Path Dependence, Competition, and Succession in the Dynamics of Scientific Revolution

John D. Sterman • Jason Wittenberg
Sloan School of Management, E53-351, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142
Department of Political Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142

Abstract
What is the relative importance of structural versus contextual forces in the birth and death of scientific theories? We describe a formal dynamic model of the birth, evolution, and death of scientific paradigms based on Kuhn's *Structure of Scientific Revolutions.* The model represents scientific activity as a changing set of coupled institutions; a simulated ecology of interacting paradigms in which the creation of new theories is stochastic and endogenous. The model captures the sociological dynamics of paradigms as they compete against one another for members, solve puzzles, and recognize anomalies. We use sensitivity tests and regression to examine the role of intrinsic versus contextual factors in determining paradigm success. We find that situational factors attending the birth of a paradigm largely determine its probability of rising to dominance, while the intrinsic explanatory power of a paradigm is only weakly related to the likelihood of success. For those paradigms surviving the emergence phase, greater explanatory power is significantly related to longevity. However, the relationship between a paradigm's "strength" and the duration of normal science is also contingent on the competitive environment during the emergence phase. Analysis of the model shows the dynamics of competition and succession among paradigms to be conditioned by many positive feedback loops. These self-reinforcing processes amplify intrinsically unobservable microlevel perturbations in the environment—the local conditions of science, society, and self faced by the creators of a new theory—until they reach macroscopic significance. Such path dependent dynamics are the hallmark of self-organizing evolutionary systems. We consider the implications of these results for the rise and fall of new ideas in contexts outside the natural sciences such as management fads.
(*Complexity*; *Simulation*; *Competition*; *Sociology of Science*; *Scientific Revolution*)

1. Introduction
The late Thomas Kuhn's *Structure of Scientific Revolutions* heralded a radically new conception of science. In the traditional view science applies universally-accepted norms of logical inquiry, and scientific development is seen as the cumulative triumph of an ever more truthful and encompassing understanding of reality. In contrast, Kuhn (1970, pp. 84–85) argues that new theories replace old ones rather than building on them, revolutionizing science's very image of itself. For Kuhn, scientific development is fraught with errors, blind alleys, and intense competition among competing worldviews, proceeding "as a succession of tradition-bound periods punctuated by non-cumulative breaks" (Kuhn 1970, p. 208). Of course, Kuhn can be situated in a long tradition of scholars who reacted against realist and positivistic views of science, perception, and knowledge, including Paul Feyerabend (1975), and famously, Ludwik Fleck, whose 1935 theory of the social construction of scientific facts was cited by Kuhn (1970, p. vii) as "an essay that anticipates many of my own ideas." Yet no other formulation captured the imagination—or generated the ire—that *Structure* did.

Indeed, the idea that individual, social, and historic contingencies play a role in scientific development rivaling or exceeding a theory's intellectual content has elated many social scientists and historians as much as it has infuriated many philosophers and scientists. (The literature is massive. For a representative sample of the debate over the social construction of science, see e.g., Bijker et al. 1987, Collins 1985, Cushing 1994, Donovan et al. 1988, Gross and Levitt 1994, Martin 1996, Pickering 1984, 1992, Pinch 1986, and Roth and Barrett 1990). For many social scientists Kuhn's theory legitimated resistance to the century-old attempt to make the study of society, politics, and culture more like Newtonian physics. For others Kuhn's attempt to historicize scientific endeavor was reckless and heretical. Yet whether as prophecy or apostasy, his ideas continue to stimulate vigorous debate about the evolution of science (e.g., Hoyningen-Huene 1993, Lightman and Gingerich 1992).

The controversy over the nature of scientific progress

1047-7039/99/1003/0322/$05.00

parallels the debate in the social sciences over whether organizational change is gradual and evolutionary or episodic and revolutionary. Though Kuhn (1970, pp. 208–210) cautioned against the applicability of his model to the social sciences, it is nevertheless widely cited by social scientists as descriptive of organizational behavior and cognitive shifts in contexts far beyond the natural sciences. Organization theorists argue that the pattern of punctuated equilibrium Kuhn finds in the history of science also characterizes many instances of organizational change (see Gersick 1991, Tushman and Anderson 1986, Sastry 1997). Tushman and Romanelli (1985, p. 171) propose a model of organizational change in which "[o]rganizations evolve through convergent periods punctuated by reorientations . . . which demark and set bearings for the next convergent period." Gersick (1991) shows there are many domains and levels of analysis, from paleontology (Gould 1990, Eldredge and Gould 1972) to group dynamics, in which change can be characterized by long periods of stasis or gradualism punctuated by sudden upheavals and revolutions. Some argue that change is often a more continuous and adaptive process (e.g., Orlikowski 1996), while still others argue that organizational adaptation is rare, with selection and evolution occurring at the population level (Hannan and Freeman 1989; Van de Ven and Poole 1995 provide a survey of change research).

The similarity between the dynamics of science and of organizations is not mere coincidence. Scientific activity is not primarily the work of solitary geniuses, but is embedded in a wide range of organizations, from the small group level of researchers and graduate students in a lab, to departments and universities, to the network of funding agencies, journal boards, and professional societies that constitute the "invisible colleges" defining a community of practice (Crane 1972). There are of course differences between the institutions of science and organizations in other domains such as business firms. Yet the institutions of science are among the most influential in our society, deserving of study in their own right. Additionally, the dynamics of scientific revolution may shed light on organizational evolution in general. Why is it that some scientific paradigms last for centuries while others quickly wither? How do intellectual, structural, and contextual forces interact to shape and constrain the development of new paradigms (and organizations)? What determines whether a new theory (or organization) survives its founding and becomes dominant? Do structural or contingent forces dominate the dynamics of social systems?

We address these questions with a formal dynamic model of paradigm emergence and competition. The model creates a simulated ecology of interacting paradigms in which the genesis of new paradigms is stochastic and endogenous. The model captures the sociological dynamics of paradigms as their members formulate and solve "puzzles," recognize and react to anomalies, and alter their beliefs and behavior. Competition for membership and resources is explicit. The model is used to investigate the relative importance of structural versus contextual factors in determining the fate of new ideas.

Although the model is based on Kuhn's work, we do not claim to capture his theory fully. Translating the theory from its highly abstract written form into an internally consistent formal model has involved simplifications and the introduction of auxiliary hypotheses (for a discussion and critique see Wittenberg 1992, Sterman 1992, Radzicki 1992, Barlas 1992). Nonetheless, formalization has advantages. Most discussions of Kuhn's theory are based on ambiguous mental models, and Kuhn's text itself is rich with ambiguity, multiple meanings, and implicit assumptions (Masterman 1970). More importantly, Kuhn offers no calculus by which one can assess whether the dynamics he describes can be produced by the causal factors he postulates. Formalization helps to surface implicit assumptions so they can be debated and tested (see Gorman 1992, Rappa and Debackere 1993, Sastry 1997, Stewart 1986, Thagard 1968, and Turner 1987 for examples). Formalization is complementary to the work of historians, sociologists, and philosophers of science working to develop and test theories of scientific change and institutional upheaval.

Finally, the model applies nonlinear dynamics to social phenomena and human behavior. Modern theories of nonlinear, far from equilibrium systems, though emerging in the physical sciences, have great potential to illuminate evolutionary behavior in social, economic, and other human systems (e.g., Anderson et al. 1988; Arthur 1989, 1994; Bruckner et al. 1989, 1990; Ebeling 1991; Lomi and Larsen 1996). The full potential of these tools in the social sciences will be realized, we believe, only when they are used to develop and test formal models. The merely metaphorical use of concepts from nonlinear dynamics, while provocative, is not sufficient, a point also stressed by Carley and Wallace (1995) and Richardson (1996, 1991). In addition, useful models will draw on experimental and field studies of human behavior to specify the decision rules governing the behavior of the simulated agents (see Carley 1991, 1995; Cyert and March 1963/1992; Forrester 1961; Hall 1976; Haxholdt et al. 1995; Lant and Mezias 1992; Morecroft 1985; Nelson and Winter 1982; Radzicki and Sterman 1994; and Sterman 1988, 1989 for discussion and examples). Here we develop a stochastic, nonlinear, disequilibrium, behavioral model of the evolution of scientific theories, grounded in Kuhn's theory. As will be seen, the dynamics exhibit self-organization and path dependence, two common modes of behavior in complex systems.

2. A Theory of Paradigm Evolution and Succession

We assume familiarity with Kuhn's work and the many interpretations and alternatives to it (see Lakatos and Musgrave 1970 for classic critiques and Hoyningen-Huene 1993 for a comprehensive survey and bibliography). A core concept in Kuhn's theory is the life cycle of a paradigm. Kuhn describes a sequence of four stages: emergence, normal science, crisis, and revolution (followed by the emergence of a new paradigm). The emergence phase is characterized by the absence of commonly accepted beliefs or standards governing scientific activity. Conflict among paradigm-candidates arises from incompatible metaphysical beliefs and logics of inquiry, as Kuhn (1970, pp. 13–15) illustrates with the state of electrical research before Franklin and his colleagues provided the field with a paradigm. As a theory attracts nearly every scientist in the field—thereby becoming the dominant paradigm—normal science begins. Debate over fundamental assumptions dwindles, and, convinced their paradigm is the proper way to characterize reality, scientists proceed to apply it to nature's puzzles. During normal science, clashes between theory and data are often resolved in favor of theory—it is often presumed that any anomalous observations are wrong, or the calculations erroneous, so that further puzzle-solving effort will resolve the anomaly, a behavior Kuhn (1970, p. 81) illustrates by citing anomalies facing Newtonian mechanics involving the speed of sound, the moon's perigee, and the precession of the orbit of Mercury. This is observed today in the debate over the value of the Hubble constant and the age of the universe (Chaisson 1997).

Normal science continues until a crisis arises. A paradigm can enter crisis when enough unsolved puzzles become recognized by practitioners as important anomalies, persuading them that the theory must, after all, be questioned. As persistent anomalies accumulate, increasing numbers of scientists will devote their time to solving them rather than addressing new puzzles, and some may propose radical solutions. A revolution occurs when a new paradigm based on such a radical reconceptualization gains wide acceptance, and science is reconstructed from new fundamentals. Obviously the timing, character, and context of each stage differ from case to case. A dominant paradigm in crisis may quickly be replaced, or crisis may deepen for decades as new theories fail to sprout or flower. Kuhn provides little guidance into the forces that cause one new idea to triumph and another to fail, or the determinants of the longevity of those new paradigms that survive their founding and become dominant. The central debate has been the relative importance of intrinsic explanatory power—the "truth" of a new theory—versus contingencies external to science such as the social, political, and cultural context of emergence, or even chance factors—the existence of an Einstein, Bohr, or Keynes—in conditioning which paradigm candidates flourish and which perish.

3. The Model

We construct a multiparadigm model in which existing theories compete for membership and resources and in which the creation of new theories is stochastic and endogenous. The model can be thought of as a set of interacting agents (the communities loyal to a particular theory or paradigm). Like other agent-based models (e.g., Holland 1995, Weisbuch 1991), the collective system dynamics emerge from the interaction of the individual agents over time. Unlike some agent-based models, the number of agents is not fixed; new theories are created with a probability that varies endogenously as conditions change. Also unlike some agent-based models, the individual paradigms have a rich internal structure representing the activities of each community, including the belief structure of the members, recruitment and defection, scientific activity such as puzzle solving and anomaly recognition, and the flows of people and information that couple the different paradigms competing against one another at any given time.[1] In what follows we provide an overview of the model; complete documentation is available from the authors.

The heart of the model is the identification of the metaphysical and epistemological facets of paradigms with metaphors, limited representations of reality that generate anomaly when pushed too far. Four properties of metaphor that are also properties of paradigms bear particular mention. First, metaphor is everywhere. Goodman (1968, p. 80) argues that "metaphor permeates all discourse, ordinary and special, and we should have a hard time finding a purely literal paragraph anywhere." Turbayne (1970) goes further, suggesting metaphor permeates our thought as well as our language. Similarly, Kuhn (1970, p. 113) stresses the priority of paradigms, suspecting that "something like a paradigm is prerequisite to perception itself." Second, metaphor involves a "transfer of schema" from one area of experience to another (Goodman 1968, pp. 71–80). Consider the metaphor "the brain is a computer." The characteristics of computers are transferred, via the metaphor, to our image of the brain. The metaphor works because the characteristics of computers are well known and carry a constellation of meanings and examples that illuminate certain characteristics of the brain. For Kuhn paradigms operate similarly: scientists are taught to transfer familiar models to new puzzles, to "grasp the analogy" (p. 189). Third, metaphors filter reality. Because metaphors are inevitably inexact, as are all models, they highlight certain relationships and obscure others. Metaphors focus our attention on particular facts and relations while others are pushed into the background. The filtering power of paradigms is central to Kuhn's theory as well: "In the absence of a paradigm . . . all the facts that could possibly pertain to the development of a given science are likely to seem equally relevant" (p. 15). Finally, metaphors define reality. Max Black (1954–1955, pp. 284–285) notes that "[i]t would be more illuminating in some of these cases to say that the metaphor creates the similarity than to say that it formulates some similarity already existing." Kuhn (1970, p. 111) attributes the same power to paradigms:

> The historian of science may be tempted to exclaim that when paradigms change, the world itself changes with them. Led by a new paradigm, scientists adopt new instruments and look in new places. . . . [They] see new and different things when looking with familiar instruments in places they have looked before. Insofar as their only recourse to the world is through what they see and do, we may want to say that after a revolution scientists are responding to a different world.

Yet metaphors are imperfect models, and if pushed too hard crack and fail. Consider the brain-as-computer metaphor. Applying this metaphor might yield statements that generate insight, motivate theory, or suggest experiments, such as "people transfer information from long-term to working memory with characteristic seek times." Eventually, however, overextension of the metaphor yields absurdities such as "brains run Microsoft Excel" or "brains are composed of silicon semiconductors." The accumulation of such anomalous claims undermines the appeal of a metaphor, and can send it to its grave, disgraced as falsehood. Kuhn views the life cycle of paradigms in a similar way. The elaboration and extension of a paradigm to new domains can lead to the accumulation of anomalies. As an "almost entirely typical" example Kuhn cites the accumulation of anomalies in Newtonian mechanics, such as the repeated failure to detect drift through the ether resulting from the effort to provide a Newtonian foundation for Maxwell's theory of electromagnetic radiation. As a result, "Maxwell's theory, despite its Newtonian origin, ultimately produced a crisis for the paradigm from which it had sprung" (p. 74).

Thus the central dynamic hypothesis of the model draws on the notion that paradigms are extended metaphors, and that metaphors are not unlimited in their applicability to reality. Specifically, we assume that the average difficulty of the puzzles faced by the practitioners of a paradigm increases as the cumulative number of puzzles they have solved grows. This "paradigm depletion" represents the idea that each paradigm is a limited model of reality that may apply well in the domain of phenomena it was originally formulated to explain, but will be harder and harder to apply as scientists extend it to new domains. The formalization of this hypothesis is described below.

The model creates a simulated ecology of interacting paradigms, each representing a community of practitioners; recruitment and defection from that community; and the intellectual activities of the members such as formulating and solving puzzles, recognizing and trying to reconcile anomalies, and conceiving new theories. The model simulates the attitudes and beliefs of the practitioners within each paradigm through constructs such as "confidence in the paradigm" and the time required to perceive unexplained phenomena as anomalies which challenge the theory. The major sectors of the model and the linkages among paradigms are shown in Figure 1; we will use causal diagrams to illustrate the feedback processes and stock-and-flow structure of the model (Richardson and Pugh 1981, Sterman 1985, Weick 1979). Each paradigm has the same internal structure; for clarity we display only paradigms *i* and *j*.

3.1. Confidence in the Paradigm

The focal point of the model is a construct called "confidence." Confidence captures the basic beliefs of practitioners regarding the epistemological status of their paradigm—is it seen as a provisional model or revealed truth? Encompassing logical, cultural, and emotional factors, confidence influences how anomalies are perceived, how practitioners allocate research effort to different activities (puzzle solving versus anomaly resolution, for example), and recruitment to and defection from the paradigm. It is defined from 0 (absolute conviction the paradigm is false, nonsensical) through 0.5 (maximum uncertainty as to its truth) to 1 (absolute conviction the paradigm is truth). Pressures leading confidence to change arise both from within a paradigm and from comparisons with other paradigms (Figure 2). Confidence rises when puzzle-solving progress is high and when anomalies are low. The impact of anomalies and progress is mediated by the level of confidence itself. Extreme levels of confidence hinder rapid changes in confidence because practitioners, utterly certain of the truth, dismiss

Figure 1 Overview of Model Structure

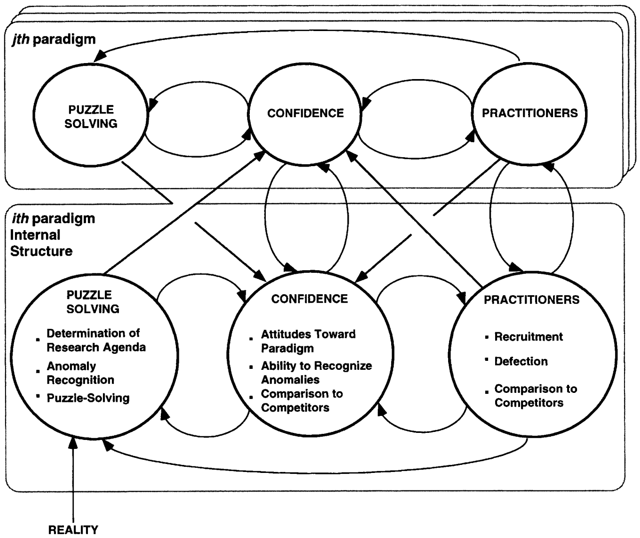

Figure 2 Internal and External Determinants of Confidence in a Paradigm

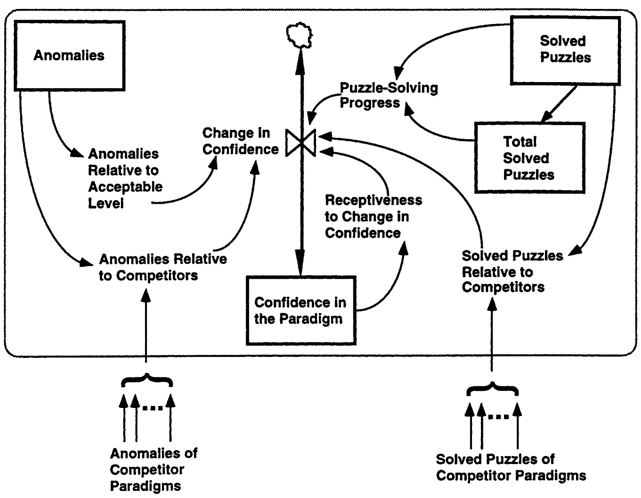

any evidence contrary to their beliefs. Practitioners with only lukewarm commitment, lacking firm reasons to accept or reject the paradigm, are far more likely to alter their beliefs in the face of anomalies.

The external factors affecting confidence encompass the way in which practitioners in one paradigm view the accomplishments and claims of other paradigms against which they may be competing. We distinguish between the dominant paradigm, defined as the school of thought that has set the norms of inquiry and commands the allegiance of the most practitioners, and alternative paradigms, the upstart contenders. The confidence of practitioners in a new paradigm tends to increase if its anomalies are less than those of the dominant paradigm, or if it has greater explanatory power, as measured by cumulative solved puzzles. Confidence tends to decrease if the dominant paradigm has fewer anomalies or more solved puzzles. Practitioners in alternative paradigms assess their paradigms against one another as well as against the dominant paradigm. Confidence in an alternative paradigm tends to decrease (increase) if it has more (fewer) anomalies or fewer (more) solved puzzles than the most successful of its competitors.

3.2. Puzzle Solving

The determinants of puzzle solving are shown in Figure 3. Three categories of puzzles are distinguished. Solved puzzles are puzzles that have already been integrated into the corpus of theory and data constituting the paradigm. Anomalies are unsolved puzzles which have come to be recognized as serious challenges to the theory. The third category, puzzles under attack, consists of those puzzles that are actively under study, but which have neither been solved nor yet recognized as anomalies. Four flows connect the different categories. Most puzzles, once formulated and attacked, will be solved, adding to the cumulative stockpile of knowledge generated by the paradigm. Such puzzles flow into the class of solved puzzles via the puzzle-solving rate. But as the intrinsic difficulty of puzzles grows, a growing number will resist solution long enough to be recognized as anomalies. Anomalies may sometimes be resolved, adding to the stock of solved puzzles via the anomaly resolution rate. The shifting balance between these flows determines the behavior of the system.

Figure 3 The Puzzle-Solving Sector

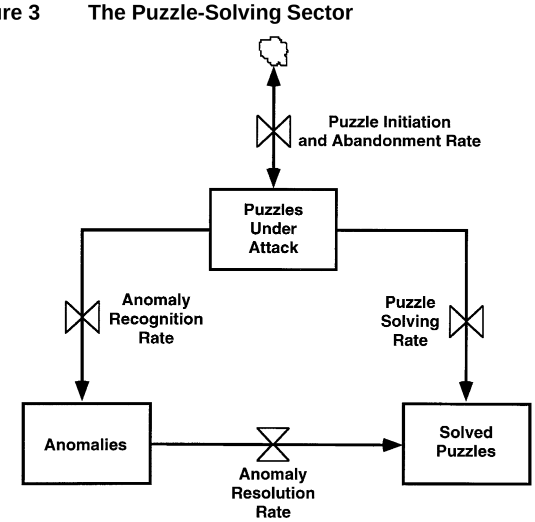

3.3. Puzzle Formulation and Puzzle-Solving Rates

The rate at which scientists formulate and solve puzzles depends on the number of puzzles under study, the fraction of practitioners involved in puzzle solving, the fraction of their time devoted to puzzle solving, and the average difficulty of the puzzles (Figure 4).

The average difficulty of new puzzles depends on how far the root metaphor defining the paradigm has been extended. As described above, the average difficulty of puzzles is assumed to rise as the paradigm is applied to phenomena increasingly removed from the original domain for which the paradigm was formulated. Specifically, the average difficulty of new puzzles to be solved, D, rises as the number of puzzles the paradigm has solved grows. We assume

$$D = (S/C)^{\gamma}, \tag{1}$$

where S is the cumulative number of solved puzzles. The nominal solved puzzle reference, C, represents the intrinsic capability of each paradigm, and γ is the rate at which difficulty rises with cumulative progress. When $\gamma < 1$, the rate at which puzzle difficulty rises with cumulative progress becomes progressively smaller, while $\gamma > 1$ indicates the difficulty of puzzles on the margin rises ever faster. For parsimony we assume $\gamma = 1$. Small values of the reference capability C mean a paradigm's intrinsic explanatory power is low—the difficulty of new puzzles rises rapidly as normal science proceeds. Large values indicate a more powerful paradigm, one that could encompass a wider array of phenomena. Note that our formulation differs from that of Masterman (1970), who viewed paradigms as analogous to nonrenewable resources, arguing that the domain of applicability for any paradigm is finite, so all attempts to extend it further would yield only anomaly. Her assumption would mean puzzle-solving difficulty in the model would become infinite when the stock of cumulative solved puzzles reached some finite value, just as no amount of effort can bring any diamonds out of a mine once it is played out. We make the less restrictive assumption that the puzzle-solving potential of paradigms is infinite, though it rises continuously on the margin as solved puzzles accumulate.

Figure 4 **Determinants of the Puzzle-Solving Rate**

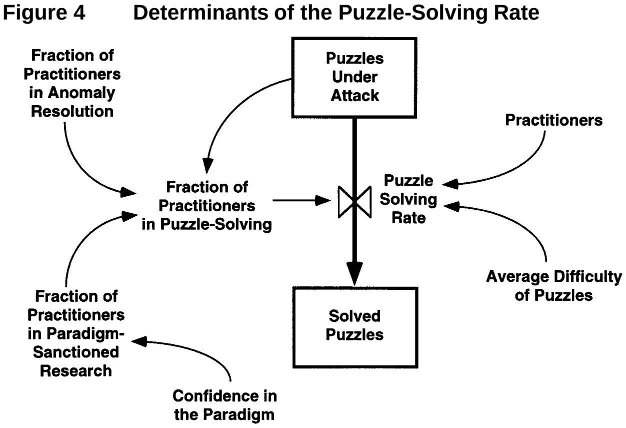

As the difficulty of puzzles grows, puzzle solving may slow and more unsolved puzzles may become recognized as anomalies. If the stock of anomalies grows too large, the confidence practitioners have in the "truth" or utility of the paradigm may fall. The collapse of confidence is self-reinforcing: anomalies erode confidence, and falling confidence increases the ability and willingness of practitioners to perceive the gaps in the theory.

The majority of practitioners will usually be involved in puzzle solving, while some will be working to resolve anomalies and others try to generate alternatives or engage in other activities such as administration or popularization. The distribution of practitioner effort among these three categories is a function of confidence in the paradigm. The higher the confidence, the greater the fraction of practitioners involved in normal science. As confidence falls, more practitioners turn their attention to anomaly resolution or altogether away from the normal science they increasingly come to doubt.

3.4. Anomaly Recognition Rate

Anomaly recognition is a subtle psychological process (Lightman and Gingerich 1992). Kuhn notes that anomalies are not simply experiments that run counter to expectation, as there are always disagreements between data and theory. Rather, a puzzle becomes recognized as an anomaly when normal science repeatedly fails to resolve the differences. Kuhn (1970, p. 82) argues that "One source of the crisis that confronted Copernicus was the mere length of time during which astronomers had wrestled unsuccessfully with the residual discrepancies in Ptolemy's system." Similarly, we assume that the longer an unsolved puzzle has resisted solution, the greater the chance it will be recognized as an anomaly. Thus, the probability a puzzle is recognized as an anomaly rises as the average difficulty of puzzles rises. However, recognition of anomalies also depends on the degree to which practitioners are conditioned to see reality as consistent with their paradigm. Kuhn cites the Bruner-Postman playing card experiments to illustrate how a paradigm conditions perception, concluding "In science, as in the playing card experiment, novelty emerges only with difficulty, manifested by resistance, against a background provided by expectation" (1970, p. 62ff). Thus in the model, the average time required to recognize an unsolved puzzle as an anomaly depends on practitioners' level of confidence in the paradigm. High levels of confidence slow the recognition of anomalies as practitioners' expectations, behaviors, and even perceptions become increasingly conditioned to be consistent with the paradigm. Decreases in confidence will cause more of the puzzles under attack to be considered anomalous as practitioners' skepticism and doubts grow.

3.5. Anomaly Resolution Rate

The rate at which anomalies are resolved depends on the number of practitioners in sanctioned research, the fraction of those involved in anomaly resolution, and the average difficulty of anomalies (Figure 5). Anomalies are assumed to be more difficult to solve than puzzles, and as the difficulty of puzzles increases, the difficulty of anomalies rises as well. The fraction of practitioners involved in anomaly resolution depends on the balance between the number of anomalies and the acceptable number. The acceptable number of anomalies is the number that can be tolerated without losing confidence in the paradigm. If the number of anomalies increases, additional practitioners are drawn into anomaly resolution in an attempt to solve the major outstanding problems challenging the theory. This negative feedback is comparatively weak, however: Kuhn argues that most practitioners are reluctant to work on anomalies, preferring instead the relative safety and professional rewards of puzzle-solving. The belief that anomaly hunting may be hazardous to your career is widespread among scientists today and often reinforced in the professional journals. Examples abound: a 1996 news article in *Science* reports Nobel laureate Martin Perl's efforts to detect free quarks, a phenomenon counter to the predictions of quantum chromodynamics, the long-successful theory of the strong force pioneered by Murray Gell-Mann and George Zweig in the 1960s. Though Perl asserts "a positive finding would overturn 30 years of our thinking about strong interactions," he "as a tenured Nobel laureate, has the 'luxury' of continuing the search." Others caution that "a younger scientist trying to make a reputation would be well-advised to avoid this line of work." (Nadis 1996, pp. 1361–1362).

3.6. Practitioner Population

The population of practitioners committed to each paradigm is endogenous, increasing with recruitment and decreasing with retirement of elder scientists and defection of others to competing paradigms. Without loss of generality we assume the total population of scientists is constant: scientists who leave one paradigm enter another; and entry of young scientists is balanced by retirement of the old. The assumption of constant total population simplifies the interpretation of the results but is in no way essential to the main conclusions. Practitioners defect based on their confidence relative to the confidence of those in the dominant paradigm (Figure 6). The greater the (negative) discrepancy between a challenger's confidence and confidence in the dominant paradigm, the larger the proportion of the challenger's practitioners that will defect. Recruitment is proportional to a paradigm's relative attractiveness and its total number of practitioners. The greater a paradigm's attractiveness, the greater the proportion of defectors from other paradigms it will recruit. Attractiveness is proportional to the number of practitioners since large paradigms are assumed to get more funding, train more students, and have a larger voice

Figure 5 Determinants of the Anomaly Resolution Rate

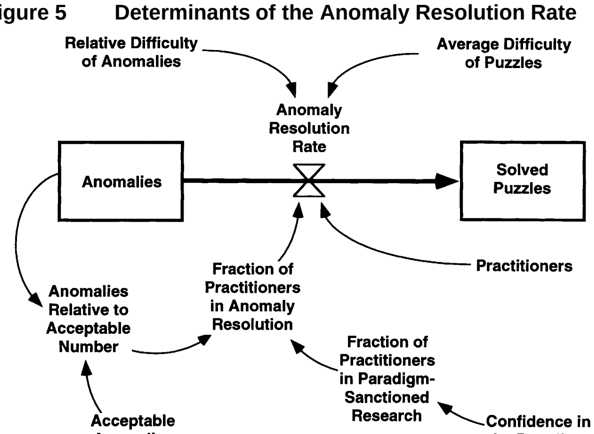

Figure 6 Internal and External Determinants of Practitioner Recruitment and Defection

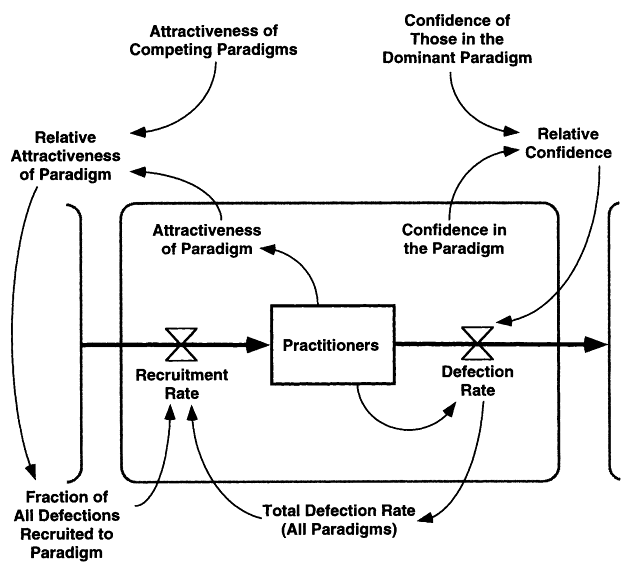

in tenure and other peer-career decisions than small paradigms. Attractiveness also depends on the confidence of the paradigm's practitioners, capturing the competence of the members, the capability of their tools, and the excitement and enthusiasm flowing from a successful endeavor.

3.7. The Creation of New Paradigms

We model the creation of a new paradigm as a stochastic event whose probability depends upon the distribution of practitioner activities in the currently dominant paradigm. Practitioners may toil in normal science (puzzle solving), anomaly resolution (the attempt to reconcile anomalies with the current paradigm), and other activities (described by Kuhn as including philosophical reconsideration of the paradigm and other activities not sanctioned by the dominant paradigm). In general, each of these activities may result in the creation of a new paradigm, but the probability that a new paradigm is created as a result of a practitioner year of effort devoted to each activity may differ. Thus:

$$P(\text{Creation})_t = \sum_{i} \pi_i^* N_{i,t},\ i \in \{\text{PS, AR, OA}\}, \quad (2)$$

where

$P(\text{Creation})_t$	=	probability a new paradigm is created at time t;
$N_{i,t}$	=	number of practitioners in the dominant paradigm engaged in activity i at time t;
π_i	=	probability of creating a new paradigm per practitioner year of effort in activity i;
{PS, AR, OA}	=	Activities: Puzzle Solving, Anomaly Resolution, Other Activities, respectively.

Following Kuhn, we assume $\pi_{AR} > \pi_{OA} > \pi_{PS}$: Normal science is unlikely to produce new paradigms, focused as it is on solving puzzles within the context of the existing paradigm. Other activities are more likely to produce a new paradigm, while effort devoted to anomaly resolution is most likely to result in the creation of radical new theories (the values of these parameters are small enough that the overall probability of creating a new paradigm in any given year is low). In the model, the distribution of effort among these three activities is endogenous. Thus the probability that a new paradigm will be created in any time period is endogenous and will vary as practitioner effort changes in response to the changing health of the dominant paradigm.

Once a new paradigm is created, we assume it begins with a small number of practitioners, a confidence level of 0.5 (neutral), a very small stock of solved puzzles, and no initial anomalies. The newly launched paradigm must then compete for members against the dominant paradigm. During a period of crisis the probability of creating a new paradigm may rise and remain high long enough for more than one new paradigm to emerge. In this case the newly created paradigms will vie for ascendancy not only against the dominant paradigm but against one another.

4. Exploring the Dynamics of Paradigm Development

We begin by simulating the model with fully endogenous competition among paradigms. The initially dominant theory (Paradigm 1 [P1]) is initialized in the midst of normal science, and new theories are created stochastically, with a probability depending upon the vitality of the dominant paradigm as specified by Equation (2). The intrinsic capability of each new paradigm is determined by a host of factors including the richness of the theoretical constructs emerging from the paradigm's root metaphor and of course the particular genius of the paradigm's creators. Thus, the rate at which puzzle solving becomes difficult as solved puzzles accumulate (the paradigm's inherent potential, C) is stochastic. Specifically, C is drawn from a lognormal distribution (truncated such that $C \leq 800$). Otherwise all paradigms have identical structure and parameters.

Figures 7a and 7b show the first 1400 years of a representative simulation. New paradigms are created stochastically, but the probability of creation is endogenous,

Figure 7 **A Typical Simulation Showing Competition and Succession among Paradigms: Random Potential Explanatory Power**

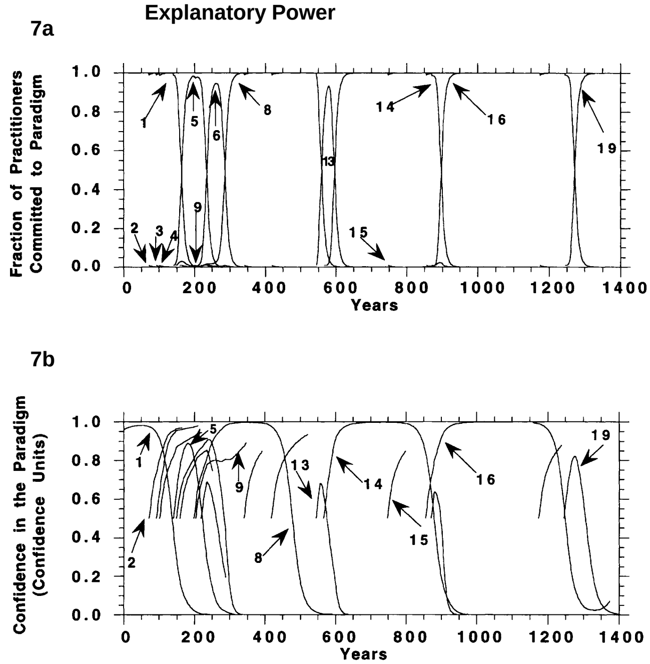

as specified in Equation (2). Each new paradigm is endowed with a randomly-selected intrinsic explanatory power [the parameter C in Equation (1)]. The simulation yields a succession of dominant paradigms in which the initial paradigm gives way to challengers, each of which goes through the typical life cycle as described by Kuhn, though with variations in length and timing. Because all paradigms have identical structure and parameters, all differences in outcomes are due only to two factors: (1) the intrinsic capability with which each is endowed; and (2) the competitive environment (number and state of other paradigms) at the time of their founding.

What is most interesting is not what the figures display but what they conceal. Most new theories face early extinction. As evident in Figure 7a, paradigms 2–4, 7, 9–12, 15, and 17–18 never become dominant, illustrating what Kuhn (1970, pp. 136–143) calls the invisibility of revolutions, where the linear and cumulative character of normal science portrayed in the textbooks conceals the contentious character of actual scientific practice. The simulation replicates the "punctuated equilibrium" pattern described by Kuhn and observed in many other fields, including organizational theory (Gersick 1991, Tushman and Anderson 1986).

The endogenous forces underlying a paradigm's evolution are best illustrated by focusing on the life cycle of a particular paradigm. Figure 8 enlarges that portion of Figure 7a portraying the life cycle of P14. Around year 500, paradigm 8 is in the full flower of normal science, with 100% of the practitioners, a high level of confidence, and few anomalies. Paradigm candidates 9–12 are, by chance, created during the period of normal science and quickly perish. However, the continued success of the dominant theory P8 leads practitioners to apply it to more and more phenomena. Anomalies slowly accumulate as puzzles gradually become more difficult to solve, eventually leading to crisis and a drop in confidence. Paradigms 13 and 14 both arise during the crisis of paradigm 8 (around years 545 and 566, respectively). By chance, P13 has very low inherent potential. Its rapid rise around year 580 is matched by an equally rapid drop as its practitioners quickly exhaust the limited potential of its underlying metaphor, making way for paradigm 14. Figure 9 illustrates the details of P14's life cycle. In the early period ($\approx$ years 560 to 610), confidence rises dramatically, since puzzle-solving progress is rapid and anomalies are low. The paradigm, initially untested, proves capable of solving puzzles, and thus attracts more practitioners, further boosting confidence.

Figure 8 The Rise and Fall of Paradigm 14 (From Figure 7)

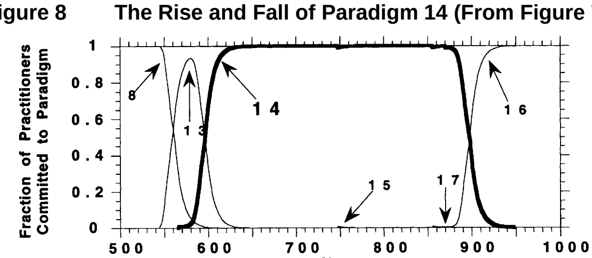

The simulation illustrates how multiple positive feedback processes cause the self-reinforcing rise of a new theory. Figure 10 shows a causal diagram highlighting two of the positive feedback loops that cause an initially unorganized and weakly committed group of practitioners to coalesce into a highly focused paradigm (for clarity negative loops are not shown). In causal diagrams, arrows indicate the direction of causality. Signs ("$+$" or "$-$") at arrow heads indicate the polarity of relationships: a "$+$" indicates that an increase in the independent variable causes the dependent variable to increase above what it would have been, ceteris paribus (and a decrease causes a decrease). A "$-$" indicates that an increase in the independent variable causes the dependent variable to decrease below what it would have been. That is, $X \rightarrow {}^{+}Y \Rightarrow (\partial Y/\partial X) > 0$ and $X \rightarrow {}^{-}Y \Rightarrow (\partial Y/\partial X) < 0$. Positive loop polarity, denoted by $(+)$ in the loop identifier, indicates a self-reinforcing (positive feedback) process. Negative $(-)$ loop polarity indicates a self-regulating (negative feedback) process (Richardson and Pugh 1981). Rising confidence and successful puzzle-solving boost practitioner confidence, leading to more focused and successful effort, articulation and improvement of theory and technique, and still greater success in puzzle solving, further boosting confidence and attracting still more members. Rising confidence, skill, and familiarity with the paradigm increasingly condition practitioner perceptions and expectations, suppressing the recognition of anomalies; a low level of anomalies further increases practitioners' confidence in and commitment to the theory. These and other positive feedbacks (shown in Figure 10) bootstrap paradigm 14 into dominance by around year 625, its metaphor, method and metaphysics triumphant over the now-discredited P13.

4.1 Normal Science

During the successful period of normal science (approximately years 620 to 830) practitioners focus their efforts on puzzle solving and are blinded to potential anomalies by their faith in the paradigm. The probability a new paradigm is created falls [see Equation (2)]. In this fashion, success suppresses the generation of new competitors which might challenge the dominant paradigm, leading

Figure 9 **The Life Cycle of Paradigm 14**

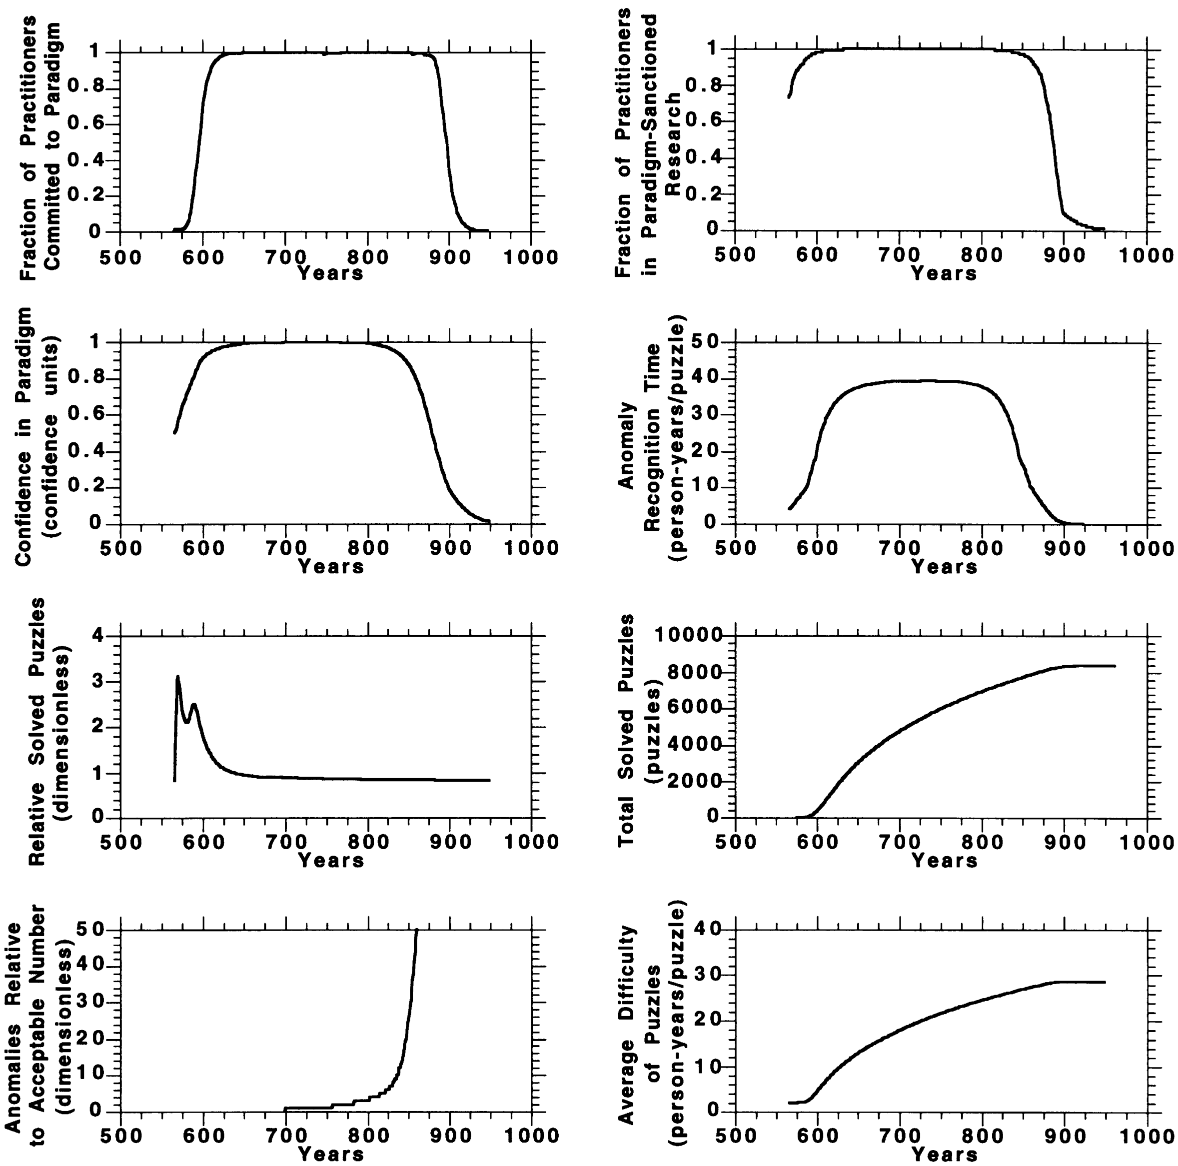

to further success. Through this self-reinforcing feedback a successful theory alters its own environment in ways that provide further advantage. This important dynamic operates through the training of graduate students, which reproduces the worldview and prejudices of the dominant theory and socializes them in the accepted canon of prior work, through the control of institutions via appointments and tenure, through resource allocation via peer review of grant proposals, and through access to journals via control of editorial boards and the selection of referees. However, occasionally a new theory does emerge during periods of normal science, such as Paradigm 15 just before year 750 (Figure 8). Such challengers usually perish in the face of competition with the still successful dominant paradigm. Indeed, P15 vanishes within a few years.

4.2 Crisis

As Paradigm 14 is elaborated and extended beyond the scope of its root metaphor, puzzles gradually become more difficult to solve. Anomalies begin to accumulate. Confidence begins to fall, slowly, around year 780. As anomalies increase, a few practitioners leave puzzle-solving, eroding progress and decreasing confidence further. Practitioners, increasingly sensitive to the paradigm's limitations, become more apt to see difficult

Figure 10 Some Positive Loops Driving Path-Dependent Behavior

Note: Shows two of the positive loops that cause initially uncommitted and unorganized practitioners to coalesce into a highly focused paradigm. (Negative loops are not shown.)

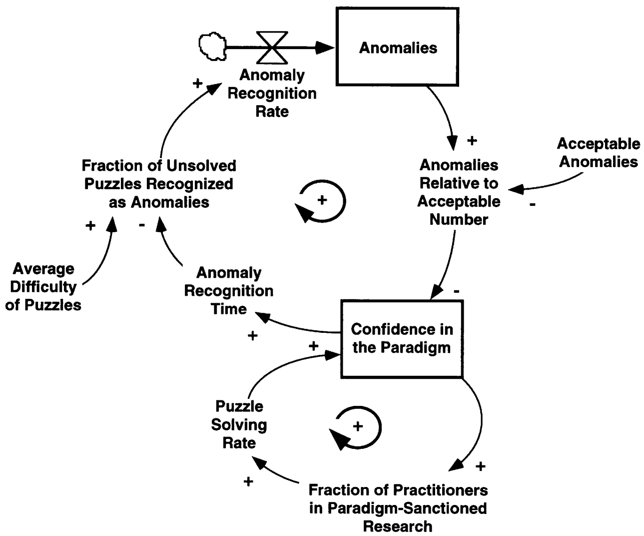

puzzles as anomalies, further increasing anomalies and decreasing confidence. The positive feedbacks that previously caused membership to rise now cause accelerating collapse. By year 850 the paradigm is in crisis.

As the number of practitioners engaged in normal science falls, and those seeking to resolve anomalies grows, the probability that a new paradigm will be created rises. Around year 855 a new paradigm is in fact created (P16 in Figure 8). Because the new theory emerges during the crisis of Paradigm 14, it quickly gains adherents while P14 loses members. Confidence and membership in P16 then accelerate sharply through the same positive feedbacks which earlier led to the success of P14. The cycle is completed as Paradigm 14's confidence and membership eventually fall to 0, while P16 grows to dominate the field. What was once uncontested "truth" is now seen as primitive error. Paradigm 17, created around year 870, is quickly crushed by the now dominant P16.

4.3 Positive Feedback and Path Dependence

The many positive feedbacks described above create the self-organizing dynamic by which uncommitted and unorganized practitioners coalesce into a highly focused paradigm with a productive program of normal science. Through these feedbacks a successful paradigm alters its environment by suppressing the creation of competitors and rapidly starving any that do emerge of the resources they would need to succeed. The same feedback processes operate in the opposite direction during the crisis period to accelerate the collapse of a paradigm which has accumulated sufficient anomalies for confidence to begin falling.

The simulations raise a number of important questions. Why do some paradigms rise to dominance while others quickly wither? Does the fate of a new paradigm depend on its intrinsic capability to explain nature or on situational contingencies surrounding its birth? Does 'truth' eventually triumph as better theories defeat inferior ones, or is timing everything?

There is evidence for both positions in the results. Supporting the view that intrinsic explanatory power is critical are examples such as Paradigm candidate 13, which rapidly exhausts its low intrinsic potential and never achieves dominance. However, intrinsic capability does not explain the fate of many others. Consider Paradigms 8 and 9 in Figures 7a and 7b, launched around years 199 and 203, respectively. Although they emerge only about four years apart, during the crisis of Paradigm 5, P8 comes to dominate the field, while P9 eventually perishes. Here the contingency of outcomes on situational factors is decisive. Paradigm 8 does not succeed because of a head start in attracting practitioners: between years 212 and 215 it actually has the same number as P9. Nor is Paradigm 8's success a result of superior explanatory power: by chance, P9 is endowed with a potential 13% *greater* than P8. The difference in their destinies lies in their levels of confidence. In the year 212 Paradigm 8, though equal in size to P9, is slightly more attractive because its adherents, having had a 4-year lead over P9 in solving puzzles, have been able to articulate their paradigm more coherently and persuasively than their chief rivals. The small advantage held by P8 is amplified as success begets success through the many positive loops surrounding the emergence process (Figure 10). Paradigm 8 eventually dominates science, while Paradigm 9 slowly fades into obscurity, to be remembered, if at all, as a blind alley, foolish error, or curiosity.

The simulations illustrate the subtle interplay between endogenous feedback processes and contextual, situational factors in determining the dynamics and succession of paradigms. The basic life cycle of paradigms is determined by the recursive, reflexive feedback loop structure discussed above. Figure 11 shows some of the positive feedback loops that act to differentiate competing paradigms even when they are initially quite similar (the many negative feedbacks are not shown). These positive feedbacks boost confidence and rapidly generate a focused community from a promising but unexplored new idea. They give a paradigm with an initial advantage an edge in recruitment of new members, leading to still greater

Figure 11 Some Positive Feedback Loops that Create Path-Dependent Behavior

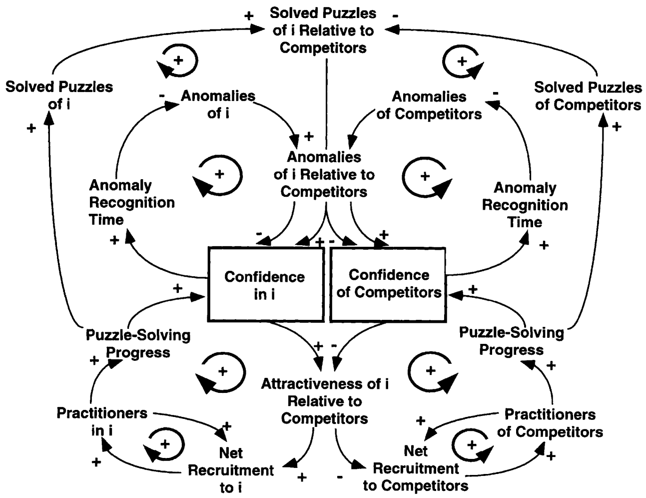

advantage, amplifying small fluctuations in local conditions to macroscopic significance, and leading to path dependence. Consider Paradigm *i* in Figure 11. If the number of anomalies and solved puzzles in Paradigm *i* compare favorably with the accomplishments of competitor paradigms, the confidence of practitioners in *i* will rise and the confidence of those in its competitors will fall. The attractiveness of *i* relative to others grows, thus strengthening *i* and weakening its competitors. The net flow of practitioners into Paradigm *i* will increase the gap in solved puzzles between *i* and its competitors, causing the gap in confidence to widen still further. The self-reinforcing differentiation continues until one paradigm emerges dominant and the others become extinct. These same loops are responsible for the resistance of the dominant paradigms to challenges, as high confidence suppresses the creation and retards the progress of new theories. High confidence leads to normal science and low anomalies, suppressing the type of inquiry likely to lead to the creation of new paradigms [Equation (2)]. And should by chance a new theory be created, the high confidence and low anomalies of a dominant paradigm make it unlikely a new theory can succeed, even if it has high intrinsic explanatory potential. Note that once a dominant paradigm begins to experience depletion of its root metaphor, these same loops operate as vicious cycles, accelerating the collapse.

In the early phase of a competition between two or more paradigm candidates, when the differences among the competing theories are small, chance events can perturb the system sufficiently to shift the advantage to a previously weaker rival. Such random events might include factors related to the theory, such as the announcement of an important experimental result, but can also include events wholly outside of science, such as the illness of the candidate's champion or political upheavals that disrupt the work of key people. However, as the positive loops confer greater and greater advantage to one of the contending theories, the likelihood that particular events can overcome the advantage of the leader rapidly diminishes, until the system has effectively "locked in" to a solution. Once such lock-in has occurred, the dominance of the winning theory is assured (until its own crisis). Yet which particular theory becomes dominant can be a matter of chance events and small perturbations early in the emergence phase.

The prevalence of positive feedback processes in the dynamics means that historical contingencies attending the creation and early years of a new theory strongly condition their fate. While it is obvious that the creation of a new theory is intrinsically unpredictable, the simulation shows clearly that, once created, the likelihood any given new paradigm survives its founding and grows to dominance is strongly contingent on the environment into which it is launched—an environment that in turn depends on the history of the paradigms preceding it. The prevalence of positive feedback processes in paradigm development means that the evolution of the system as a whole is strongly path-dependent.

The ability of positive feedback processes to create path-dependent lock-in to particular equilibria from an initially undifferentiated choice set has been amply documented in biological, economic, technological, and other systems. Examples beyond the familiar QWERTY keyboard include the universal left-handed chirality of proteins throughout the plant and animal kingdom, the choice of technological standards such as the gauge for a railroad or the shape of electrical plugs, the designation of Greenwich as the prime meridian, the length of the standard meter in Paris (or the choice of the metric over the English system), the dominance of the IBM/Microsoft Windows architecture for personal computers, and the growing dominance of the major world languages while the languages of small indigenous peoples become extinct.

Even when all choices are equally attractive ex ante as in the choice of the length of the standard meter or the shape of electrical plugs, the symmetry is broken by microscopic noise or external perturbations. The positive feedbacks then amplify these small initial differences to macroscopic significance. Once a dominant design has emerged, the costs of switching become prohibitive, so the equilibrium is self-enforcing, at least until there is an architectural shift that renders the dominant design obsolete (Henderson and Clark 1991), as in the replacement of analog broadcast television by HDTV.

5. Intrinsic Capability or Historical Contingency?

To test the argument above and quantify the roles of intrinsic versus contingent factors, we analyzed the pooled results of 57 2000-year model runs. The only parameters varied were the paradigm's intrinsic explanatory power and the random number seed affecting the launch of new paradigms. To eliminate initial transients and end effects, the first and last five paradigms of each simulation are eliminated from the analysis. There are 350 dominant paradigms and 676 never-dominant paradigms in the sample.[2]

We consider a LOGIT model with three explanatory variables: intrinsic capability (C), the confidence in the dominant paradigm at the time the new paradigm is launched (CP^{dom}), and the number of competitor paradigms (not including the dominant paradigm) each new paradigm faces when launched. Since the probability of success need not depend linearly on the number of competitors, we treat the number of competitors as a categorical variable. Thus, the dummy variable $COMPET_i = 1$ if the number of competitors equals i at the time each paradigm is founded, and zero otherwise, for situations of up to four competitors:

$$P_t(\text{Dom}) = 1/(1 + \exp(-(b_0 + b_1 C + b_2 CP_t^{\text{dom}} + \sum_{i=1}^{4} w_i COMPET_{i,t}))), \quad (3)$$

where the subscript t indicates that the probability is calculated in the year each new paradigm is created.

Table 1 shows how well the model predicts successes and failures, where an estimated probability greater than 0.5 is interpreted as a prediction that the paradigm becomes dominant and estimated probabilities <0.5 are interpreted as predictions of failure. Overall, 83% of the cases are predicted correctly. The sensitivity and specificity of the model are roughly equal (both ≥ 0.83), indicating the model's error rate is about the same for predictions of dominance when the paradigm in fact fails versus predictions of failure when the paradigm in fact succeeds. The statistic $\lambda_b = 1 -$ ((errors|model)/(errors|no model)) measures how much the model improves prediction success compared to the chance success rate. In the absence of the model, the best guess is that any paradigm picked at random fails, since fully two-thirds of the paradigms in the sample never become dominant.

Table 1 Ability of LOGIT Model To Predict a Given Paradigm's Rise to Dominance

The model reduces the error rate in predicting dominance by half compared to chance.

		Predicted		
		Nondominant	Dominant	Total
Actual	Nondominant	641	35	676
	Dominant	135	215	350
	Total	776	250	1026

$\lambda_b = 0.51$; *Proportion correct* = 0.83; *Sensitivity* = 0.86; *Specificity* = 0.83

The model reduces the error rate by half compared to chance.

The regression results (Table 2) show that all estimated coefficients have the predicted signs. A new paradigm's chances of success rise with greater intrinsic capability, a weaker dominant paradigm, and a smaller number of competitors. However, the effect of a paradigm's intrinsic explanatory potential, C, on its probability of success is not significant, while the contextual variables are highly significant. In particular, the confidence level of the dominant paradigm at the time a new contender is created has a strong effect on the challenger's likelihood of success. Similarly, the chances of success fall precipitously as the number of competitors rises. The estimated coefficients illustrate the weak role of intrinsic capability in comparison to the contextual factors in determining whether a paradigm becomes dominant. Table 3 shows how the

Table 2 LOGIT Regression Comparing Intrinsic and Contextual Factors in Likelihood of Success

Variables characterizing the competitive environment have a strong impact on the likelihood of success while the intrinsic explanatory power of a paradigm (C) has only a weak effect.

Indep. Variable	Estimated Coeff.	Standard Error	t-statistic
Constant	5.44*	0.52	10.42
C	6.86e-4	4.34e-4	1.58
CP^{dom}	−7.27*	0.55	−13.19
$COMPET_1$	−1.43*	0.23	−6.17
$COMPET_2$	−4.99*	0.52	−9.54
$COMPET_3$	−13.52	50.00	−0.27
$COMPET_4$	−14.65	147.91	−.099

$N = 1026$
$^*P < 0.05$

Table 3 The Influence of Intrinsic and Contextual Factors on the Probability a New Paradigm Becomes Dominant

Conditions at Emergence	Probability of Dominance (with 95% confidence interval)
One Competitor	$0.18 \leq 0.25 \leq 0.40$
Two Competitors	$0.004 \leq 0.01 \leq 0.05$
$C = 100$, No Competitors	$0.44 \leq 0.53 \leq 0.65$
$C = 800$, No Competitors	$0.52 \leq 0.63 \leq 0.79$
$C = 100$, One Competitor	$0.15 \leq 0.22 \leq 0.38$
$C = 800$, One Competitor	$0.21 \leq 0.31 \leq 0.57$

probability that a given paradigm rises to dominance depends on its intrinsic capability compared to the contingent factors. Each row shows the probability a new paradigm becomes dominant given the conditions listed and assuming all other explanatory variables take their mean values, along with the 95% confidence interval for the probability.[3] The table clearly illustrates the relatively weak influence of intrinsic capability (C) on the probability of becoming dominant. On average, a paradigm launched in the presence of one competitor has only a 25% chance of succeeding. If there are two competitors the probability drops to only 1%. Even if a new paradigm is endowed with the maximum amount of intrinsic capability ($C = 800$), the consequences for dominance are still strongly mediated by the contingent factors at emergence. Thus, for example, while the probability a new paradigm becomes dominant reaches 0.63 if it faces no competitors, it has only a 0.31 chance of surviving when it faces one additional competitor. Furthermore, increasing the intrinsic capability of a new paradigm by a factor of eight boosts the probability of success by only about ten percentage points. Contingent factors at the time of emergence far outweigh the influence of intrinsic capability.

The relative importance of intrinsic capability C versus the contextual factors CP^{dom} and the number of competitors $COMPET_i$ is further illustrated in Figure 12. Each point in the plot represents the probability of dominance of a particular paradigm, as predicted by its intrinsic capability, the number of competitors it faces at birth (excluding the dominant paradigm), and the confidence of the dominant paradigm it faces. The smooth curves plot the predicted probability of dominance as CP^{dom} varies over the [0,1] interval, for each number of competitors and assuming intrinsic capability takes on its mean value $C_{\mathrm{avg}} = 371.4$; that is:

Figure 12 Logit Model Results

The probability a given paradigm rises to dominance as it depends on confidence in the dominant paradigm and the number of competitors at the time it is created.

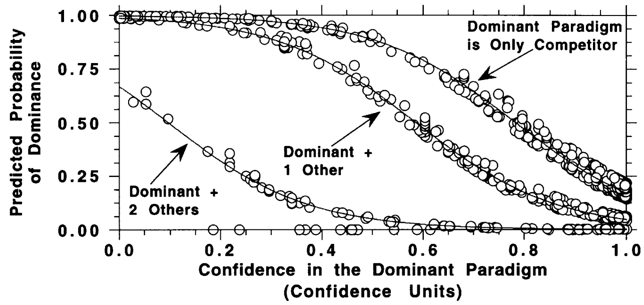

$$P_t(\mathrm{Dom}) = 1/(1 + \exp(-(5.44 + 0.000686C_{\mathrm{avg}} - 7.27CP^{\mathrm{dom}} + w_i))). \quad (4)$$

For new paradigms competing only against the dominant paradigm, the probability of dominance is given by the curve in the upper right. Curves are also displayed for environments with two and three competitors. The curve for four competitors has probabilities ≈ 0. For all but the smallest values of CP^{dom}, the greater the number of competitors, the less likely a new paradigm becomes dominant. Likewise, the greater the value of CP^{dom}, the less likely a new paradigm is to become dominant. The regression results and Figure 12 show the number of competitors existing at the time a new paradigm is created strongly influences its fate. When CP^{dom} is between about 0.1 and 0.6, a new paradigm stands a better than even chance of becoming dominant if it faces a total of two competitors or less, and will likely fail if there are three or more competitors. When CP^{dom} is between about 0.6 and 0.8, the new paradigm is more likely than not to become dominant if it faces only the dominant paradigm, likely to fail if it faces two competitors, and almost sure to die if faces three or more competitors.

Thus the likelihood that a new paradigm will rise to dominance in the model is overwhelmingly determined by historical contingencies and only weakly influenced by its intrinsic explanatory power. The relative importance of inherently unpredictable situational factors is not particularly sensitive to the parameters. Rather it is a consequence of the many positive feedbacks by which paradigms bootstrap themselves from doubt to normal science (Figure 10).

But how do internal and contextual factors interact to determine the longevity of those paradigms that survive their founding and go on to dominate their field? Do intrinsically powerful paradigms remain dominant longer than their weaker counterparts? Here one would expect that the paradigms with greater explanatory power should survive longer. Figure 13 shows a paradigm's longevity as a function of its intrinsic capability only for those paradigms that went on to become dominant. The figure shows those successful paradigms that emerged when the dominant paradigm against which they had to compete was strong ($CP^{dom} \geq 0.75$; $N = 131$) and those that emerged when the dominant paradigm was weak ($CP^{dom} \leq 0.25$; $N = 104$).[4] As expected, for those paradigms surviving their founding, longevity is significantly related to intrinsic capability. In both cases, longevity roughly follows a power law in capability $L = \alpha C^{\beta}$. Such power law scaling is common in a wide range of dynamical systems (Schroeder 1991).

However, Figure 13 shows that even for successful paradigms, historical contingencies matter greatly to their longevity. Those paradigms emerging when the dominant paradigm is very strong (with confidence ≥ 0.75) actually survive significantly *longer* than those emerging when their principal competitor is weak (confidence ≤ 0.25). The median longevity for those emerging when $CP^{dom} \geq 0.75$ is more than twice as great than that for those emerging when $CP^{dom} \leq 0.25$. The differences in outcomes arise from differences in the circumstances attending the birth of these successful theories. Paradigms emerging when confidence in the dominant paradigm is relatively high face strong competition. Most scientists are still satisfied with the dominant paradigm, so the rate of recruitment to the new paradigm is relatively slow. During this time, however, the few adherents of the new paradigm are able to solidify the foundations of their theory and develop skill with their tools and techniques. Anomalies remain low as practitioners solve the relatively easy puzzles for which their paradigm is well suited. Their confidence rises. By the time the crisis of the dominant paradigm deepens and its members become disaffected, the initial adherents of the new theory will have articulated it well enough to provide an attractive and viable alternative. With high confidence, skill, and a productive agenda to focus research on the puzzle solving of normal science, the new paradigm is poised to realize its intrinsic potential. For these paradigms, longevity follows the power law scaling with intrinsic capability fairly closely ($L \propto C^{.78}$, $R^2 = .64$).

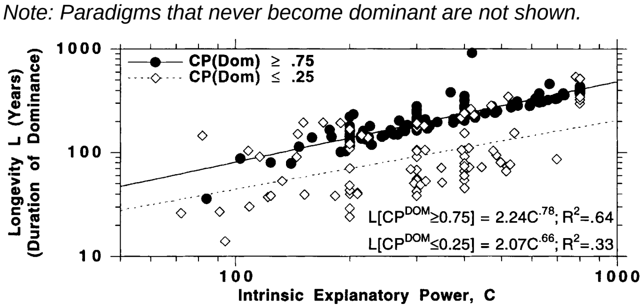

Figure 13 **Relationship between Longevity (L) of Successful Paradigms and Intrinsic Explanatory Power (C)**

Note: Paradigms that never become dominant are not shown.

Paradigms emerging when confidence in the dominant paradigm is low face a competitor in crisis. Thus, as new and unproven as the new paradigm is, it nonetheless quickly wins new members. The rapid influx of new practitioners means the rate of effort is high. Rapid growth in activity means the average difficulty of puzzles rises quickly, increasing the number of unresolved puzzles likely to be seen as anomalies. Most important, the influx of new practitioners occurs when confidence is low, meaning basic disagreements about methods, data, and criteria for validity still persist. Without the learning and skill experience afford, without the acculturation and perceptual filters provided by a well-articulated paradigm, disagreements and anomalies arise at an alarming rate. If enough anomalies accumulate, confidence can fall. Falling confidence causes people to perceive anomalies still more readily, further decreasing confidence. The new paradigm rapidly disintegrates, its high intrinsic potential largely unrealized. For paradigms emerging when their principal competitor is weak, longevity scales with intrinsic capability only as $L \propto C^{.66}$, and the variance of longevity around the best fit is much greater.

The results show the strong role of contingent, historical factors even for those paradigms that become dominant. As expected, the probability of surviving the founding period and becoming dominant is negatively related to the intensity of the competitive environment. However, counter to what one might expect, the more intense the competition, the *longer* the expected life of the successful theories. There are two reasons. First, strong selective pressures during the emergence phase ensure that only those paradigm candidates with high intrinsic capability can survive. When selection pressure is weak, some paradigm candidates with low intrinsic potential can become dominant. Second, and even more insidiously, when competition is weak many paradigm candidates with high intrinsic potential die young as they grow too rapidly, overextending themselves before their members develop enough skill, understanding, and confidence to prevent the accumulation of anomalies. Historical contingencies not only determine which paradigms succeed but also how long those that succeed may thrive.

6. Discussion: Guru Dynamics and Management Fads

The dynamics generated by the model resemble the life cycle of intellectual fads. Often a promising new idea rapidly becomes fashionable through excessive optimism, aggressive marketing, media hype, and popularization by gurus. Many times the rapid influx of poorly trained practitioners, or the lack of established protocols and methods, causes expectations to outrun achievements, leading to a backlash and disaffection. Such fads are commonplace, especially in (quack) medicine and most particularly in the world of business, where "new paradigms" are routinely touted in the pages of popular journals of management, only to be displaced in the next issue by what many business people have cynically come to call the next "flavor of the month" (see Abrahamson 1996). No doubt many such fads have no intrinsic merit (in our terms, intrinsic capability C is low) so their rapid demise is the desired and rational outcome (similar to the fate of the low potential Paradigm 13 in Figure 8). However, too many of these fads achieve broad acceptance and lead to large expenditures, only to suffer a backlash when they fail to live up to their promise.

The theory developed here helps explain how this occurs. Typically, a guru proposes a new theory, tool, or process promising to address persistent problems facing businesses (that is, a new paradigm claiming to solve the anomalies that have undermined the old paradigm.) The early adopters of the guru's method spread the word and initiate some projects. Even in cases where the ideas of the guru have little merit, the energy and enthusiasm a team can bring to bear on a problem, coupled with Hawthorne and placebo effects and the existence of "low hanging fruit" will often lead to some successes, both real and apparent. Proponents rapidly attribute these successes to the use of the guru's ideas. Positive word of mouth then leads to additional adoption of the guru's ideas. (Of course, failures are covered up and explained away; as in science there is the occasional fraud as well.) Media attention further spreads the word about the apparent successes, further boosting the credibility and prestige of the guru and stimulating additional adoption.

As people become increasingly convinced that the guru's ideas work, they are less and less likely to seek or attend to disconfirming evidence. Management gurus and their followers, like many scientists, develop strong personal, professional, and financial stakes in the success of their theories, and are tempted to selectively present favorable and suppress unfavorable data, just as scientists grow increasingly unable to recognize anomalies as their familiarity with and confidence in their paradigm grows. Positive feedback processes dominate the dynamics, leading to rapid adoption of those new ideas lucky enough to gain a sufficient initial following. Hirshleifer (1995) and Bikhchandani et al. (1992) present similar models of fads caused by positive feedbacks, and Sastry (1998) and Sastry and Coen (1998) discuss positive feedbacks in organizations. Of course formal models of innovation diffusion as a process driven by positive feedback go back at least to Bass (1969), and conceptual models of such positive feedback processes can be traced to Myrdal's (1944) "principle of cumulative causation," Merton's (1948) theory of the self-fulfilling prophecy, and J. S. Mill's (1848) theory of speculative bubbles (see Richardson 1991 for the history of feedback theories in the social sciences). More recent work discusses the differences between the diffusion of ideas and of technologies, and the role of social networks and other factors in conditioning the strength of the positive loops driving adoption, e.g., Rogers (1995), Valente (1995), and Kaufer and Carley (1993).

The wide range of positive feedbacks identified above can lead to the swift and broad diffusion of an idea with little intrinsic merit because the negative feedbacks that might reveal that the tools don't work operate with very long delays compared to the positive loops generating the growth. In science there are often long delays between the initial success of a theory and the execution and interpretation of experiments that can test it. In the world of social action, the delays are often even longer. Rigorous follow up studies to assess the effectiveness of a new management tool are notoriously difficult because of the inability to conduct controlled experiments in social systems, the essential participation of human beings in the interventions, and the ambiguity of outcomes. The combination of strong positive feedbacks promoting the growth of new management ideas and slow, weak negative feedbacks revealing which are wheat and which chaff predisposes the world of management to a succession of highly touted "new paradigms," each shining brilliantly for a few brief years only to be discarded once the negative feedbacks of follow-up evaluation lead to disaffection and the advent of a new guru with a new, more attractive theory. The same positive feedbacks can also lead to inflated expectations, insufficient practitioner skill, overly broad scope of application, and inadequate time to resolve anomalies, causing some theories with high intrinsic capability to be abandoned too soon, as seen in the simulations.

As discussed, our model shows that the likelihood a new theory will be created and gain significant popularity is endogenous, rising as confidence in existing theories falls. Thus we would predict a higher incidence of management fads during times of economic and social stress, when confidence in existing institutions and their motivating ideologies falters. Indeed, the rise of management fads has coincided with the slow growth, downsizing, globalization, rapid technical change, and other pressures of the past few decades. These stresses constitute the anomalies eroding confidence in existing organizational structures and political ideologies. At the same time, our model predicts that low confidence in existing institutions increases both the number of new theories lacking intrinsic merit that gain significant popularity and the number of high potential ideas that die young as a result of the skill dilution and insufficient learning caused by rapid growth. Though economic stress may stimulate management innovation, it also increases the probability businesses will both embrace useless theories *and* prematurely discard potentially useful ones.

The results of our model suggest that the long-term success of new theories can be enhanced by slowing the positive feedback processes, such as word of mouth, marketing, media hype, and extravagant claims of efficacy by which new theories can grow, and strengthening the processes of theory articulation and testing, which can enhance learning and puzzle-solving capability.

So-called "chaos" or "complexity theory" itself provides a recent example. The practical value of nonlinear dynamics has repeatedly been demonstrated in physics and the life sciences (see Chin et al. 1996, Costantino et al. 1997, and Sturis et al. 1991 for recent examples). However, rapid growth, fed by successful popularization (e.g., Gleick 1987, Waldrop 1992) and ill-advised claims for the universality of "complexity" as a "new paradigm" for the reconstruction of the social as well as natural sciences have already led to a backlash (for example, Horgan 1995). Developing the full potential of complexity theory, especially in the social sciences, requires more rigorous theory development and fewer popular articles extolling the virtues of the "new paradigm", more studies testing the new theories and fewer anecdotal claims of efficacy, greater development of tools tailored for particular contexts, and fewer claims of universality. Without such rigor, social scientists face the danger that, despite its high potential, "complexity theory" will soon be discarded, perhaps prematurely, as yet another unfortunate case of physics envy.

Testing our theory against real-world examples such as the emergence of complexity theory poses daunting but not insurmountable challenges. Testing the model empirically requires measuring model constructs such as "confidence," "anomaly," and "average difficulty of puzzles." Confidence might be measured through surveys or interviews with relevant researchers, asking them to rate their degree of belief in the theory. Content analysis of publications in the field would also reveal the strength and universality of the claims made by key practitioners, indicating their confidence level. Content analysis might also be used to analyze critical reactions to particularly thorny problems, thus identifying potential anomalies. Bibliometric techniques could be used to determine how long a research problem ("puzzle") has gone unsolved and gauge the number of researchers working on it, to yield a measure of the difficulty of puzzles. Donovan et al. (1988), and Jacobsen and Bronson (1995) discuss the practical difficulties involved in such empirical tests; Rappa and Debackere (1993) use survey and bibliometric tools to shed light on the demographics and attitudes of scientists in several fields, illustrating how the constructs in the model might be measured.

7. Conclusion

Before turning to the conclusions, we pause to consider the limitations of the model. All models (formal or otherwise) are inevitably less than the world their authors seek to portray. We agree with Cartwright (1983, p. 153) that models "are a work of fiction." Of course the model is not comprehensive, nor does it capture all the subtleties of Kuhn's theory. Rather, we seek to demonstrate that it is both desirable and possible to portray in a formal model the causal hypotheses embodied in written theories of scientific endeavor and test whether they can generate the dynamics as those authors see them. The process of formalizing such hypotheses helps to identify inconsistencies, implicit assumptions, glosses, and errors in the mental simulations authors necessarily perform to infer the dynamics of science from their theories of its structure. Such an endeavor is worthwhile as a complement to historical and sociological studies. Complete documentation of the model is available; we invite others to replicate, critique, revise, and extend the model to test views of scientific development different from ours.

The simulations suggest an important role for situational contingencies in the evolution of science. We find that the fate of a particular new theory or paradigm is strongly conditioned by the circumstances surrounding its creation, and only weakly influenced by its explanatory power or logical force (at least for theories above a minimum threshold of explanatory power). Environmental conditions at the time a new theory is created, such as the morale and confidence of practitioners in the old paradigm and the number of contending alternative new theories, powerfully determine whether a new theory will rise to dominance or quickly perish. In particular, the simulations show new theories with great explanatory power frequently fail to attract a critical mass of adherents, while weaker ones often triumph. The frequent eclipse of the strong by the weak is not a pathological outcome, but rather a normal consequence of scientific activity as we have modeled it.

The interplay between intrinsic explanatory potential and historical contingency is quite subtle. A paradigm's inherent potential—its logical force and power to explain nature—does influence its future development: of those paradigms surviving their youth, those with high intrinsic capability do remain dominant longer, on average, than those that are weaker. But the impact of intrinsic capability on the longevity of any given paradigm is mediated by the competitive conditions in the emergence period. In particular, weak competitive environments make it more likely a new paradigm will rise to dominance, but can condemn even powerful paradigms to early deaths as they are extended too far and too fast, generating anomalies and prematurely destroying confidence. On the other hand, though competition reduces the likelihood of survival, competition gives those that do survive time to bootstrap themselves into normal science, insulating them against mere disconfirmation, and ensuring they persist until the anomalies ultimately causing revolution, in Kuhn's words, "penetrate existing knowledge to the core" (Kuhn, 1970, p. 65).

Most important, however, competition does *not* serve to weed out the weak paradigms so the strong may grow. On the contrary, competition decimates the strong and weak alike—we found that intrinsic capability has but a weak effect on survival. The mortality rate for paradigms seems to depend almost entirely on the environmental conditions surrounding their birth. This is a sobering result, since we can never know the microlevel contingencies of history that can prove decisive; here favoring an intrinsically weak paradigm, there killing an intrinsically strong theory.[5] These characteristics of the competition among paradigms are consequences of the powerful positive feedback processes operating within and among paradigms. These positive loops can amplify microscopic perturbations in the environment—the local conditions of science, society, and self faced by the creators of a new theory—until they reach macroscopic significance. Such dynamics are the hallmark of path dependent evolutionary systems.

Contemplating the reflexive feedbacks between people and the world, Kuhn (1990, p. 7), in "The road since *Structure*," captured the essence of path-dependence in evolutionary systems, arguing that "scientific development must be seen as a process driven from behind, not pulled from ahead—as evolution from, rather than evolution towards." Yet while acknowledging the role of the biological, cognitive, and social in the evolution of science, Kuhn (1990, p. 10) argues forcefully that path-dependence does not mean the course of scientific development is entirely arbitrary or reality merely a social construction:

> . . . the world is not invented or constructed. . . . [It] has been experientially given, in part to the new inhabitants directly, and in part indirectly, by inheritance, embodying the experience of their forebears. As such, it is entirely solid: not in the least respectful of an observer's wishes and desires; quite capable of providing decisive evidence against invented hypotheses which fail to match its behavior. Creatures born into it must take it as they find it. They can, of course, interact with it, altering both it and themselves in the process, and the populated world thus altered is the one that will be found in place by the generation which follows.

Through these feedbacks the world we inhabit is made; it is a world of nonlinear, disequilibrium dynamics in which, as Kuhn (1990, p. 12) says, "small changes . . . can have large-scale effects."

Acknowledgments

The authors thank Kathleen Carley, David Edge, Erik Mosekilde, Michael Rappa, Anjali Sastry, Frank Sulloway, the referees, and seminar participants at MIT for helpful comments.

Endnotes

[1]Sterman (1985) provides a formal model of Kuhn's theory representing the life cycle of a single paradigm; full documentation is provided. In this paper we extend the original model to allow for explicit competition among different theories.

[2]In most of the simulations, intrinsic capability, C, for each paradigm was drawn randomly from a lognormal distribution truncated such that $C \leq 800$. In some runs all paradigms had identical intrinsic capabilities, with $C = 200$, 300, or 400, to further reduce the variance and isolate the role of historical contingencies. These restrictions do not affect the model's qualitative behavior.

[3]To compute the 95% confidence intervals we drew 1,000 simulated parameter estimates from a multivariate normal distribution defined by the estimated coefficients and variance–covariance matrix, sorted the resulting estimates of the probability of dominance, and extracted the probabilities from the 25th and 975th values (see King et al. 1998).

[4]We omit from Figure 13 those paradigms emerging when the confidence level of the dominant paradigm was between 0.25 and 0.75 to simplify the presentation; including the full sample does not significantly alter the power law scaling or the result that longevity depends on contingencies including the confidence of the principal competitor.

[5]See Gould (1990) for a similar view applied to the evolution of life.

References

Abrahamson, E. 1996. Management fashion. *Acad. Management Rev.* **21**(1) 254–285.

Anderson, P., K. Arrow, D. Pines, eds. 1988. *The Economy as an Evolving Complex System.* Addison-Wesley, Reading, MA.

Arthur, B. 1989. Competing technologies, increasing returns, and lock-in by historical events. *Econom. J.* **99** 116–131.

———. 1994. *Increasing Returns and Path Dependence in the Economy.* University of Michigan Press, Ann Arbor.

Bass, F. 1969. A new product growth model for consumer durables. *Management Sci.* **15** 215–227.

Barlas, Y. 1992. Comments on "On the very idea of a system dynamics model of Kuhnian science" *System Dynam. Rev.* **8**(1) 43–47.

Bijker, W. E., T. P. Hughes, T. J. Pinch, eds. 1987. *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology.* MIT Press, Cambridge, MA.

Bikhchandani, S., D. Hirshleifer, I. Welch. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Political Econom.* **100** (5) 992–1026.

Black, M. 1954–1955. Metaphor. *Proc. Aristotelian Soc.* **55** 284–285.

Bruckner, E., W. Ebeling, A. Scharnhorst. 1989. Stochastic dynamics of instabilities in evolutionary systems. *System Dynam. Rev.* **5**(2) 176–191.

———. 1990. The application of evolution models in scientometrics. *Scientometrics* **18** (1–2) 21–41.

Carley, K. 1991. A theory of group stability. *Amer. Sociological Rev.* **56** 331–354.

———. 1995. Computational and mathematical organization theory: Perspective and directions. *Comput. Math. Organization Theory* **1**(1), 39–56.

———, W. Wallace. 1995. Editorial. *Comput. Math. Organ. Theory* **1**(1) 5–7.

Cartwright, N. 1983. *How the Laws of Physics Lie.* Clarendon Press, Oxford.

Chaisson, E. 1997. Cosmic age controversy is overstated. *Sci.* **276** (16 May) 1089–1090.

Chin, G., L. Senesac, W. Blass, J. Hillman. 1996. Stabilizing lead-salt diode lasers: Understanding and controlling chaotic frequency emissions. *Sci.* **274** (29 November) 1498–1501.

Collins, H. 1985. *Changing Order: Replication and Induction in Scientific Practice.* Sage Publications, Beverly Hills.

Costantino, R., R. Desharnais, J. Cusing, D. Dennis. 1997. Chaotic dynamics in an insect population. *Sci.* **275** (17 January) 389–391.

Crane, D. 1972. *Invisible Colleges: Diffusion of Knowledge in Scientific Communities.* University of Chicago Press, Chicago.

Cushing, J. 1994. *Quantum Mechanics: Historical Contingency and the Copenhagen Hegemony.* University of Chicago Press, Chicago.

Cyert, R., J. March. 1963/1992. *A Behavioral Theory of the Firm,* 2nd ed. Blackwell, Cambridge, MA.

Donovan, A., L. Laudan, R. Laudan, eds. 1988. *Scrutinizing Science.* Kluwer Academic Publishers, Dordrecht, The Netherlands.

Ebeling, W. 1991. Stochastic models of competition processes in nonphysical systems. *Syst. Anal. Model. Simul.* **8** (1) 3–17.

Eldredge, N., S. Gould. 1972. Schopf, T., ed. *Models in Paleobiology.* Freeman, San Francisco. 82–115.

Feyerabend, P. 1975. *Against Method.* New Left Books, London.

Forrester, J. 1961. *Industrial Dynamics.* MIT Press, Cambridge MA.

Gersick, C. 1991. Revolutionary change theories: A multilevel exploration of the punctuated equilibrium paradigm. *Acad. Management Rev.* **16** (1) 10–36.

Gleick, J. 1987. *Chaos: Making a New Science.* Viking, New York.

Goodman, N. 1968. *Languages of Art.* Bobbs-Merrill, Indianapolis.

Gorman, M. 1992. *Simulating Science: Heuristics, Mental Models and Technoscientific Thinking.* Indiana University Press, Bloomington.

Gould, S. J. 1990. *Wonderful Life.* W. W. Norton, New York.

Gross, P., N. Levitt. 1994. *Higher Superstition: The Academic Left and Its Quarrels with Science.* Johns Hopkins University Press, Baltimore, MD.

Hall, R. 1976. The system pathology of an organization: The rise and fall of the old *Saturday Evening Post. Admin. Sci. Quart.* **33** 1–23.

Hannan, M. and J. Freeman. 1989. *Organizational Ecology.* Harvard University Press, Cambridge, MA.

Haxholdt, C., C. Kampmann, E. Mosekilde, J. D. Sterman. 1995. Mode locking and entrainment of endogenous economic cycles. *System Dynam. Rev.* **11**(3) 177–198.

Henderson, R., K. Clark. 1991. Architectural innovation: The reconfiguration of existing technologies and the failure of established firms. *Admin. Sci. Quart.* **35**(1).

Hirshleifer, D. 1995. The blind leading the blind: Social influence, fads, and informational cascades. M. Tomassi, K. Ierulli, eds. *The New Economics of Human Behavior.* Cambridge University Press, Cambridge. 188–215.

Holland, J. 1995. *Hidden Order: How Adaptation Builds Complexity.* Addison-Wesley, Reading, MA.

Horgan, J. 1995. From complexity to perplexity. *Sci. Amer.* (June) 104–109.

Hoyningen-Huene, P. 1993. *Reconstructing Scientific Revolutions: Thomas S. Kuhn's Philosophy of Science.* University of Chicago Press, Chicago.

Jacobsen, C., R. Bronson. 1995. Computer simulations and empirical testing of sociological theory. *Sociological Methods and Res.* **23**(4) Sage, Beverly Hills. 479–506.

Kaufer, D. S., K. M. Carley. 1993. *Communication at a Distance: The Influence of Print on Sociocultural Organization and Change.* Lawrence Erlbaum Associates, Hillsdale, NJ.

King, G., M. Tomz, J. Wittenberg. 1998. Making the most of statistical analyses: Improving interpretation and presentation. Working paper. Department of Government, Harvard University.

Kuhn, T. S. 1970. *The Structure of Scientific Revolutions,* 2nd ed. University of Chicago Press, Chicago.

———. 1990. The road since structure. *PSA* **2**(7) 3–13.

Lakatos, I., A. Musgrave. 1970. *Criticism and the Growth of Knowledge.* Cambridge University Press, Cambridge.

Lant, T., S. Mezias. 1992. An organizational learning model of convergence and reorientation. *Organ. Sci.* **3**(1) 47–71.

Lightman, A., O. Gingerich. 1992. When do anomalies begin? *Sci.* **255** (7 February) 690–695.

Lomi, A., E. Larsen. 1996. Interacting locally and evolving globally: A computational approach to the dynamics of organizational populations. *Acad. Management J.* **39**(4) 1287–1321.

Martin, B. 1996. Social construction of an "attack on science." *Soc. Stud. Sci.* **26** 161–173.

Masterman, M. 1970. The nature of a paradigm. I. Lakatos, A. Musgrave, eds. *Criticism and the Growth of Knowledge.* Cambridge University Press, Cambridge.

Merton, R. 1948. The self-fulfilling prophecy. *Social Theory and Social Structure.* Free Press, New York.

Mill, J. S. 1848. *Principles of Political Economy.* J. W. Parker, London.

Morecroft, J. 1985. Rationality in the analysis of behavioral simulation models. *Management Sci.* **31**(7) 900–916.

Myrdal, G. 1944. *An American Dilemma.* Harper, New York.

Nelson, R., S. Winter. 1982. *An Evolutionary Theory of Economic Change.* Belknap, Harvard University, Cambridge, MA.

Nadis, S. 1996. Oil drop trap is set for lone quark. *Sci.* **271** (8 March) 1361–1362.

Orlikowski, W. J. 1996. Improvising organizational transformation over time: A situated change perspective. *Inform. Systems Res.* **7**(1) 63–92.

Pickering, A. 1984. *Constructing Quarks: A Sociological History of Particle Physics.* University of Chicago Press, Chicago.

———, ed. 1992. *Science as Practice and Culture.* University of Chicago Press, Chicago.

Pinch, T. 1986. *Confronting Nature: The Sociology of Solar-Neutrino Detection.* D. Reidel, Dordrecht, Holland.

Radzicki, M. 1992. Reflections on "On the very idea of a system dynamics model of Kuhnian science." *System Dynam. Rev.* **8** (1) 49–53.

———, J. D. Sterman. 1994. Evolutionary economics and system dynamics. R. England. ed. *Evolutionary Concepts in Contemporary Economics.* University of Michigan Press, Ann Arbor. 61–89.

Rappa, M., K. Debackere. 1993. Youth and scientific innovation: The role of young scientists in the development of a new field. *Minerva* **31** (1) 1–20.

Richardson, G. 1991. *Feedback Thought in Social Science and Systems Theory.* University of Pennsylvania Press, Philadelphia.

———. 1996. Problems for the future of system dynamics. *System Dynam. Rev.* **12**(2) 141–157.

———, A. Pugh. 1981. *Introduction to System Dynamics Modeling with DYNAMO.* Productivity Press, Portland, OR.

Rogers, E. M. 1995. *Diffusion of Innovations.* The Free Press, New York.

Roth, P., R. Barrett. 1990. Deconstructing quarks. *Soc. Stud. Sci.* **20** 579–632.

Sastry, M. A. 1997. Problems and paradoxes in a model of punctuated organizational change. *Admin. Sci. Quart.* **42**(2) 237–275.

———. 1998. Archetypal self-reinforcing structures in organizations: A system dynamics perspective of cognitive, social, and institutional processes. Proceedings of the 1998 International System Dynamics Conference, Quebec City, Quebec.

———, C. Coen. 1998. Tracing the hand of the past: Towards an integrated theory of how history matters in organizations. Working paper, University of Michigan Business School.

Schroeder, M. 1991. *Fractals, Chaos, Power Laws.* W. H. Freeman, New York.

Sterman, J. D. 1985. The growth of knowledge: Testing a theory of scientific revolutions with a formal model. *Tech. Forecasting and Soc. Change* **28** (2) 93–122.

———. 1988. Deterministic chaos in models of human behavior: Methodological issues and experimental results. *System Dynam. Rev.* **4** 148–178.

———. 1989. Deterministic chaos in an experimental economic system. *J. Econom. Behavior and Organ.* **12** 1–28.

———. 1992. Response to "On the very idea of a system dynamics model of Kuhnian science." *System Dynam. Rev.* **8** (1) 35–42.

Stewart, J. 1986. Drifting continents and colliding interests: A quantitative application of the interests perspective. *Soc. Stud. Sci.* **16** 261–279.

Sturis, J., K. S. Polonsky, E. Mosekilde, E. Van Cauter. 1991. Computer model for mechanisms underlying ultradian oscillations of insulin and glucose. *Amer. J. Physiology* **260** (Endocrinol. Metab. 23), E801–E809.

Thagard, P. 1968. *Computational Philosophy of Science.* MIT Press, Cambridge, MA.

Turner, R. S. 1987. Paradigms and productivity: The case of physiological optics, 1840–94. *Soc. Stud. Sci.* **17** 35–68.

Turbayne, C. M. 1970. *The Myth of Metaphor*, rev. ed. University of South Carolina Press, Columbia, SC.

Tushman, M., P. Anderson. 1986. Technological discontinuities and organizational environments. *Admin. Sci. Quart.* **31** 439–465.

———, E. Romanelli. 1985. Organizational evolution: A metamorphosis model of convergence and reorientation. *Res. Organizational Behavior* **7** 171–222.

Valente, T. W. 1995. *Network Models of the Diffusion of Innovations.* Hampton Press, Cresskill, N.J.

Van de Ven, A., M. Poole. 1995. Explaining development and change in organizations. *Acad. Management Rev.* **20**(3) 510–540.

Waldrop, M. 1992. *Complexity: The Emerging Science at the Edge of Order and Chaos.* Simon & Schuster, New York.

Weick, K. E. 1979. *The Social Psychology of Organizing*, 2nd ed. Random House, New York.

Weisbuch, G. 1991. *Complex Systems Dynamics: An Introduction to Automata Networks.* Addison Wesley, Reading, MA.

Wittenberg, J. 1992. On the very idea of a system dynamics model of Kuhnian science. *System Dynam. Rev.* **8** (1) 21–33.

Accepted by Kathleen M. Carley; received June 27, 1997. This paper has been with the authors for one revision.