Watts 2003). Without nonlinear dynamics, one must often explain such phenomena exogenously. Neoclassical economics and human behavioral ecology, for example, make extensive use of static optimization. Nonlinearities are pervasive in both theories, but without dynamics the effects of these nonlinearities do not feed back into the system. Behavioral variation thus becomes a question of how exogenous shocks, exogenous preferences, or exogenous differences in the environment affect behavior. Although exogenous forces may be important in some settings, and although static optimization may be a perfectly adequate theoretical framework in some settings, we suggest that economics and human evolutionary ecology would do well to tip the balance considerably more toward nonlinear dynamics. Indeed, we suspect that many of the truly difficult questions will respond to no other treatment.

Nonlinearity, however, will place a greater burden on empiricism and increase the difficulty of integrating theory and data. Our experience with nonlinear dynamics suggests to us that, with a few key nonlinearities in one's model, all hell can effectively break loose and leave one with a model that supports a daunting array of outcomes. The resulting theory can be illuminating, but the complexity and excess of the system require the authority of data. Unfortunately, the technical demands will only increase, as will the need for new methods in all aspects of experimental design and data analysis (Bowles 2004). And so we call this paper a prolegomenon. Virtually all the work to be done yet remains.

Acknowledgements

We would like to thank the following individuals for comments on an earlier draft of this paper: Kim Sterelny, Brian Paciotti, William Baum, Ed Edsten, Mark Lubell, Richard McElreath, and Tim Waring.

References

- Akaike H. 1973. Information theory as an extension of the maximum likelihood principle. In: Petrov B.N. and Csaki F. (eds.), *Second International Symposium on Information Theory*, Akademiai Kiado, Budapest, pp. 267–281.
- Alvard M.S. 2004. The ultimatum game, fairness, and cooperation among big game hunters. In: Henrich J., Boyd R., Bowles S., Camerer C., Fehr E. and Gintis H. (eds.), *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*, Oxford University Press, Oxford, pp. 413–435.
- Barro R.J. and Sala-i-Martin X. 2004. *Economic Growth*. 2nd ed., The MIT Press, Cambridge.
- Baum W., Richerson P.J., Efferson C.M. and Paciotti B.M. 2004. Cultural evolution in laboratory microsocieties including traditions of rule giving and rule following. *Evol. Human Behav.* 25: 305–326.

- Bolton G.E. and Ockenfels A. 2000. ERC: a theory of equity, reciprocity, and competition. *Am. Econ. Rev.* 90: 166–193.
- Bowles S. 2004. *Microeconomics: Behavior, Institutions, and Evolution*. Russell Sage, New York.
- Bowles S. and Gintis H. 2000. Walrasian economics in retrospect. *Quart. J. Econ.* 115: 1411–1439.
- Boyd R. and Richerson P.J. 1985. *Culture and the Evolutionary Process*. University of Chicago Press, Chicago.
- Boyd R. and Richerson P.J. 2005. *The Origin and Evolution of Cultures*. Oxford University Press, Oxford.
- Boyd R. and Richerson P.J. 1992. How microevolutionary processes give rise to history. In: Nitecki M.H. and Nitecki D.V. (eds.), *History and Evolution*, SUNY Press, Albany, pp. 178–209.
- Burnham K.P. and Anderson D.R. 2002. *Model Selection and Multi-Model Inference: A Practical Information-Theoretic Approach*. 2nd ed., Springer-Verlag, New York.
- Camerer C. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton.
- Camerer C. and Ho T. 1999. Experience-weighted attraction learning in normal form games. *Econometrica* 67: 827–874.
- Cavalli-Sforza L.L. and Feldman M.W. 1981. *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton UP, Princeton.
- Charness G. and Rabin M. 2002. Understanding social preferences with simple tests. *Quart. J. Econ.* 117: 817–869.
- Cushing J.M., Costantino R.F., Dennis B., Desharnais R.A. and Henson S.M. 2003. *Chaos in Ecology: Experimental Nonlinear Dynamics*. Academic Press.
- Diamond J. 1997. *Guns, Germs, and Steel: The Fates of Human Societies*. W. W. Norton, New York.
- Diamond J. 2005. *Collapse: How Societies Choose to Fail or Succeed*. Viking, New York.
- Efferson C., Richerson P.J., McElreath R., Lubell M., Edsten E., Waring T., Paciotti B. and Baum W.M. Learning, noise, productivity: an experimental study of cultural transmission on the Bolivian altiplano (submitted).
- Ensminger J. 2004. Market integration and fairness: evidence from ultimatum, dictator, and public goods experiments in East Africa. In: Henrich J., Boyd R., Bowles S., Camerer C., Fehr E. and Gintis H. (eds.), *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*, Oxford UP, Oxford, pp. 356–381.
- Frank R.H. and Cook P.J. 1995. *The Winner-Take-All Society: Why the Few at the Top Get So Much More Than the Rest of Us*. Penguin, New York.
- Fehr E. and Gächter S. 2002. Altruistic punishment in humans. *Nature* 415: 137–140.
- Fehr E. and Fischbacher U. 2003. The nature of human altruism. *Nature* 425: 785–791.
- Fehr E. and Fischbacher U. 2004. Social norms and human cooperation. *Trends Cogn. Sci.* 8: 185–190.
- Fehr E. and Schmidt K.M. 1999. A theory of fairness, competition, and cooperation. *Quart. J. Econ.* 114: 817–868.
- Forster M. and Sober E. 1994. How to tell when simpler, more unified or less ad hoc theories will provide more accurate predictions. *Br. J. Philos. Sci.* 45: 1–35.
- Gurney W.S.C. and Nisbet R.M. 1998. *Ecological Dynamics*. Oxford University Press, Oxford.
- Hasting A. 2004. Transients: the key to long-term ecological understanding. *Trends Ecol. Evol.* 19: 39–45.
- Henrich J. 2001. Cultural transmission and the diffusion of innovations: adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *Am. Anthropol.* 103: 992–1013.
- Henrich J. 2004. Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses – the Tasmanian case. *Am. Antiq.* 69: 197–214.
- Henrich J., Boyd R., Bowles S., Camerer C., Fehr E. and Gintis H. 2004a. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford UP, Oxford.

- Henrich J., Boyd R., Bowles S., Camerer C., Fehr E., Gintis H. and McElreath R. 2004b. Overview and synthesis. In: Henrich J., Boyd R., Bowles S., Camerer C., Fehr E. and Gintis H. (eds.), *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*, Oxford University Press, Oxford, pp. 8–54.
- Henrich J. and Gil-White F.J. 2001. The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evol. Human Behav.* 22: 165–196.
- Henrich J. and Smith N. 2004. Comparative experimental evidence from Machiguenga, Mapuche, Huinca, and American populations. In: Henrich J., Boyd R., Bowles S., Camerer C., Fehr E. and Gintis H. (eds.), *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*, Oxford University Press, Oxford, pp. 125–167.
- Kagel J.H. and Roth A.E. 1995. *The Handbook of Experimental Economics*. Princeton University Press, Princeton.
- McElreath R., Lubell M., Richerson P.J., Waring T.M., Baum W., Edsten E., Efferson C. and Paciotti B. 2005. Applying evolutionary models to the laboratory study of social learning. *Evol. Human Behav.* 26: 483–508.
- Murdoch W.W., Briggs C.J. and Nisbet R.M. 2003. *Consumer-Resource Dynamics*. Princeton UP, Princeton.
- Nisbet R.M. and Gurney W.S.C. 1982. *Modeling Fluctuating Populations*. Wiley, New York.
- Richerson P.J. and Boyd R. 2005. *Not by Genes Alone: How Culture Transformed Human Evolution*. University of Chicago Press, Chicago.
- Richerson P.J., Boyd R. and Bettinger R.L. 2001. Was agriculture impossible during the Pleistocene but mandatory during the Holocene?: a climate change hypothesis. *Am. Antiq.* 66: 387–411.
- Tooby J. and Cosmides L. 1992. The psychological foundations of culture. In: Barkow J., Cosmides L. and Tooby J. (eds.), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, Oxford University Press, New York.
- Tracer D.P. 2004. Market integration, reciprocity, and fairness in rural Papua New Guinea: results from a two-village ultimatum game experiment. In: Henrich J., Boyd R., Bowles S., Camerer C., Fehr E. and Gintis H. (eds.), *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*, Oxford UP, Oxford, pp. 232–259.
- Turchin P. 2003. *Historical Dynamics: Why States Rise and Fall*. Princeton UP, Princeton.
- Watts D.J. 2003. *Six Degrees: The Science of a Connected Age*. W.W. Norton, New York.