In Search of Excellence[1]

David Strang and Michael W. Macy
Cornell University

The faddishness of the business community is often noted and lamented but not well understood by standard models of innovation and diffusion. We combine arguments about organizational cognition and institutional mimicry to develop a model of adaptive emulation, where firms respond to perceived failure by imitating their most successful peers. Computational experiments show that this process generates empirically plausible cascades of adoption, even if innovations are entirely worthless. Faddish cycles are most robust across alternative treatments of managerial decision making where innovations have modest positive effects on outcomes. These results have broad implications for the faddishness of a business community increasingly marked by media-driven accounts of success, and for the properties of organizational practices that are hot one day and cold the next.

> "The Atheist, when he was in Samothrace . . . was shown many vows and votive portraits from those who have survived shipwrecks and was then asked, 'You, there, who think that the gods are indifferent to human affairs, what have you to say about so many men saved by their grace?'—'It is like this,' he replied, 'there are no portraits here of those who stayed and drowned—and they are more numerous!'" (Montaigne, p. 44)

q1   The diffusion of innovations is a central problem within the social sciences (for reviews, see Rogers 1995; Strang and Soule 1998). Diffusion is a particularly valuable research site for sociology because it provides an opportunity to observe social influence. Sociologists have developed re-

lational accounts of the way adoption is channeled by social networks
(Coleman, Katz, and Menzel 1966; Burt 1987), interpretive accounts of
the discourse that catalyzes flow (Hirsch 1986; Snow 1993), and institu-
tional accounts of the growth of legitimacy and taken-for-grantedness
(DiMaggio and Powell 1983; Tolbert and Zucker 1983).

q2
q3

However, conventional models of diffusion present striking limitations.
While much work emphasizes the impact of *adoptions* elsewhere, there
is little attention to how actors respond to the *results* experienced by
others. Empirically, this is unfortunate since evidence suggests that success
is especially contagious (Holden 1986; Conell and Cohn 1995). It is also
theoretically costly, because ideas about rationality and effectiveness come
to be cast in opposition to ideas about imitation (DiMaggio and Powell
1983; Banerjee 1992).

Formal models of diffusion also exhibit fundamental predictive limi-
tations. They are well posed to explain rises in the number of adopters
but poorly equipped to account for almost anything else. This stems from
a theoretical and empirical focus on the boom periods of ultimately pop-
ular practices. But many practices never take off, and others experience
a short heyday followed by a dramatic decline in credibility and usage.
Such "faddish cycles," in particular, present an explanatory challenge that
is not met by conventional diffusion models.

This article addresses both problems by proposing a performance-
driven model that can generate cycles of adoption and abandonment. We
describe a process of "adaptive emulation" where actors respond to per-
ceived failure by imitating their most successful peers. This process is
illustrated within the corporate context, where faddish cycles are prom-
inent and where the behavioral supports for our argument are in evidence.
The model may be relevant to other contexts as well and may be useful
for understanding not only fads but also other possible diffusion trajec-
tories. We elaborate these extensions in the article's conclusion.

Our approach is motivated by a salient empirical feature of business
discourse on innovation: the centrality of the "success story." The biases
involved when successes but not failures are attended to, we contend, can
generate faddish cycles. This hypothesis directly challenges the conven-
tional wisdom that fads are a manifestation of lemming-like conformity.
We believe fads can result from *overattention* to the bottom line.

This article can be understood best, not as offering a novel model of
organizational cognition, but as investigating some of the most important
and influential diffusion arguments in sociology. Adaptive emulation com-

bines problem-driven models of search (March and Simon 1958) with DiMaggio and Powell's (1983) notion of organizational mimicry. Analysis of the behavior generated by this model permits a closer understanding of the mutual relevance of institutional and choice-theoretic formulations.

Our strategy is to use computational experiments to ask what patterns of adoption and abandonment would arise if firms acted on the basis of adaptive emulation. We then explore how these dynamics depend upon the effectiveness of competing innovations and the stratification of the market in which a group of businesses compete. The results suggest quite different relationships between competitive effectiveness, mimicry, and isomorphism than are generally assumed within the organizations literature. While the results confirm our initial intuition about the dynamical consequences of a bias toward success stories, we also find that the implications of adaptive emulation are more complex and interesting than simple faddishness.

THE FADDISH CAREER OF THE QUALITY CIRCLE

We begin with a familiar illustration of cyclical dynamics: the career of the quality circle, a prominent organizational innovation of the 1980s (for detailed discussions, see Cole 1989; Strang 1997; Abrahamson and Fairchild 1999).[2] Quality (control) circles were initially developed in Japan to involve foremen and line workers in quality control on an autonomous, team basis. Widely adopted within Japanese industry by the mid-1960s, quality circles were reinterpreted in the United States as voluntary groups that made workplace recommendations to management. They experienced a slow incubation period in American industry (the earliest reported adopters include Smith-Kline Instruments and Lockheed Aerospace). More than a decade later, American corporations using quality circles could still be enumerated (Cole 1979).

Quality circles took off in the early 1980s, due to their perceived role as a "management secret" responsible for Japanese economic success. While in the 1970s circles had spread between corporations doing business with one another, now they were profiled in core business journals, bestselling management books like Ouchi's *Theory Z*, and, most influentially, in a CBS documentary entitled "If Japan Can, Why Can't We?" A 1982 survey found quality circles in 14% of firms with more than 100 employees

---

[2] Although "innovation" can connote "improvement," that is not our usage. Rather, we follow Kimberly (1981) in defining a management innovation as a change in the state of the art of management at the time it is introduced. Rather than assume that new practices represent progress, we take the impact of the innovation on performance as a key parameter and study its impact on adoption trajectories.

and in 43% of manufacturing firms with more than 10,000 employees (Freund and Epstein 1982).

From their base in high-tech manufacturing, quality circles spread to industries like health care, financial services, and even education and public administration. Lawler and Mohrman (1985) reported that virtually every member of the Fortune 500 had given quality circles a try, and quality circles spread down the organizational size hierarchy as smaller firms moved to imitate corporate giants.

Quality circles did not become institutionalized within adopting firms, however. Three years after Lawler and Mohrman (1985) declared in the *Harvard Business Review* that "as everyone knows, it's a fad," Castorina and Wood (1988) reported that 80% of quality circle efforts in the Fortune 500 had been abandoned. Longitudinal research confirms a picture of considerable instability: Drago (1988) found that 23 of 81 Milwaukee-area circle programs failed in two years, and Eaton's (1994) analysis of employee involvement found a failure rate of 20%–40% over three years. Abandonment occurred even among the most highly publicized success stories: Lockheed disbanded its circles in 1979 after their founders left to become consultants, and Honeywell's 600 circles in 1982 were "practically extinct" three years later (Geber 1986).

We are not able to graph measures of quality circle adoption over time. A variety of cross-sectional surveys are available, but these use different definitions and study different populations. Also, the high-level managers who are typically surveyed appear to overestimate program retention. Gray (1993) calculated 60% quality circle penetration from a conventional survey study, but detailed interviews found that only 1 in 8 firms had a quality circle program. Hill (1991) received permission from several firms to study their circles, only to discover that no groups were still in operation. And Eaton (1994) found that managers were less likely to report program termination than union representatives were, even in the same firms.

Quality circle usage can be tracked, however, by following the "carriers"—the consultants who organize circles and train employees in quality circle techniques. Table 1 gives numbers of consultants listed in the relevant national directory who offered quality circle training between 1978 and 1991. Only six time points are available, but the shape of the curve seems clear. Consulting activities first rose dramatically, from two consulting firms with 11 full-time consultants in 1980 to 60 firms employing 469 consultants in 1983. Five years later, two-thirds of these firms had exited the market. Underlining the collapse in the quality circle market was the national association's decision in 1987 to drop the term "quality circle" from its name, as consultants regrouped around total quality management (TQM) and other involvement practices. The number of con-

TABLE 1
QUALITY CIRCLE CONSULTANTS IN THE UNITED STATES.

Year	Consulting Firms	Consulting FTEs
1978 ...	2	
1980 ...	2	11
1983 ...	60	469
1988 ...	21	264
1991 ...	13	91
1994 ...	5	60

SOURCE—Training and Development Organizations Directory (Gale Publishing, J. McLean, ed.), various years.

sultants continued to decline in the 1990s, with only five consulting firms remaining in 1994.

Drawing from the example of quality circles, a stylized picture of a faddish cycle might include the following elements:

1. A potentially extensive incubation period where few firms utilize the innovation.
2. A take-off period where popularity rises explosively.
3. A short period of ascendancy marked by very high levels of innovation usage.
4. A period of rapid decline leading to a low equilibrium level of usage.

While we certainly do not contend that all innovations experience faddish careers, it appears a common and characteristic dynamic. "Quality of work life" programs were highly popular in the mid-1970s, but Goodman (1980) reports a 75% failure rate over 5 years. Mintzberg (1994) discusses the rise and fall of strategic planning, which promptly rose again (see Prahalad and Hamel 1994). Abrahamson and Fairchild (1999) note the wave-like sequencing not only of quality circles but also of "job enrichment," "TQM," and "reengineering," and argue that new innovations are in part precipitated by the collapse of a prior innovation.

The volatility of innovation efforts is also visible in the experience of individual firms. Cole (1999) describes Hewlett-Packard's two-decade movement from quality circles to quality management systems (QMS) to hoshin management to "quality 1 on 1." Strang and Still (1999) discuss how benchmarking in a financial services firm stimulated a host of innovation efforts with short life spans, including corporate culture, benchmarking, and quality initiatives. These studies suggest that managers develop organizational capacities while moving rapidly among the hot innovations of the day, a process that Cole (1999) describes felicitously as "managing quality fads."

Finally, we note that faddish cycles occur in the absence of clear-cut evidence of the effectiveness of the innovation. Modifying organizational

strategies to take advantage of new information technology, refocusing on customer satisfaction, and empowering employees can be seen as the way American business has remade itself since the 1970s. But rhetoric and reality may be tenuously connected, as Hackman and Wageman (1995) argue for the case of TQM. Rather than build in assumptions about the impact of innovative efforts, our strategy is to examine the implications of effectiveness in the context of a formally specified model of organizational adoption and abandonment.

THEORETICAL MOTIVATION

Two lines of argument dominate the literature on diffusion. "Rational accounts" point to the virtues of the innovation as the motive force behind adoption. Evolutionary versions argue that selection forces winnow out less adaptive alternatives; in the more common optimizing story, prescient decision makers appraise alternative practices and make optimal choices. For example, Chandler (1962) treats the multidivisional form as solving the coordination problems faced by large firms operating in multiple markets. Kraatz and Zajac (1996) point to resource needs in explaining why poor, failing liberal arts schools adopt professional programs. Across the diffusion literature, Rogers (1995) notes that cost effectiveness is the most frequently observed correlate of innovation spread.

Contagion forms the major alternative explanation for the spread of innovation, where contact with prior adopters generates further diffusion. Much sociological work traces mechanisms and channels. These include socialization arising within face-to-face relations (Davis 1990; Haunschild 1993), the spread of news via weak ties (Granovetter 1973), competition among structural equivalents (Burt 1987), pressure from powerful exchange partners (Cole 1989), and conformity to popular or normative practice (Tolbert and Zucker 1983; Hannan and Freeman 1988).

Some empirical work links the two logics. Holden (1986) and Conell and Cohn (1995) find that successful hijackings and strikes spark more mimicry than do unsuccessful attempts. Tolbert and Zucker (1983) argue that rationality and contagion operate at different historical stages. Their analysis of civil service reform indicates rationally motivated adoption (where innovation use was matched to local conditions) during an early, embattled period. As civil service procedures became broadly legitimated, however, all sorts of organizations adopted. Westphal, Gulati, and Shortell (1997) contend that early adopters customize practices to their needs, while late adopters conform to standardized practices.

But little real theoretical integration of the two approaches has occurred, and that mainly from a choice-theoretic viewpoint. Banerjee (1992)

PROOF 6

models "herd" dynamics generated by rational decision makers' reliance on rates of adoption as indicators of the private information available to others (also see Bikhchandani, Hirshleifer, and Welch 1992). For example, one both scrutinizes the menu of an untried restaurant and looks to see how many others have decided to eat there. Similarly, Abrahamson and Rosenkopf (1993) treat bandwagons as driven by a combination of internal assessments of efficacy and sensitivity to the adoptions of others. They demonstrate that contagion permits diffusion even when individual assessments are generally low—a result echoing the related work of Granovetter (1978) on threshold models of collective behavior.

Rationality and contagion are key components of diffusion analyses, but when applied independently as abstract principles, their behavioral assumptions are often implausible and their empirical implications restrictive. Most important, neither idea accounts for the abrupt declines in popularity that often follow an initial wave of adoption. We outline these difficulties and then propose an alternative model that integrates insights into both the rationality of adaptation and the institutional basis of contagion.

In our view, standard diffusion arguments miscast the problems faced by intelligent, performance-oriented managers who seek to evaluate innovations. Choice-theoretic models are "overrationalized," treating the merits of an innovation as accessible to internal calculation from first principles. In ambiguous or complex domains, such feats are heroic. Models of contagion, on the other hand, are "underrationalized." They contend that managers pay close attention to what others do, while lacking interest in what happens when they do it. But business discourse focuses intently on performance (Strang 1997), not popularity, and it is here that interpretive accounts of organizational action are on firmest ground.[3]

On the predictive side, rationality and contagion have trouble explaining downturns in diffusion cycles. Consider the dilemma confronting choice-theoretic formulations. If innovations are better mousetraps, why are the hottest new practices quickly discredited and discarded? And if innovations are gimmicks, how did they previously attract cascades of growing excitement and enthusiasm among far-sighted maximizers? Contagion models also have difficulty accounting for collapse. While they

[3] Evolutionary and learning-theoretic versions of the rationality argument are not prone to the two horns of this critique; they allow decision makers to be intendedly rational but not necessarily prescient. Selection pressures, not conscious calculation, do the work of optimizing, allowing "backward-looking" managers to act "as if" they purposefully design effective strategies (Macy 1997). But such models do not explain how practices of dubious utility can reproduce and spread. This article instead proposes an extension of the evolutionary approach that produces convergence in the absence of strong selection pressures.

explain sudden fad-like convergence, contagion models suggest that popular practices will enjoy tremendous staying power, even if the innovation is worthless. But as the case of quality circles illustrates, hot innovations can experience a rapid decline in popularity.

This explanatory difficulty is obscured by conventional research designs. Diffusion studies typically examine the up phase of ultimately popular practices, asking why some adopt early while others lag behind (Rogers 1995; Strang and Soule 1998).[4] They neglect the drop in attention and usage that often follows prominence, and so they fail to consider the explanatory limits of standard accounts.

We seek to avoid the extremes of over- and underrationalized accounts and to better capture the vicissitudes of innovation. To do so, we integrate backward-looking models of organizational decision making with institutional arguments about imitation.

From work on organizational decision making and backward-looking rationality, we take the notion of adaptive agents engaged in problem-driven gradient search (March and Simon 1958; Cyert and March 1963). Managers are assumed unable to assess innovations from first principles but instead seek to learn from the coincidence of innovative strategies and subsequent outcomes. Most important, firms tend to stick with present practice if performance exceeds aspirations, while failure stimulates search for an alternative.

But when firms search, where do they look? As DiMaggio and Powell note, "Organizations tend to model themselves after similar organization in their field that they perceive to be more legitimate or successful" (1983, p. 152).[5] Peters and Waterman's (1982) runaway best-seller *In Search of Excellence* provides a canonical example, where highly successful firms become models to be imitated. Empirical analyses include Fligstein's (1990) discussion of how new conceptions of control diffuse and Haveman's (1993) examination of imitative entry into new markets.

Problem-driven search is generally assumed to lead to the evolution of sensible behavior and individual adaptation, while imitation of successful others is a behavioral foundation of institutional isomorphism (DiMaggio and Powell 1983). Suppose organizations embody both tendencies? We argue that adaptive emulation tends to produce not stable isomorphic

---

[4] To our knowledge, Burns and Wholey's (1993) examination of matrix management stands out in empirically modeling both adoption and abandonment. Greve (1995) examines the abandonment of "easy listening" formats by radio stations.

[5] DiMaggio and Powell ground their argument in the decision-making tradition, contending that emulation is sensible when means-ends relationships are ambiguous. As Mizruchi and Lisa (1999) point out, mimetic isomorphism is the most rationalist and most utilized of the three sources of institutional isomorphism that DiMaggio and Powell (1983) delineate.

PROOF 8

equilibrium, but faddish cycles. Although faddish behavior is often taken as evidence of lemming-like conformity, a very different dynamic may be at work. We suggest that a preoccupation with performance can paradoxically generate waves of adoption of innovations that are worthless, or nearly so, followed by waves of abandonment.

Our model is empirically grounded in the prominence of the "success story" in business communication (Strang 1997). Success stories convey information about the returns to innovation through vivid accounts of action taken by one or a few successful firms. Attention to these stories belies the assumption in contagion models that firms imitate blindly, attending to popularity rather than performance. At the same time, reliance on success stories falls far short of assumptions about effective decision making proposed in models of rational choice. The business press and its readers succumb to confirmation bias, a tendency that social psychologists have observed in most human decision making (Gilovich 1993). Extensive research shows that people tend to favor confirming over disconfirming evidence in evaluating the truth content of a belief.[6]

Cognitive bias is compounded by a social bias in access to good and bad news. Consultants and managers are more likely to broadcast the details of their approach when blessed with success than when plagued by failure. Mass media coverage of organizational innovation (Bums 1982; Strang 1997; Abrahamson and Fairchild 1999) is laudatory, particularly where anecdotal discussion is concerned. For example, Strang (1997) found that 85% of discussions of quality circles were favorable in content and that all but two mentions of specific companies cited positive outcomes.

Success stories dominate business discourse to the virtual exclusion of close theoretical or comparative analyses, particularly in the take-off phase. For example, the case for quality circles depended on their perceived role in Japanese competitive success. Between 1979 and 1981, a flurry of articles appeared in the business press tracing the productivity advantage, dramatic quality gains, and harmonious human relations of Japanese firms to quality circles. Much discussion also rehearsed the benefits of quality circles in American exemplars like Lockheed, reporting cost savings of $3 million, tenfold defect reduction, a 6:1 return on in-

q10

---

[6] Macy and Strang (2000) discuss the cognitive basis of confirmation bias in detail, considering results from the Wason selection task. They present a model of adoption and abandonment that is expanded here to incorporate effects of market stratification and elaborated to better capture variation in managerial decision making.

vestment, and 90% employee satisfaction. During the 1974–81 upswing when quality circles became hot, no articles traced quality circle failures.[7]

Preoccupation with success stories produces a self-reinforcing dynamic. Suppose that every firm uses a different innovation, all of which are entirely worthless, and that all firms have an equal shot at success in a highly competitive market. It is then very unlikely that randomly selected winners will cluster around particular innovations. But it is only a matter of time until, by chance, consecutive winners happen to be using the same innovation. Even if only a few daring firms copy the leaders, the odds have now been tilted slightly toward a third winner using this same innovation. Eventually there will be three consecutive wins, which may turn the heads even of the skeptics. The more repetition, the more converts, and the more converts, the stronger the odds of yet another repetition, followed perhaps by a rush to jump on the bandwagon.

The dynamic generated by this positive feedback is vitiated by a second consideration. Firms evaluate innovations based not only on the experience of highly visible "top performers," but also on their individual experience. As Burns and Wholey comment, "Although the adoption decision is a function of network influences . . . the abandonment decision seems to be based on information peculiar to an institution's direct experience" (1993, p. 133). In line with work on problem-driven search, the orientation to successful others is strong when considering novel alternatives but weak when more immediate, vivid, and relevant evidence is available from personal experience. Firms will abandon a practice associated with success elsewhere if their own experience is disappointing, even if the practice is popular with others.

In short, individual experience can be expected to drive the distribution of worthless strategies toward uniformity (all strategies have equal probability of adoption), while vicarious experience can be expected to periodically drive the distribution toward convergence. These opposing tendencies make model dynamics difficult to calculate in closed form. The implications of adaptive emulation are further complicated when we consider even a minimal parameter set. Diffusion trajectories should be sensitive to the efficacy of the innovations under consideration, an issue often finessed in empirical research. Will adaptive agents converge on innovations that are entirely worthless, and is such convergence stable (observed isomorphism) or unstable (observed faddishness)? What happens if the innovations are worthwhile, improving the competitive position of

---

[7] Eventually, research involving experimental controls and close measurement provided a more balanced assessment. This work was published after the business community had moved away from quality circles, however, and was written and read mainly by academics.

firms? What if the population is stratified, making some firms consistently more likely to succeed than others? And how sensitive is the model to different assumptions about managerial skepticism and inertia?

To find out, we move from verbal accounts to a formal model of firm decision making. We then present the results of computational experiments that test its implications.[8] Computational models are useful in applications to evolutionary systems where interest centers on dynamic rather than equilibrium properties (Axelrod 1998, p. 4; see also Carley and Prietula 1998). Computational models of organizational behavior have been used with impressive success by Carley and Svoboda (1996) and Carroll and Harrison (1998); for a recent collection, see Prietula, Carley, and Gasser (1998). These agent-based models are object oriented (Zeggelink 1994) in that the units are actors, not attributes. As noted by Cohen (1998, p. x), "This is a far more congenial framework within which to express intuitions about organizational processes." Accordingly, we program a population of adaptive agents whose individual attributes are either exogenous (fixed or experimentally manipulated) or consequent upon collective choices. The program is written in Delphi and the source code is available from the authors on request.

COMPUTATIONAL MODEL AND EXPERIMENTAL DESIGN

We model a population of 100 firms that seek to attain valued economic outcomes, such as revenues, profit, return on investment, dividends, stock appreciation, and growth in market share.[9] We abstract these outcomes as a single score that indicates each firm's performance in each round. This score depends on three factors: *market position, choice of innovation,* and *luck.* Market position reflects the firm's location in a stratified population, where some firms possess exogenous structural advantages that

---

[8] Empiricists may question the use of "test" in this context, and more broadly, the external validity of highly abstract computational models. Methodologically, our approach parallels Schelling's (1978) formalization of neighborhood segregation. Shelling modeled a highly stylized world in which people lived on a checkerboard and made decisions about moving and staying without regard to collective outcomes. Nevertheless, his experiments demonstrated that stable patterns of segregation could arise in the absence of segregative intent. Using an equally abstract (and "unrealistic") model, we show that bandwagons can rise and fall in the absence of conformist or distancing intentions.

[9] We experimented with larger numbers of firms (up to 1,000) and found no important effects of population size.

are not affected by the choice of innovation.[10] Innovation decisions and random noise also contribute to outcomes, but unlike market position, these can change from round to round. In the first round, firms select randomly from a pool of innovations. After that, firms either maintain their current innovation or select a new one, as described below.

Formally, outcome $O(0 \leq O \leq 1)$ for firm f at time t is given by:

$$O_{ft}(i) = \alpha K_f + \beta V_i + (1 - \alpha - \beta)\epsilon_{ft} \tag{1}$$

where K_f is the market position of firm f, V_i is the performance value of innovation i, and ϵ is noise (luck). α and β are weights that range from 0 to 1 inclusive, such that $\alpha + \beta \leq 1$, while ϵ, V, and K are standard normal deviates whose range is truncated to the unit interval.

Our experiments manipulate the relative impact of market position, choice of innovation, and luck on firm performance as a weighted combination, where weights sum to unity.[11] We begin by investigating the simplest case, where all innovations are worthless and market positions are equal ($\alpha = \beta = 0$). Here, success depends entirely on luck. We then examine how adoption trajectories change where innovation choice affects outcomes ($\beta > 0$) and where markets are stratified ($\alpha > 0$).

We cross the manipulation of α and β with two decision-making factors. *Inertia* indexes the firm's tendency toward experimentation and change. Volatile firms are quick to try new innovations, even when they are doing well; inert firms are slow to search, even when they are doing badly. *Skepticism* indexes the firm's readiness to draw conclusions from the record of outcomes it observes. Naive firms are quick to infer that today's success story can be replicated. Skeptical firms treat the association between innovation and performance as potentially spurious. They want to see a consistent pattern of outcomes before investing company resources in a new idea. We employ a free parameter $\lambda(0 \leq \lambda \leq 1)$ to calibrate the decision function so that a firm with "average" skepticism ($S = 0.5$) would

---

[10] While real firms that experience poor outcomes may exit the population, we leave the study of selection effects for future research. Selection reduces structural inequality over time, making it difficult to test the consequences of market stratification. Strong selection pressures also increase the relative frequency of innovations used by successful firms, while we are interested in the choices of both winners and losers. In this article, we therefore hold the population constant so as not to confound the effects of abandonment and adoption with those of bankruptcy and start-up.

[11] This approach allows us to manipulate the relative contribution of each distribution while holding their variances constant. A different approach would manipulate the variance of K and V to modulate their impact on performance. However, this would confound relative contributions with changes in the distribution of O.

be highly unlikely to believe the association between innovation and performance based on a single observation.[12]

For computational and theoretical simplicity, we assume discrete time and simultaneous updating. Managers base strategic decisions on information from the previous time period. Alternatively, we could have treated decisions as based on information from the current time period, with the order of decision randomized at each iteration. However, simultaneous updating is not only simpler but also captures the intuition that firms do not have immediate access to information about the innovations and performance of other firms.[13]

The firm's decisions are made in two steps. First, the firm decides whether to abandon its current innovation and search for an alternative. If the firm decides to retain its current innovation, the decision process is over for that firm in that round. If the firm abandons its innovation, it then moves on to the second step, the selection of a new plan.

The abandonment decision.—Firms are treated as selecting innovations purely to improve performance (here ignoring for analytic purposes the many other motives that impel real firms to act). We assume that firms compare current outcomes to an "aspiration level" (Levinthal and March 1981; Lant 1992). Aspirations can be based on the firm's own recent experience ("We may be doing poorly, but are we improving?") or on a fixed standard that reflects the range of observed outcomes ("How are we doing relative to the current state of the art?"). Extensive experimentation showed nearly identical results with fixed or floating aspiration levels or their combination. We therefore report results using a simple model with fixed expectations, in which firms compare outcomes achieved under their current innovation to the best possible performance (a perfect score of 1). The probability that firm f will abandon its present innovation at time t is thus:

$$\text{Prob}(D_{fit}) = [1 - \bar{O}_{ft}(i,M)]C_{fit}. \tag{2}$$

[12] We used $\lambda = 0.25$, but experimentation showed that any midrange value gives similar results. Extreme values undermine the effect of skepticism; as λ approaches zero, skeptical firms demand an unrealistic number of repetitions, and as it approaches one, they demand too few.

[13] As a partial test of the robustness of this approach, we randomly programmed up to half the firms to respond to information from prior time periods, making their responses asynchronous. Collective patterns of adoption and abandonment were robust, although convergence levels diminished the more firms were out of phase.

Here $\bar{O}_{ft}(i,M)$ gives the mean performance of the firm over each use of innovation i in the past M iterations.[14]

Observed outcomes are weighted by C_{fit}, the firm's confidence in its ability to assess the effect of innovation i. Confidence is a function of the amount of experience that f has had with innovation i and f's level of inertia. The greater the inertia, the more evidence is required to persuade the firm to abandon an ineffective innovation. Having invested in an innovation, highly inert firms disregard a few bad outcomes as possible flukes. They require disappointment to cumulate over several periods before concluding that an investment was ill-advised.[15]

To test for a fluke, inert firms weight mean performance by the number of observations on which the evaluation is based. Let $n_t(i)$ indicate the number of trials within recent memory when a firm has used innovation i, while I_f denotes firm f's inertia ($0 \le I \le 1$). We have

$$C_{fit} = 1 - I_f^{\lambda n_t(i)}. \tag{3}$$

If f is volatile ($I_f = 0$), then $C_{fit} = 1$ without regard to the extent of f's experience with i. If f is highly inert ($I_f = 1$), C_{fit} is zero regardless of f's experience with i. For $0 < I_f < 1$, C_{fit} is a concave function of experience.

The adoption decision.—If a firm chooses to drop its current innovation, it then looks to "best practice." In our model, this is the "winning innovation" w used by the firm with the highest score in the last round.[16] We assume that all firms know what innovation was used by the top performer in each round and that they store this information for later recall. If a firm abandons its current practice but elects not to emulate the top performer, it innovates without regard to what other firms are doing. (In natural settings this might correspond to internal research and development.) We assume that such choices are random over the population of possibilities.

The probability that a firm will emulate (rather than draw a new innovation at random) is a function of the firm's skepticism and of what it knows about the track record of the winning innovation. Skeptical firms

[14] M is a population parameter that is constant across all firms. Experiments with $10 \le M \le 100$ showed that longer time horizons have no important effect on results. We therefore used $M = 10$.

[15] This specification of inertia follows the spirit of Hannan and Freeman's (1984) notion of structural inertia as internal adaptations that trail environmental change. In the more constrained context of our simulation experiments, exploration of alternative innovations becomes the firm's only form of adaptation, and experiential lessons replace environmental shifts.

[16] An alternative formulation would make firms aware of a number of successful others, or limit attention to the performance of firms occupying the same market niche or geographic area. We leave the analysis of these complications, including network structures, for future investigation.

assume that the "win" may be a fluke and that what works for others may not work for firms such as itself. The probability that a firm adopts the winning innovation thus multiplies two components: one related to the innovation's "external" success with others (R_{fwt}^{ext}) and the other related to "internal" experience the firm may have had with the winning innovation (R_{fwt}^{int}):

$$\text{Prob}[A_{ft}(w)|D_{fit}] = R_{fwt}^{ext}\, R_{fwt}^{int}. \tag{4}$$

To test for a spurious association with external success, managers check to see how many different firms have succeeded with w. The more firms that have won with w, the higher the probability that it, and not something else, lies behind their success. Take $N_t(w)$ as the number of different firms that have won with w between t and $t - M$. We then have

$$R_{fw}^{ext} = 1 - S_f^{\lambda N_t(w)}. \tag{5}$$

The more skeptical the firm, the higher $N_t(w)$ must be to rule out a fluke. If the firm is not skeptical at all $(S_f = 0)$, $R_{fw}^{ext} = 1$ regardless of past experience.

Skeptical firms also consider their own experience by discounting innovations that they have previously tried and abandoned. The more skeptical the firm and the more recently the innovation was abandoned, the larger the discount. If we take $t^*(w)_f$ as the number of rounds since firm f abandoned w, we have

$$R_{fw}^{int} = \begin{cases} 1 & \text{if } w \text{ has not been abandoned by } f \text{ within memory} \\ 1 - S_f^{\lambda t^*(w)} & \text{if } w \text{ was abandoned by } f \text{ between } t \text{ and } t - M \end{cases}$$

$$\tag{6}$$

RESULTS

Experiment 1: Worthless Innovations

We begin with the simplest case, where outcomes are entirely random and firms are naive and volatile. This scenario resembles some portrayals of managerial fashion: firms capriciously jump between "flavors of the month." We begin with this condition not because it is realistic, but because it establishes a baseline for later experiments that study the effects of nonrandom outcomes among more conservative decision makers.

In this "lottery" condition, firms try to interpret inscrutable patterns of success. Where outcomes are entirely random, all innovations are equally ineffective (or effective) and all firms have the same probability of success. Yesterday's best practice is unlikely to make the headlines tomorrow,

making it difficult to single out a "hot" innovation. Further, even if a bandwagon starts to roll, emulators are likely to jump off since there is a low probability that they will replicate the past success of others.

Computational experiments show that this regime produces faddish cycles. To illustrate, figure 1 charts the history of a world where innovations are entirely worthless ($\beta = 0$), and firms are volatile ($I = 0$), naive ($S = 0$), and equally competitive ($\alpha = 0$). This trial was selected from a set of 100 replications as the most representative, based on two measures that track the most frequently chosen innovations in each iteration (which we refer to as "leading innovations"): *popularity* counts the average number of firms using the leading innovation in each iteration; *turnover* counts the number of changes in the identity of the leading innovation per 100 iterations.[17] We utilize these two measures throughout to provide insight into fad-like diffusion patterns, which we would define informally as a condition where a small number of innovations become highly popular, but for brief periods of time.

Figure 1 sketches the percentage of firms following the six most popular innovations in the trial. The graph shows a characteristic bell-shaped pattern as some innovations rise to prominence rapidly, maintain support for modest periods of time, and then disappear. In their heyday, the most popular are able to gain the support of almost two-thirds of all firms. In some intervals, there are two or three popular innovations in contention, while in others, a lucky innovation enjoys a run of dominance before finally succumbing. The mean duration of the six cascades attaining great popularity was just about 15 iterations. A world of worthless innovations is thus a world where stars enjoy a Warhollian "15 iterations of fame."

Figure 1 offers an important existence proof. *An artificial world of actors preoccupied with performance via success stories is a world of fad-like waves of adoption and abandonment.* This tendency to jump on and off bandwagons is not grounded in lemming-like pressures to conform or the false belief that popularity is a proxy for merit, as in Banerjee's (1992) theory of herd dynamics. On the contrary, firms adopt and abandon innovations based strictly on their observed performance. Nor were fads induced by real but fleeting effects of innovations on outcomes (as in the theory of successively "better mousetraps"). On the contrary, all innovations were entirely void of impact at all times. It is remarkable that, in

---

[17] For the selected trial, mean popularity and turnover were 47.7 and 53.5. Over 100 replications of the experiment, mean popularity was 47.6 and mean turnover was 51.6. The between-trial variances in popularity and turnover were so small (0.33 and 3.63, respectively) that a randomly chosen trial would not differ much from the one illustrated in figure 1.

FIG. 1.—Adoption and abandonment of six leading innovations with $\alpha = \beta = S = I = 0$ (innovations are numbered 1–6). Selected as the most representative of 100 trials, with 200 iterations per trial (for visual clarity, only the last 100 iterations are depicted).

a world of worthless innovations, performance-driven firms converge even temporarily in their choices.

Experiment 2: Reluctant Innovators

In experiment 2, we tested the robustness of the faddish pattern observed above across different assumptions about managerial decision making. Figure 2 indexes the popularity and turnover of leading innovations (the twin dimensions of faddishness defined above) as inertia and skepticism vary between 0 and 1 (our coupled indicators of faddishness defined above) as inertia and skepticism are varied between 0 and 1. On the left, firms are naive or impetuous; on the right, skeptical or inert. The individual graphs are labeled to indicate which factor was manipulated (the factor not manipulated was set to 0). Results are based on 100 replications at each of 11 levels of S and I, in steps of 0.1 from 0 to 1.

Skepticism and inertia strongly dampen cascades of support for worthless innovations. On the extreme left, we see the condition illustrated by figure 1, where more than 60% of firms converge each round on a shifting set of leading innovations. But the popularity of the leading innovation plunges as firms become inert and skeptical. Skepticism has particularly strong and fast-acting effects on bandwagons, since all but the most naive

FIG. 2.—Popularity and turnover in leading innovations, by skepticism and inertia ($\alpha = \beta = 0$; S and I incremented in steps of 10 from 0 to 100, with 100 firms, 200 iterations per trial and 10 trials at each step).

firms demand levels of consistency that are not often met in a world where outcomes are determined by chance. By contrast, inertia has little effect on popularity through most of its range. Only where inertia is near its maximum are firms so slow to change innovations that bandwagons fail to emerge.

Turnover is affected by both skepticism and inertia, but in different directions. Skepticism increases the rate at which new leading innovations emerge, since firms quickly determine that a popular innovation's success is a fluke. Inertia decreases turnover—firms are slow to abandon their current practice, and so collective shifts in relative popularity also occur sluggishly.

Experiments 1 and 2 thus provide mixed evidence concerning the possible sources of organizational fads. Experiment 1 demonstrates that a bias in access to evidence of success (relative to failure) can lead large numbers of firms to converge on worthless innovations. Experiment 2 shows that skepticism and inertia tend to snuff out these faddish cascades. While managers are often caricatured as impetuous and gullible (as in the Dilbert series), it seems more plausible to assume that relatively few fit the cartoon, while a few others are ossified and skeptical, and most are somewhere in-between.

PROOF 18

Experiment 3: Effective Innovations

We next experiment with innovations that possess intrinsic merit, so outcomes are affected by the firm's choice of innovation, as well as by chance. Intrinsically effective innovations are likely to be associated with high levels of performance, and firms adopting effective innovations are likely to see their fortunes improve. This should act to stabilize the cycles we saw in the worthless innovations case.

Figure 3 provides data from an experiment in which intrinsic merit was systematically increased across trials, graphing the same summary statistics shown in figure 2. At the extreme left, we have again the case of worthless innovations; at the extreme right, we have a regime where outcomes are entirely determined by the effectiveness of the innovation, with no element of luck. Skepticism and inertia are normally distributed around expected values of 0.5, a point at which cascades were largely snuffed out in experiment 2.

Both popularity and turnover shift systematically as the impact of innovations on outcomes grows. As the popularity of leading innovations rises, more firms converge on a few innovations. The rate of turnover in leading innovations declines as popular innovations remain in favor for longer periods of time. The results show that fad-like patterns with unstable convergence are possible even when innovations affect performance, assuming modest levels of skepticism and inertia.

Figure 4 indicates how performance rises with β, the parameter that weights the effect of innovations on outcomes. Three summary measures are graphed: mean outcomes across all firms, the optimality frequency (the fraction of rounds where the leading innovation is the optimal innovation),[18] and the optimality ratio (the average value of leading innovations divided by the value of the optimal innovation). Each data point averages results for five trials generated for that parameter set.

The graph demonstrates that adaptive emulation is in fact a fairly good strategy, at least in a simple world where innovations affect all firms in the same way. Mean performance rises in a roughly linear fashion with β, as more firms adopt increasingly "sure bets." The average value of leading innovations goes up more sharply, equaling about 95% of the value of the optimal innovation for β's of 20% or more. If some innovations promote substantially improved performance, adaptive emulation tends to find them.

Nevertheless, adaptive emulation does not produce convergence around *optimal* practices. As has been noted in studies of technological diffusion

q13

[18] The optimal innovation is the innovation with the largest performance value V_i (see eq. [1]) in the trial.

FIG. 3.—Popularity and turnover of leading innovations by value of innovation ($\alpha = 0$, normally distributed skepticism and inertia with $E(S) = E(I) = 0.5$; 100 firms, 200 iterations per trial, and 10 trials at each step, from 0 to 100 by 10).
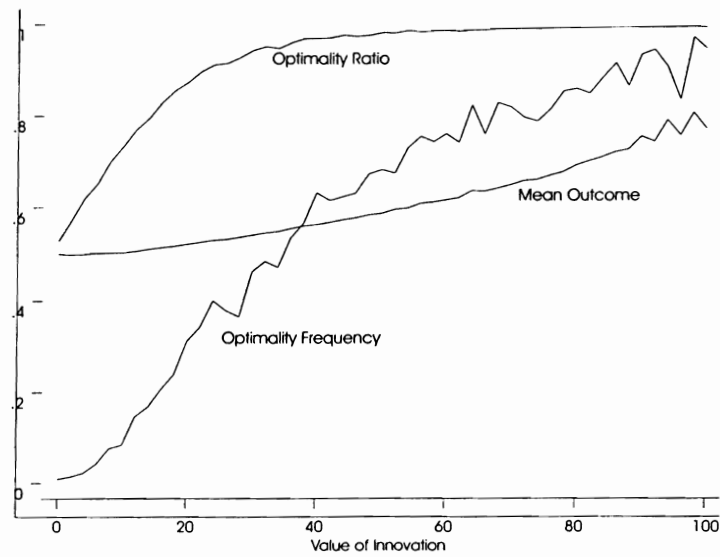

FIG. 4.—Changes in mean outcomes and the optimality of leading innovations as the value of innovation increases from 0 to 100.

PROOF 20

(i.e., VHS versus Beta video formats, QWERTY versus Dvorak keyboards), effective but suboptimal practices that enjoy early popularity can lock out technically superior innovations (David 1985; Arthur 1989). As figure 4 indicates, the frequency with which leading innovations are in fact optimal is rather low and variable, even when innovations are responsible for more than 50% of outcomes.

It seems anomalous that fully deterministic outcomes ($\beta = 100$) do not produce permanent convergence on a single innovation. We thus experimented with decision functions that minimize the probabilistic element in firm decision making. However, such decision functions turn out to be very poor performers, illustrating the trade-off between exploitation and exploration (March 1991). Competency traps (or suboptimal equilibria) are evaded only if firms continue to experiment even when performance is strong.

Most remarkable, however, is how sensitively collective dynamics respond to very modest levels of innovation effectiveness. As figure 3 indicates, popularity and turnover are concave functions of the performance impact of innovations. Most of the surge in popularity and the drop in turnover is concentrated in the region below $\beta = 20\%$. In this range, we see faddish dynamics, that is, innovations have sufficient merit to induce the skeptical and the inert to join the bandwagon but insufficient merit to prevent eventual abandonment and collapse.

As the impact of innovations on performance becomes substantial, it becomes very difficult to dislodge highly popular innovations. An innovation that is both popular and effective tends to "win" so frequently that it is unlikely to ever be supplanted. For example, where β equals 20% or more, we generally find that the majority of firms converge on a single innovation that persists as the dominant choice for more than half the rounds of the trial. When innovations strongly affect outcomes, one sees not fads but institutionalization and, thus, the convergence pattern predicted by standard models of diffusion. Our model incorporates the equilibrium outcome as a special case (highly effective innovations) of a more general process that can also explain the unstable cascades characteristic of fad-like innovation.

It is of further interest that skepticism and inertia become impediments to effective action as the performance impact of innovations increases. Mean performance scores decline by 6% as skepticism and inertia increase from 0 to 0.5, and by another 18% when raised to their maxima of 1. Where popular innovations boost performance, those who hesitate are lost.

We conclude this exploration of the effects of decision-making parameters by noting the robustness of these results. We find faddish behavior for some range of innovative efficacy under every mix of decision-making

parameters. For example, even where skepticism and inertia equal 0.7, fad-like diffusion cycles emerge for levels of β between 5% and 15%. This is a key result. It means that *fads arise regardless of decision-making parameters, if the performance impact of innovations is high enough to generate cascades without being so high as to produce convergence.*

Experiment 4: Effective Firms

Experiment 4 tests the effects of market stratification among firms who differ in their market position. This is a world where Coca-Cola and Microsoft repeatedly outperform competitors due to the comparative advantages they possess, such as a superior product or technology, market leadership based on prior strategic action, or an inspiring corporate culture. Of course, such advantages come from somewhere and are reproduced somehow, but for our purposes, they are exogenous and stable. Most important, the market positions of firms are orthogonal to the innovations whose adoption is being modeled.

Firm-specific market advantages permit innovations to become consistently but spuriously linked to high performance. While innovations adopted by market leaders may have little or no impact on outcomes, outsiders observe that firms using those innovations do rather well. Unlike the condition where outcomes are purely random, an innovation's record of success can be durable (unless the market leader shifts innovations). At the same time, the low rate of turnover among winning firms should make it more difficult to locate effective innovations, even when they are present. These considerations suggest that market stratification may act to moderate the impact of effective innovations.

Figure 5 indicates little effect of market stratification on diffusion dynamics when innovations are worthless. The popularity and turnover of leading innovations is virtually constant across almost the entire range of α, the parameter that weights the effect of market position on outcomes. Only where market position virtually determines outcomes ($\alpha > 80\%$) do we see a sharp decline in the rate at which new bandwagons arise. At this point, the difference between "stars" and "dogs" is so large that dogs find themselves drawn repeatedly to the same empty innovations of the stars.

But substantial effects of market stratification do appear when innovations are effective. Figure 6 demonstrates this relationship for β set to 20. Popularity declines by more than a third as the degree of market advantage shifts from 0% to 40%. Turnover moves in the opposite direction, nearly doubling as the impact of market position rises from 0% to 50%.

This interaction effect implies that market stratification expands the

PROOF 22

FIG. 5.—Popularity and turnover of leading innovations by market stratification ($\beta = 0$, normally distributed skepticism and inertia with $E(S) = E(I) = 0.5$; 100 firms, 200 iterations per trial, and 10 trials at each step, from 0 to 100 by 10).

faddish region across a wider range of innovation effectiveness. Consider the case represented in figure 6, where innovation choice contributes 20% toward performance. If firms are not stratified, this level of β generally produces convergence around a single innovation, but in stratified markets, we see faddish dynamics instead. In fact, high levels of market stratification can preclude convergence up to the point where innovations account for nearly 40% of outcomes.

Stratified markets are more faddish because they produce confusing records of success that make it hard for firms to locate the best innovations. When market position strongly affects outcomes, well-positioned firms are likely to be top performers, even when they choose mediocre innovations. These firms are also more likely to be satisfied with suboptimal innovations, since their market advantage alone ensures such strong performance that they are not motivated to search.[19] As markets become more stratified, we are thus more likely to see "satisficers" at the top whose innovations contribute little to their success.

Repeat winners using suboptimal innovations provide poor guidance

[19] The business press often notes this tendency and reserves its warmest praise for the rare firm that remains hungry and innovative in the face of success.

FIG. 6.—Popularity and turnover of leading innovations by market stratification ($\beta =$ 20, normally distributed skepticism and inertia with $E(S) = E(I) = 0.5$; 100 firms, 200 iterations per trial, and 10 trials at each step, from 0 to 80 by 10).

about what innovations to adopt. In worlds without firm-specific effects, the innovations used by top performing firms are likely to help others as well. That is not true, however, when some firms are "born winners" that send misleading clues about the secrets of success. Now if an average firm emulates a success story, it is liable to find its performance little improved. Emulators are likely to abandon such innovations, reducing their popularity and destabilizing bandwagons.

Even if repeat winners happen to be using highly effective innovations, skeptical observers will have difficulty confirming the causal relationship. Skeptics assume the relationship is a fluke until proven otherwise, and proof requires observation that multiple firms report the same "secret of success." When the same firm always wins, this confirmation is unavailable.

These tendencies are documented in table 2, which presents a regression analysis of market position and innovation effectiveness on average scores per trial. Popularity rises with the value of innovation and declines with market stratification, while turnover moves in the opposite direction. Magnitudes are greater for the effects of innovation than for market position (about four times larger in both cases), but all effects are substantial. There is also a negative interaction between the two factors, so the degree

TABLE 2
COEFFICIENTS FROM LEAST-SQUARES REGRESSION ANALYSES OF POPULARITY AND
TURNOVER IN LEADING INNOVATIONS

	Popularity		Turnover	
Value of innovation	.44	(.003)	−.30	(.006)
Market stratification	.09	(.003)	.07	(.006)
Value of innovation × market stratification ...	−.22	(.011)	−.47	(.20)
B_0	.18	(.001)	.23	(.002)
R^2		.87		.61

NOTE.—SEs are given in parentheses. Data were generated for five trials for every combination of the value of innovation and market stratification from 0 to 100 in steps of 2 (= 6,630 trials). All coefficients are statistically significant at the .001 level.

to which market position diminishes popularity and increases turnover is particularly great at high levels of innovation effectiveness.

The bottom line: *our model predicts faddish dynamics where markets are stratified, and innovations are better than snake oil but far short of "a sure bet."* How much better must innovations be? The more skeptical and inert the population, the greater the intrinsic merit needed to overcome the reluctance of firms to join bandwagons. And the more stratified the market, the wider the range of innovation effectiveness that leads to faddish cycles rather than stable isomorphism.

Summary of Results

Our computational experiments suggest four major results:

Even worthless innovations can spread in a population of performance-driven actors.—Consider the pure lottery in experiment 1, where innovations have no impact on randomly generated outcomes. About half of all firms converge on a single practice in each round. This convergence occurs not because of mindless conformity, but because of the very preoccupation with performance that is commonly regarded as its antidote.

Adaptive emulation generates faddish cycles.—Convergence is often transitory, however. We find that popularity is highly unstable if innovations are worthless or modestly worthwhile, even if firms are highly skeptical and inert. Across a broad range of behavioral and structural conditions, bandwagons collapse as dissatisfaction with performance restores the appetite for new success stories.

Our computational experiments further show that stable isomorphism arises only if innovations strongly affect outcomes. Only then does private experience align with public signals such that innovations used by highly successful firms also tend to work well for emulators. The innovations that become institutionalized are nonoptimal but are effective enough to

build on early success to gain an insuperable advantage over possible competitors.

Fad-like vicissitudes are most robust not when innovations are worthless but when they have identifiable but modest merit.—The relationship between innovation effectiveness and diffusion dynamics is highly non-linear. If innovations are worthless, firms that are skeptical or inert fail to rise to the bait of fleeting success. Further, if innovations largely determine outcomes, popular innovations never lose their hold on the popular imagination, leading to institutionalization, not fads.

Only one condition consistently produces fads across all behavioral assumptions examined here. To our surprise, that situation is one where innovations have modest positive effects on performance, accounting for about 2%–10% of a given outcome. In sum, the most robust (and least expected) implication of adaptive emulation is simply this: fickle enthusiasm and unstable fad-like cascades are most likely to occur not around "snake oil," but around innovations whose impact is modestly positive.

We have made no attempt here to measure empirically the worth of actual innovations that have run through faddish cycles in natural settings, but, a priori, it is an attractive notion that fads arise around good ideas with faint average effects. Innovations like "management by objectives," "matrix management," "reengineering," and "strategic planning" are all sensible practices whose average impact is probably small.

Market stratification promotes faddish cycles by making it more difficult for emulators to recognize effective innovations.—Where exogenous differences in firm performance are strong, innovations associated with success are less likely to be intrinsically effective. This occurs because advantaged firms are likely to outperform competitors even when the innovations they choose are mediocre, and because advantaged firms do not search extensively for the best innovations. They are more likely to present imitators with success stories that turn out to be duds. Conversely, the consistency of the winners may lead skeptics to mistake better mouse-traps for firm-specific advantages. By clouding the benefits of effective innovations, market stratification allows faddish dynamics to persist even when we would otherwise witness stable convergence around worthy innovations.

CONCLUSION

Implications

Theoretical and empirical investigations of innovation diffusion generally reason in two incongruent ways. One line of analysis makes the case for efficiency: innovations are adopted and retained because they provide

"better mousetraps." An alternative argument proposes contagious interactions that lead popularity to breed upon itself. When applied ad hoc, diffusion processes are segregated at the price of explanatory indeterminacy. "Overrationalized" arguments about optimizing performance account for the spread of effective practices, while "underrationalized" arguments about conformity account for the spread of ineffective ones.

The solution is a single model that applies to both cases. We propose a synthesis where boundedly rational actors seek to learn from limited and biased information: that of their own experience and that of successful peers. Building on analyses of organizational decision making, we assume that firms are unable to directly calculate benefits but instead develop innovations experimentally, with consequent asymmetries in adoption and abandonment. Building on institutional accounts, we argue that firms seek to learn from each other and particularly tend to emulate their most successful peers.

This model of adaptive emulation is consistent with DiMaggio and Powell's (1983) discussion of mimetic processes. However, our results turn conventional wisdom on its head. Following DiMaggio and Powell (1983), most students of organizations have presumed that mimetic processes generate stable convergence around one or a few forms and that this effect is independent of competitive and efficiency imperatives. This line of argument helps account for upswings in adoption and their stabilization but tells us little about why bandwagons might collapse.

We show that emulation can generate not only convergence but also fad-like cascades of adoption and abandonment. Instability arises because dissatisfied firms innovate by leaving the pack and charting a new course. If a deviant happens to emerge as a "winner," conceptions of best practice may shift away from popular innovations. By defecting from a common practice to emulate an uncommon (but apparently successful) one, mimics increase population diversity and promote temporal instability.

Moreover, whether convergence is stable or unstable turns primarily on the effectiveness of innovations. When innovations are worthless or nearly so, there is much room for dissatisfaction with outcomes and little reason why deviants should not outperform the competition. Novel success stories routinely arise to erode old bandwagons and start new ones. By contrast, mimetic isomorphism obtains only when the choice of innovation strongly affects outcomes. Emulation then leads an effective practice to gain enough popularity that the sustained development of rival success stories becomes extremely unlikely.

We should emphasize that we make no claim that processes built on emulation always produce the pattern of results found here. Indeed, other formal models consistent with DiMaggio and Powell's (1983) rich discussion of organizational modeling could produce different dynamics. For

example, more stability might arise if best practice were defined by cumulative rather than recent success (see Mezias and Lant 1994). Models of abandonment as a random process, or as inversely related to popularity, might also generate different results.[20]

We also recognize that many social processes besides adaptive emulation can produce cyclic dynamics. New developments may lead today's practices to be replaced by new and better ones—a technological engine of change in some industries. Cycles could emerge because fashion-setters flee their followers. Popularity may overload an initially effective product, causing "traffic jams." Or innovations may spread faster than the skills needed for their effective implementation. Such processes can be formally modeled and empirically tested.[21]

Given these alternatives, our concern is not to identify a unique process that generates fad-like cycles. Rather, we have sought to integrate insights from two major lines of organizational analysis. While institutional and choice-theoretic arguments are generally seen as opposed, we find that their combination predicts novel outcomes and relationships. In particular, we show that fad-like cascades of adoption and abandonment can occur under conditions that neither institutional nor choice-theoretic arguments would have predicted, that is, when decision makers focus on optimizing performance and when the innovations they adopt are entirely worthless or nearly so.

Assuming that managers are exceptionally intelligent and under intense pressure to get it right, and that consultants charge a premium to implement hot innovations, the faddishness of the business community is indeed curious. We show that this faddish behavior may occur not in spite of these performance pressures and high consulting fees but *because* of them. Consultants advertise their winners, not their also-rans. And managers feel pressure from executive boards and stockholders to emulate peer success. Cognitive and social biases toward confirmation of success have powerful self-reinforcing effects, even when managers are careful to test for spurious associations between innovation and performance and even when they pay careful attention to their own experience. The policy implications of our research can be distilled to Montaigne's heuristic: "Look not only at the believers who were saved, but at the even larger number who drowned." Or in more contemporary terms, "Look for best practice

---

[20] We would note, however, that some model of abandonment is needed. If firms never abandon their present practice, change is not possible and an initially differentiated population will remain differentiated forever.

[21] To our knowledge, only the first of these conjectures has been analyzed in detail, following Schumpeter's ideas about waves of technological innovation (see Anderson and Tushman 1990).

not only among the Intels and Microsofts, but also among the Wangs and Digitals."

The trend in the contemporary business community, however, is in the opposite direction. Increasingly, competitive benchmarking has been adopted as a mode of corporate planning and innovation, where firms learn from the practices of the "best in class." And increasingly, organizational analysis is externalized in consultants and the business media, creating much awareness of what other firms do but in a potentially biased fashion. The result, we think, will be a continuing amplification of faddish cycles as firms adopt and abandon "hot" innovations.

Limitations and Applications

To investigate the implications of adaptive emulation, we developed a formal model that abstracts away from many complexities. There are thus many points where the argument proposed here could be scrutinized and elaborated. We briefly enumerate some theoretical extensions and empirical tests of the model that seem promising:

Relax core assumptions of our specification.—In this article, we assume that differences in the value of innovations are randomly distributed across the population, while in reality, specific innovations fit some firms better than others, and first- (and sometimes last-) movers often have an advantage. Likewise, our simulation experiments assume that all firms attend to the same exemplars, slighting the networks that much research q14 has shown to channel diffusion (i.e., Davis 1992; Haunschild 1993). Further analysis could examine the impact of heterogeneity in diffusion channels and firm-innovation matches.

Our formalization also employs stylized assumptions about the mechanics of adoption and abandonment. These include the notion that firms first abandon practices and then search for alternatives, that rival innovations are adopted sequentially rather than simultaneously, and that the sole costs of adopting innovations are foregone opportunities to innovate in other directions. (Such costs include not only implementation expenses but also loss of morale when firms move rapidly among "flavors of the month.") The model could be enhanced to permit more flexible patterns of action by firms and to incorporate costs as well as benefits.

Elaborate the model in fundamental ways.—Our discussion of confirmation bias pointed to the tendency of management consultants to publicize their successes and avoid discussing clients who stumbled. However, the formal model captures this tendency only as it is manifested in managerial decision making. An elaborated model might explore the strategic interdependence of management consultants and their clients by modeling the intersection of two decision processes, where firms consider the wares

that consultants offer and consultants seek to anticipate what firms will want. This game-theoretic specification shifts the locus of confirmation bias from cognitive error to social interaction.

A second fundamental extension would enrich the cognitive capacities of the actors. The model of adaptive emulation developed here rests on a reinforcement version of learning, where firms respond to failure by shifting innovations and emulating top performers. A more complex approach would allow firms to devise alternative decision-making strategies (perhaps generated randomly and spreading through observation or random contact) and to "learn to learn" by revising these models. An evolutionary version of decision making might have substantial implications for findings presented here, perhaps reinforcing adaptive emulation in worlds of effective innovations but extinguishing it where innovations are worthless. While such extensions go far beyond the framework of this article, they suggest some of the potential power of formal models of organizational diffusion and learning.

Test the assumptions and implications of the model.—At the level of the individual firm, one can ask how managers assess innovations. In this article, we have built on a vigorous line of research into problem-driven search and aspiration levels (Cyert and March 1963; Lant 1992) and the prominence of success stories in organizational discourse (Burns and Wholey 1993; Strang 1997). But the formalization of these ideas highlights relatively unexplored issues for diffusion research. What biases in attention to success and failure are evident in managerial decision making? How do managers weigh information about outcomes versus information about popularity? What sorts of firms serve as models to be emulated?

At the level of a single practice spreading across firms, two considerations are suggested by the analysis here. First, we should measure not only whether relevant others have adopted, but what outcomes these others have experienced. Diffusion analysis is too often limited by the assumption that all social forces conspire to promote further spread, neglecting the sobering lessons of failure. Inattention to these effects are theoretically as well as empirically limiting. As Strang and Soule (1998, p. 269) note, "The inability to specify what is observed [by potential adopters] produces some theoretical fuzziness about the microprocesses involved in diffusion."

We also see the need for research on innovation abandonment. It is important to develop a theory of institutionalization within the firm that recognizes the fragility of efforts at change (Walton 1975; Zeitz, Mittal, and McAuley 1999). While institutional accounts have helped explain how nonoptimal organizational forms gain currency, we also need to understand how good ideas are discredited. Adaptive emulation provides one

PROOF 30

insight into this problem, demonstrating that direct experience is not a sufficiently strong guide in a world dominated by success stories.

Although conceptually straightforward, the study of abandonment is methodologically challenging. Abandonment is often a nonevent, where firms cease to supply the symbolic and material support that organizational change requires. It is much more difficult to observe than is adoption, which can often be dated through associational records or conventional organizational surveys. More intensive methodologies are needed, like Easton and Jarrell's (1998) construction of organizational innovation histories.

The key implications of adaptive emulation, however, are best examined in a comparative or population-level analysis of multiple innovations. Simulation results suggest that faddishness is powerfully and nonmonotonically affected by innovation effectiveness, with fads most likely to arise around modestly useful ideas. And they suggest an interaction effect between market stratification and innovation effectiveness, where more stratified markets make it more likely that highly effective innovations will lead to cascades and crashes rather than convergence around a "best practice." These implications are sufficiently distinctive that comparative analyses could go far in testing the model proposed here.

Examine the relevance of adaptive emulation in other settings.—A compelling example is provided by cultural genres, where a single hit often generates a wave of imitation. For example, "Seinfeld" attracted a swarm of New York-based sitcoms; "Who Wants to be a Millionaire" revived the game show; and "Whassup" has quickly entered the video lexicon of the advertising industry. Moreover, all three of these examples may have disappeared from the cultural landscape by the time you read this.

Faddish dynamics also seem rife in the financial (as opposed to the managerial) world. "Sector rotation" and "momentum plays" among equity traders may manifest the self-reinforcing dynamics of adaptive emulation, aided in part by the financial media. When chip and Internet stocks are soaring, investors and journalists swear allegiance to the "New Economy," but just as quickly scorn the naivete of this belief when the NASDAQ heads south of the Dow.

Innovation in areas like health and education also suggest adaptive emulation. A remarkable recent example involves efforts to clone KIPP Academy, a pair of public schools in poor neighborhoods whose students have achieved high scores on standardized tests (Wilgoren 2000). In a political climate intent on discovering new models of schooling, the approach taken in these schools has garnered much attention, most widely at the Republican National Convention (where students chanted "read, baby, read"). The owners of the Gap have now given $15 million to stimulate the construction of hundreds of KIPP academies around the

country. Meanwhile, a debate continues over whether KIPP's success is replicable or based on local conditions or a few dynamic leaders.

In the broadest sense, asymmetric attention to success and failure is evident in many social processes. For example, career choices are shaped in part by the successes of others, a phenomenon one sees in aggregate form today in the salience of the Internet start-up and the emulation of superstar athletes. The latter shows how confirmation bias may have tragic consequences for occupational aspirations if abandonment does not arrive until it is too late to salvage a viable career. On the demand side, we observe the feeding frenzy that often follows an expression of interest by a leading firm (or university) in hiring a rising star.

We do not want to suggest, of course, that adaptive emulation is universal. Abstractly, three scope conditions may limit its applicability. First, actors must be highly performance oriented. A preoccupation with success is likely to arise in sharply competitive settings where personal vulnerability leads to collective anxiety (Smelser 1962). Second, causal relationships must be sufficiently ambiguous that actors learn from previous outcomes rather than calculating them a priori (DiMaggio and Powell 1983). This "backward-looking" reliance on propinquity as a proxy for causality permits actors to attribute success to innovations that may have made little or no contribution (see Macy and Flache [1995] for an overview of recent studies using models of backward-looking rationality). Third, actors must be positioned to identify the same "secrets of success." This is difficult if top performers conceal their success or their practices, if communication networks are fragmented, and if the mass media do not broadcast success stories.

The argument developed here thus suggests an agenda for institutional analysis. In this article, we have implicitly criticized what Powell (1991) called the "institutionalization of institutional analysis"—the theoretical segregation of technical and institutional forces and a focus on conformity to the latter. We would argue instead for inquiry into the social construction of success. Institutional accounts can help explain why some high performers become exemplars while others are overlooked, and why some practices are interpreted as transferable solutions while others are ignored. The formal model of adaptive emulation examined here does not ask these questions, assuming for simplicity that the spotlight shines on the highest performing firm and the unique innovation it utilizes. The real world is not so transparent, however. Success stories are not raw data but cultural constructions, formulated in terms of the problems that actors perceive and modes of action they find comprehensible. We thus hope readers will look beyond the limits of our formalization to consider the broader theoretical agenda provoked by the integration of choice-theoretic and institutional ideas.

q15

REFERENCES

q16   Abrahamson, Eric. 1996. "Management Fashion." *Academy of Management Review* 21:1254–85.

Abrahamson, Eric, and Gregory Fairchild. 1999. "Management Fashion: Lifecycles, Triggers, and Collective Learning Processes." *Administrative Science Quarterly* 44: 708–40.

Abrahamson, Eric, and Lori Rosenkopf. 1993. "Institutional and Competitive Bandwagons: Using Mathematical Modeling as a Tool to Explore Innovation Diffusion." *Academy of Management Review* 18:487–517.

Anderson, Philip, and Michael L. Tushman. 1990. "Technological Discontinuities and Dominant Designs: A Cyclical Model of Technological Change." *Administrative Science Quarterly* 35:604–33.

Arthur, W. Brian. 1989. "On Competing Technologies and Lock-in by Historical Events: The Dynamics of Allocation under Increasing Returns." *Economic Journal* 99: 116–31.

Axelrod, Robert. 1998. *The Complexity of Cooperation*. Princeton N.J.: Princeton University Press.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107:797–817.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100:992–1026.

Burns, Lawton R., and Douglas R. Wholey. 1993. "Adoption and Abandonment of Matrix Management Programs: Effects of Organizational Characteristics and Interorganizational Networks." *Academy of Management Journal* 36:106–38.

Burt, Ronald S. 1987. "Social Contagion and Innovation: Cohesion versus Structural Equivalence." *American Journal of Sociology* 92:1287–1335.

Carley, Kathleen, and M. Prietula. 1998. "Computational Approach to Organizations and Organizing." In *Simulating Organizations: Computational Models of Organizations and Organizing,* edited by M. Prietula, K. Carley, and L. Gasser. Cambridge Mass.: MIT Press.

Carley, Kathleen, and D. Svoboda. 1996. "Modeling Organizational Adaptation as Simulated Annealing Process." *Sociological Methods and Research* 25:3–30.

Carroll, Glenn R., and J. Richard Harrison. 1998. "Organizational Demography and Culture: Insights from a Formal Model and Simulation." *Administrative Science Quarterly* 43:637–67.

Castorina, Paul, and Brian Wood. 1988. "Circles in the Fortune 500: Why Circles Fail." *Journal for Quality and Participation* 11:40–41.

Chandler, Alfred D., Jr. 1962. *Strategy and Structure*. Cambridge, Mass.: MIT Press.

Cohen, Michael. 1998. "Foreword." In *Simulating Organizations: Computational Models of Institutions and Groups*, edited by M. Prietula, K. Carley, and L. Gasser. Cambridge Mass.: MIT Press.

Cole, Robert E. 1979. "Made in Japan-Quality Control Circles." *Across the Board* 16: 72–78.

———. 1989. *Strategies for Learning*. Berkeley: University of California Press.

———. 1999. *Managing Quality Fads: How American Business Learned to Play the Quality Game*. New York: Oxford University Press.

Conell, Carol, and Samuel Cohn. 1995. "Learning from Other People's Actions: Environmental Variation and Diffusion in French Coal Mining Strikes, 1890–1935." *American Journal of Sociology* 101:366–403.

Cyert, Richard M., and James G. March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs, N.J.: Prentice-Hall.

American Journal of Sociology

David, Paul A. 1985. "Clio and the Economics of QWERTY."*American Economic Review* 75:32–37.

q17 Davis, Gerald F. 1991. "Agents without Principles? The Spread of the Poison Pill through the Intercorporate Network." *Administrative Science Quarterly* 36:583–613.

q18 Davis, Gerald F., and Henrich R. Greve. 1997. "Corporate Elite Networks and Governance Changes in the 1980s." *American Journal of Sociology* 103:1–37.

DiMaggio, Paul J., and Walter W. Powell. 1983. "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields." *American Sociological Review* 48:147–60.

Drago, Robert. 1988. "Quality Circle Survival: An Exploratory Analysis." *Industrial Relations* 27:336–51.

Easton, George, and Sherry Jarrell. 1998. "The Effects of Total Quality Management on Corporate Performance: An Empirical Investigation." *Journal of Business* 71: 253–307.

Eaton, Adrienne E. 1994. "The Survival of Employee Participation Programs in Unionized Settings." *Industrial and Labor Relations Review* 47:371–89.

Freund, William C., and Eugene Epstein. 1982. *People and Productivity.* Homewood, Ill.: Dow Jones-Irwin.

Geber, Beverly. 1986. "Quality Circles: The Second Generation." *Training* 23:54–61.

Gilovich, Thomas. 1993. *How We Know What Ain't So: The Fallibility of Human Reason in Everyday Life.* New York: Free Press.

Goodman, Paul S. 1980. "Quality of Work Lie Projects in the 1980s." *Labor Law Journal* 487–94.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78:1360–80.

———. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83:1420–43.

Gray, George. 1993. "Quality Circles: An Update." *SAM Advanced Management Journal* 58:41–47.

Greve, Henrich R. 1995. "Jumping Ship: The Diffusion of Strategy Abandonment." *Administrative Science Quarterly* 40:444–73.

Hackman, J. Richard, and Ruth Wageman. 1995. "Total Quality Management: Empirical, Conceptual, and Practical Issues." *Administrative Science Quarterly* 40: 309–42.

Haunschild, Pamela R. 1993. "Interorganizational Imitation: The Impact of Interlocks on Corporate Acquisition Activity." *Administrative Science Quarterly* 38:564–92.

Haveman, Heather A. 1993. "Follow the Leader: Mimetic Isomorphism and Entry into New Markets." *Administrative Science Quarterly* 38:593–627.

Hill, Stephen. 1991. "Why Quality Circles Failed but Total Quality Management Might Succeed." *British Journal of Industrial Relations* 29:541–68.

Holden, Robert T. 1986. "The Contagiousness of Aircraft Hijacking." *American Journal of Sociology* 91:874–904.

q19 Imberman, Woodruff. 1982. "Common Causes of QCC Failures." *Quality* 21:30–31.

Kimberly, John R. 1981. "Managerial Innovation." Pp. 84–104 in *Handbook of Organizational Design*, edited by W. H. Starbuck and P. C. Nystrom. New York: Oxford University Press.

Kraatz, Matthew S., and Edward J. Zajac. 1996. "Exploring the Limits of the New Institutionalism: The Causes and Consequences of Illegitimate Organizational Change." *American Sociological Review* 61:812–36.

Lant, Theresa K. 1992. "Aspiration Level Adaptation: An Empirical Exploration." *Management Science* 38:623–44.

q20 Lawler, Edward E., III, Gerald E. Ledford, and Susan A. Mohrman. 1987. *Employment Involvement in America: A Study of Contemporary Practice.* Houston: American Productivity and Quality Center.

Lawler, Edward E., III, and Susan A. Mohrman. 1985. "Quality Circles after the Fad." *Harvard Business Review* 65–71.

Levinthal, Daniel, and James G. March. 1981. "A Model of Adaptive Organizational Search." *Journal of Economic Behavior and Organization* 2:307–33.

Macy, Michael W. 1997. "Identity, Interest and Emergent Rationality." *Rationality and Society* 9:427–48.

Macy, Michael W., and Andreas Flache. 1995. "Beyond Rationality in Theories of Choice." *Annual Review of Sociology* 21:73–91.

Macy, Michael W., and David Strang. 2000. "Pluralistic Ignorance and the *Top Secret Management Handbook*: A Computational Model of Fashionable Innovation." In *Simulating Organizational Societies: Theories, Models and Ideas*, edited by Alessandro Lomi and Eric Larsen. Cambridge Mass.: MIT Press.

March, James G. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2:71–87.

March, James G., and Herbert A. Simon. 1958. *Organizations*. New York: Wiley.

q21 March, James G., Lee S. Sproull, and Michal Tamuz. 1991. "Learning from Samples of One or Fewer." *Organization Science* 2:1–13.

q22 Meyer, John W., and Brian Rowan. 1977. "Institutionalized Organizations: Formal Structure as Myth and Ceremony." *American Journal of Sociology* 83:440–63.

Mezias, Stephen J., and Theresa Lant. 1994. "Mimetic Learning and the Evolution of Organizational Populations." Pp. 174–98 in *Evolutionary Dynamics of Organizations*, edited by Joel A. C. Baum and Jitendra V. Singh. New York: Oxford University Press.

q23 Micklethwait, John, and Adrian Wooldridge. 1996. *The Witch Doctors*. New York: Random.

Mintzberg, Henry S. 1994. *The Rise and Fall of Strategic Planning*. New York: Free Press.

Mizruchi, Mark S., and Lisa C. Fein. 1999. "The Social Construction of Organizational Knowledge: A Study of the Uses of Coercive, Mimetic, and Normative Isomorphism." *Administrative Science Quarterly* 44:653–83.

q24 Osterman, Paul. 1994. "How Common Is Workplace Transformation and Who Adopts It?" *Industrial and Labor Relations Review* 47:173–88.

Peters, Thomas J., and Waterman, Robert H., Jr. 1982. *In Search of Excellence*. New York: Harper and Row.

Powell, Walter W. 1991. "Expanding the Scope of Institutional Analysis." Pp. 183–203 in *The New Institutionalism in Organizational Analysis*, edited by Walter W. Powell and Paul J. DiMaggio. Chicago: University of Chicago Press.

Prahalad, C. K., and Gary Hamel. 1994. *Competing for the Future*. Boston: Harvard Business School Press.

Prietula, M., K. Carley, and L. Gasser. 1998. *Computational Models of Institutions and Groups*. Cambridge Mass.: MIT Press.

Rogers, Everett M. 1995. *Diffusion of Innovations*. New York: Free Press.

Shelling, Thomas C. 1978. *Micromotives and Macrobehavior*. New York: W.W. Norton.

Smelser, Neil J. 1962. *Theory of Collective Behavior*. New York: Free Press.

Strang, David. 1997. "Cheap Talk: Managerial Discourse on Quality Circles as an Organizational Innovation." Presented at the annual meetings of the American Sociological Association, Toronto.

q25 Strang, David, and John W. Meyer. 1993. "Institutional Conditions for Diffusion." *Theory and Society* 22:487–512.

Strang, David, and Sarah A. Soule. 1998. "Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills." *Annual Review of Sociology* 24: 265–90.

Strang, David, and Mary C. Still. 1999. "Influence and Innovation within the Business

Community: Who Does an Elite Organization Emulate?" Presented at the annual meetings of the American Sociological Association, Chicago.

Tolbert, Pamela S., and Lynne G. Zucker. 1983. "Institutional Sources of Change in the Formal Structure of Organizations: The Diffusion of Civil Service Reform." *Administrative Science Quarterly* 28:22–39.

Walton, Richard E. 1975. "The Diffusion of New Work Structures: Explaining Why Success Didn't Take." *Organizational Behavior* 3–21.

Westphal, James D., Ranjay Gulati, and Stephen M. Shortell. 1997. "Customization or Conformity: An Institutional and Network Perspectcive on the Content and Consequences of TQM Adoption." *Administrative Science Quarterly* 42:366–94.

Wilgoren, Jodi. 2000. "Seeking to Clone Schools of Success for Poor." *New York Times*, p. 1. Wednesday, August 16.

Zeggelink, Evelyn. 1994. "Dynamics of Structure: An Individual Oriented Approach." *Social Networks* 16:295–333.

Zeitz, Gerald, Vikos Mittal, and Brian McAuley. 1999. "Distinguishing Adoption and Entrenchment of Management Practices: A Framework for Analysis." *Organizational Studies* 20:741–76.