

INVITED RESEARCH ARTICLE

Going Beyond the Dominant Paradigm for Information Technology Innovation Research: Emerging Concepts and Methods*

Robert G. Fichman

Wallace E. Carroll School of Management
Boston College
fichman@bc.edu

Abstract

Research on information technology (IT) innovation is concerned with identifying the factors that facilitate or hinder the adoption and diffusion of new IT-based processes or products. Most of this research has been conducted within the confines of a dominant paradigm wherein innovations are assumed to be beneficial, and organizations that have greater innovation-related needs and abilities are expected to exhibit a greater amount of innovative activity. This essay suggests that the dominant paradigm may be reaching the point of diminishing returns as a framework for supporting ground-breaking research, and urges researchers to adopt a more innovative approach to the study of IT innovation itself. Toward this end, I present seven opportunities for conducting new kinds of research that go beyond the dominant paradigm.

Keywords: IT innovation, IT diffusion, IT adoption, innovation research methods, innovation theory

Introduction

The IT innovation field is concerned with understanding the factors that facilitate or inhibit the adoption and diffusion of emerging IT-based processes or products within a population of potential adopters (Fichman, 2000; Swanson, 1994). These adopters can be individuals (Brancheau and Wetherbe, 1990), organizational units (Cool *et al.*, 1997), firms (Cooper and Zmud, 1990), or even groups of inter-related firms (Damsgaard and

* Sirkka Jarvenpaa was the accepting senior editor for this paper. Robert Zmud and Robert Kauffman were reviewers for this paper.

Lyttinen, 2001). Research on IT innovation has become increasingly popular as IT has continued its relentless march into almost every aspect of organizational life, and as innovation itself has taken on greater importance as a driver of organizational competitiveness (Hamel, 1998).

The majority of prior research on IT innovation, and indeed on organizational innovation in general, has been done within what I will call the *dominant paradigm*. This paradigm is typified by the desire to explain innovation using economic-rationalistic models, whereby organizations that have a greater *quantity* of what might be called “the Right Stuff” (i.e., greater innovation-related *needs* and *abilities*) are expected to exhibit a greater *quantity* of innovation (i.e., greater frequency, earliness, or extent of adoption).¹ The ultimate goal of innovation research is to provide guidance to managers on the question of “*whether, when, and how to innovate with IT*” (Swanson and Ramiller, 2004). This research has shown—to name a few of the dozens of factors that have been studied—that IT organizations that are larger, more diverse, have greater technical expertise, possess supportive senior management, operate in more competitive contexts, and perceive the innovation as more beneficial and compatible, are more likely to adopt a larger number of innovations, to adopt them earlier, and to implement them more thoroughly.

The dominant paradigm has yielded tremendous insights on the subject of how potential adopters can more effectively evaluate innovations and manage the process of assimilating them. It has also provided valuable guidance to vendors, government agencies, consultants, and other kinds of *propagating institutions* (Eveland and Tornatzky, 1990) in their efforts to promote the diffusion of innovations. Nevertheless, the dominant paradigm may be reaching the point of diminishing returns in providing additional opportunities for highly influential future research, for two reasons. First, while the 1990s produced a burst of novel theoretical approaches from within the dominant paradigm, the appearance of additional groundbreaking work from within this paradigm seems to be abating. (See Fichman (2000) for a review of these novel theoretical approaches.) Second, the shear quantity of research in this area has brought us to a point where the broad elements of how managers can promote effective innovation according to the dominant paradigm are now fairly well understood. Just as in industry, where occasional radical innovation is required to move beyond the diminishing returns associated with incremental elaborations on past breakthroughs, IT innovation researchers must also occasionally question fundamental assumptions and seek out more radical concepts and methods to reinvigorate the field.

As a result, this paper addresses the question of how future work on IT innovation can step outside the dominant paradigm in some fundamental way. Toward this end, I will begin by presenting a sketch of the dominant paradigm that identifies key independent and dependent variables, and the nature of the relationships among them. I will also

¹Abrahamson, in a similar line of argument, notes that diffusion research has been dominated by the efficient-choice perspective, which assumes “that agents, usually organizations or their top management teams, have little uncertainty about (a) their preferences or goals, be they profit maximization, market share growth, competitive advantage, or any other strategic preference, (b) innovations’ technical efficiency measured as the ratio of outputs to inputs. Therefore, given existing resource constraints, agents rationally choose the innovation that will allow them to most efficiently produce the outputs that are useful for obtaining their goals” (Abrahamson 1991, p 592).

offer a brief defense of this model. Then, I will survey several promising opportunities to go outside the dominant paradigm. To structure the discussion I will classify these opportunities as primarily affecting the choice of independent variables, as affecting the relationships between independent and dependent variables, or as affecting the choice of dependent variables. Finally, I will offer some general prescriptions for conducting future research on IT innovation both within and beyond the dominant paradigm.

It should be noted that although I'll be describing ways to go beyond the dominant innovation paradigm, I will remain firmly ensconced within the positivist paradigm more generally; there are, of course, other notable innovation threads outside the positivist paradigm, including structuration (Orlikowski, 1992), adaptive structuration (DeSanctis and Poole, 1994), and socio-technical approaches (Bijker, 1995). These streams, and other non-positivist approaches that may emerge are certainly in keeping with the goal of breaking with the dominant innovation paradigm, but are beyond the scope of the present essay.

The Dominant Paradigm in IT Innovation Research

Under the dominant paradigm, it is assumed that organizations with a greater quantity of the "Right Stuff" will exhibit a greater quantity of IT innovation (see Figure 1). Such organizations may also be said to fit the *innovator profile*. The underlying logic is that firms that fit this profile will have greater returns to innovation because they can innovate more easily, economically, or effectively—or because they have a greater opportunity to leverage the benefits the innovation provides. Representative examples of empirical work within this stream include Armstrong and Sambamurthy (1999), Cooper and Zmud (1990), Fichman and Kemerer (1997), Grover *et al.* (1997), Iacobou *et al.* (1995), Premkumar *et al.* (1994), Purvis *et al.* (2001), and Rai and Bajwa (1997). In addition, surveys of prior work on IT innovation can be found in Fichman (2000), Gallivan (2001), Prescott and Conger (1995), and Swanson (1994). Other surveys of organizational innovation more generally are provided in Damanpour (1991), Downs and Mohr (1976), Rogers (1995), Tornatzky and Klein (1982) and Wolfe (1994). While there are, of course, many different types of IT innovations, the central logic portrayed in Figure 1 has been assumed to hold across types.² That is, while the specific variables theorized to constitute the "Right Stuff" may vary across type, the assumption of a positive link between the "Right Stuff" and innovation does not.

The quantity of the "Right Stuff" has been conceptualized as the extent to which organizations possess certain characteristics—or operate in certain contexts—that increase the *need* for innovation and/or the *ability* to innovate successfully. The quantity of innovation has been conceptualized as the extent to which an organization adopts innovations often, adopts them early, and/or adopts them thoroughly (Fichman, 2001). Sometimes the focus is on a particular innovation or class of innovations, with the goal of identifying the factors that affect the adoption and diffusion of that kind of innovation. Other times the focus is on organizations, with the intention of identifying general characteristics that predispose firms to innovate across a broad range of technologies.

² The general innovation literatures has used many different typologies to distinguish different types of innovation, including: process versus product, administrative versus technical, radical versus incremental, and disruptive versus sustaining. In the IT literature, Swanson (1994) has introduced a typology distinguishing six types of IT innovations.

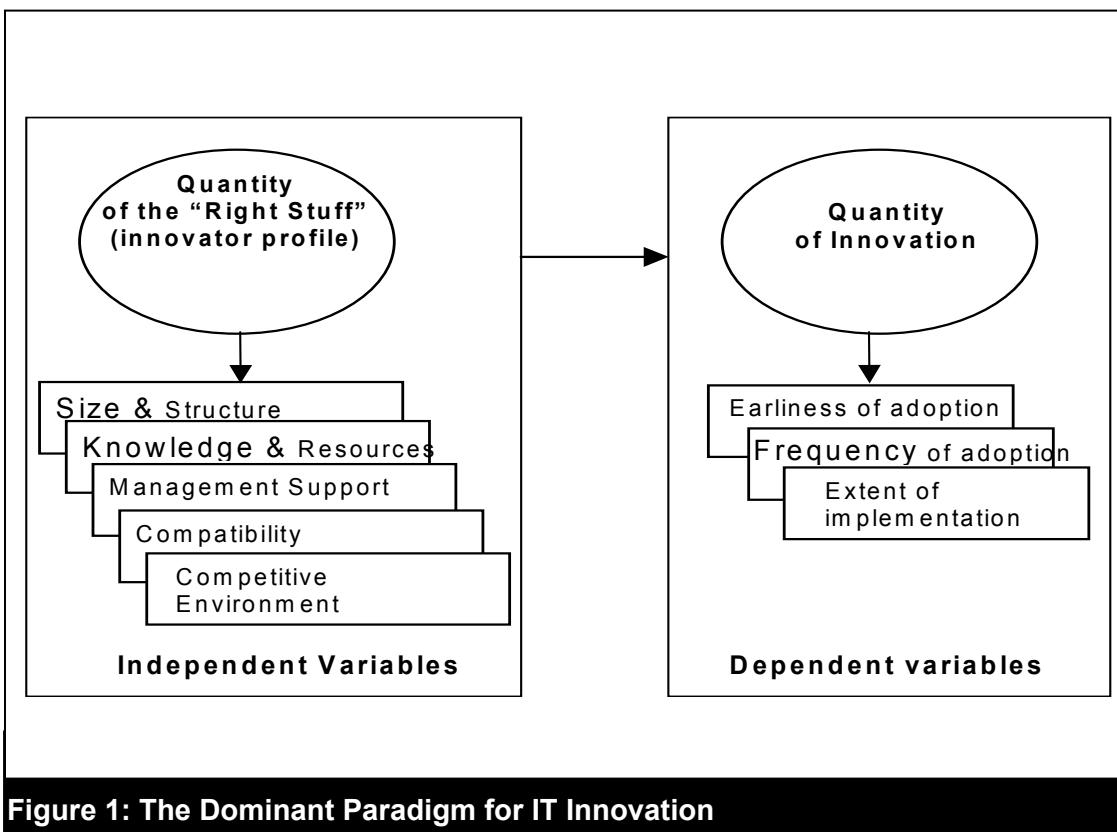


Figure 1: The Dominant Paradigm for IT Innovation

Regardless of whether the focus has been on innovations and what makes them more adoptable or on organizations and what makes them more innovative, researchers have usually employed an *economic-rationalistic* logic to link variables. The economic-rationalistic label reflects a focus on factors that affect the *economic* returns to innovation, and the assumption that managers take these factors into account in a normatively *rational* way in their innovation decisions.

The ultimate outcomes or benefits of innovation with IT are rarely considered in studies within the dominant paradigm. No doubt the difficulty of measuring impacts has played a role in this, however another important reason is the well known *pro-innovation bias*, which refers to the assumption that innovations are beneficial (Kimberly, 1981). If innovations are good, then more innovation is better, and thus the focus on quantity of innovation *per se* is more easily justified.

A Brief Defense of the Dominant Paradigm

While this paper will critique certain elements of research within the dominant paradigm, it is worth pausing to provide a defense of this work. One can begin by observing that the pro-innovation bias does have a basis in reality. Taken in aggregate, technological innovation is not just beneficial, but in fact essential to the long-term health of organizations, industries and societies (Clark and Guy, 1998; Nadler and Tushman, 1999). It has been estimated that technical progress accounts for a larger share of productivity growth than any other single factor, including investments in capital equipment (Solow, 1957). Furthermore, a large number studies have demonstrated the

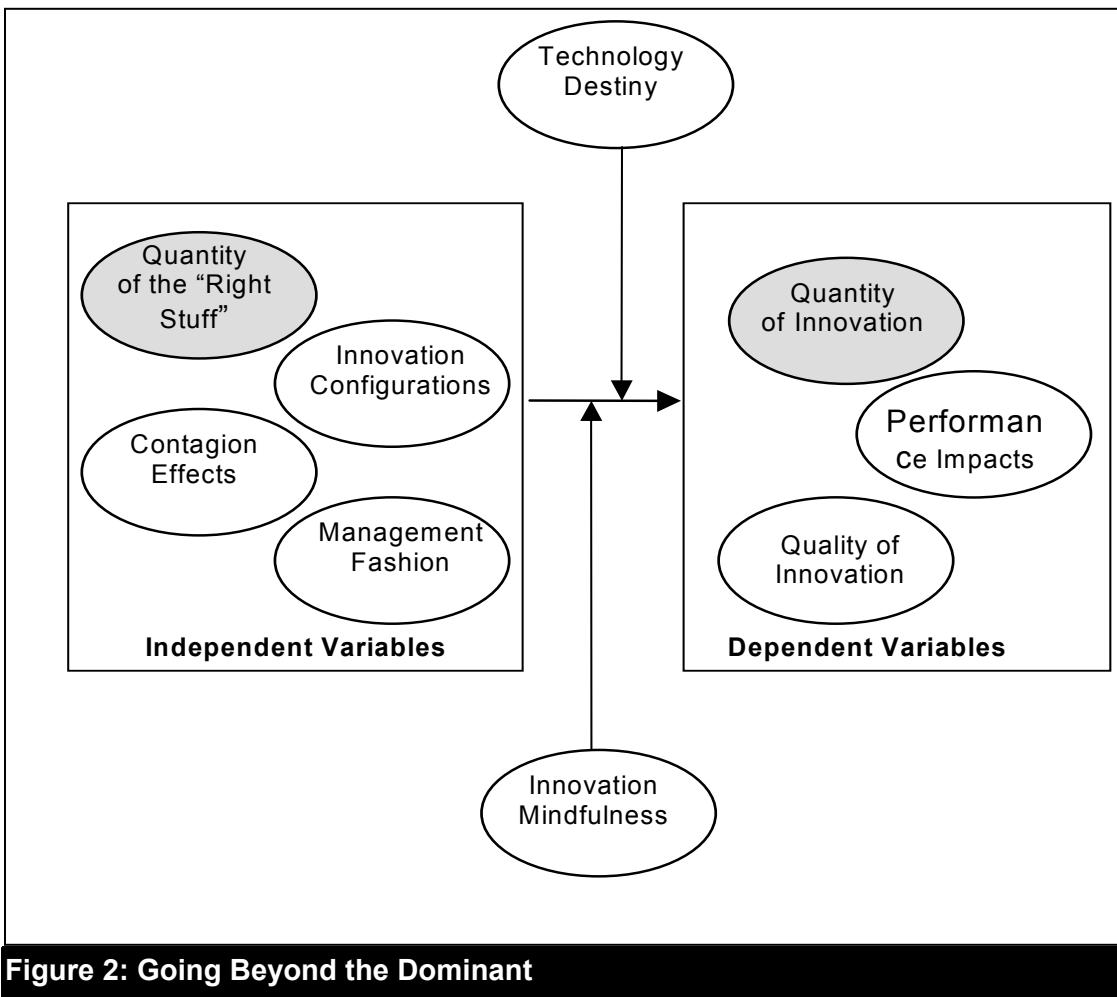
productivity benefits flowing from IT investments in general, and innovative IT applications in particular (Melville *et al.*, 2004). In fact, many observers have attributed the recent acceleration in productivity growth in the United States to the cumulative effect of major investments in IT made in the 1990s (Ives *et al.*, 2003).

In addition, the strategic choice perspective on organizational decision-making—which implicitly underlies the dominant paradigm—has a long history in management research generally and is often an accurate description of how firms operate (Nutt, 2002). Alternatively, the population ecology perspective, another venerable stream, can be seen as enforcing roughly rational behaviors based on selection from the environment (Hannan and Freeman, 1977). The idea here is that organizations that are persistent in their failure to adopt beneficial innovations will become less well adapted to their competitive environments, and will therefore be more prone to deselection by that environment.

And finally, the models developed from within the dominant paradigm generally do work reasonably well as far as answering the questions they are intended to answer. A meta-analysis by Damancpour (1991) showed that many generalizations of the dominant paradigm have held over a wide variety of studies. Furthermore, models consistent with the dominant paradigm can be highly predictive in the right circumstances. Meyer and Goes (1988), in a careful study of the assimilation of ten medical innovations among thirty hospitals, explained 74% of the variance in aggregate innovation in those hospitals. In the IT domain, Fichman (2001) developed a model that explained 49% of the variance in the propensity to assimilate software process innovations in corporate MIS departments. Thus, I am by no means suggesting that work within the dominant paradigm is misguided or should not be pursued going forward.

Beyond the dominant paradigm

Nevertheless, I will argue for a substantial increase in attention to research falling outside this paradigm. In essence, I am calling for a more innovative approach to the study of IT innovation itself. Toward this end I will survey seven promising opportunities to go beyond the dominant paradigm, as summarized in Figure 2 and Table 1. Three of these areas—*contagion effects*, *management fashion*, and *innovation mindfulness*—have already come to the attention of a few IT innovation researchers. In contrast, three other areas—*innovation configurations*, *technology destiny*, and *quality of innovation*—have yet to be studied in the IT context and may represent more of a ground-floor opportunity for future research. Finally, the seventh area, *performance impacts* of innovative IT, has been examined in several studies, but usually in order to demonstrate that positive impacts exist rather than to address the issue of whether, when, and how to innovate with IT.

**Figure 2: Going Beyond the Dominant**

Reconceptualizing the Independent Variables

This section presents three departures from standard conceptualizations of independent variables in IT innovation research. I begin by using the notion of *innovation configurations* to develop an alternative conceptualization for how independent variables interact to produce an innovation outcome. I then examine two research streams—*social contagion* and *management fashion*—that identify causal forces that depart from standard economic-rationalistic assumptions.

Innovation Configurations

The typical approach to model building within the dominant paradigm is to identify a set of direct antecedents of innovation, with each antecedent assumed to have independent effects that contribute to the variance explained in the outcome variable (such as the propensity to adopt or implement a single innovation or a group of innovations). With such approaches, the greater the weighted sum of an organization's score on the independent variables, the greater the extent of predicted innovation. The result is a *variance* model that operates under a *more is better* logic with regard to included variables (Markus and Robey, 1988).

Table 1. Summary of Seven Perspectives

Perspective/ Key Citations	Central Concept	Central Argument	Key Research Questions
Innovation Configurations (Ragin 1987; Ragin 1999)	An innovation configuration is a specific combination of factors that are collectively sufficient to produce a particular innovation-related outcome.	The factors that affect complex phenomena (such as large scale IT innovation) can interact in complex ways that go beyond simple linear interaction effects, and thus must be viewed holistically to draw valid conclusions.	Which holistic combinations of factors explain IT innovation outcomes, especially in cases where there are smaller numbers of large scale events with more extreme outcomes (i.e., dramatic success or failure)?
Social Contagion (Greve 1995; Teo <i>et al.</i> 2003)	Social contagion exists when organizations feel social pressure to adopt an innovation that increases in proportion to the extent of prior adoptions.	Potential adopters are highly affected by <i>who</i> has adopted an innovation in the past rather than just by the technical merits of the innovation taken in isolation.	What are the effects of social contagion (and the related effects of susceptibility and proximity) on IT innovation? What is the relative strength of different institutional forces (normative, coercive, mimetic)?
Management Fashion (Abrahamson and Fairchild 1999; Carson <i>et al.</i> 2000)	Management fashion waves are relatively transitory collective beliefs, disseminated by the discourse of management-knowledge entrepreneurs, that a management technique resides at the forefront of rational management progress.	The fashionable discourse disseminated by management-knowledge entrepreneurs creates and reinforces institutional forces that can cause the adoption of inefficient innovations and the rejection of efficient ones.	What triggers the emergence of IT fashion waves? What effect does fashionability have on innovation adoption and institutionalization? What factors determine whether a fashionable intervention will become a transient fashion or an enduring institution?
Mindfulness (Fiol and O'Connor 2003; Swanson and Ramiller 2004)	An organization innovates mindfully to the extent that it attends to the innovation with reasoning grounded in its own facts and specifics.	Mindful organizations are more likely to make sound judgments about whether to adopt an innovation, when, and how best to manage the assimilation process.	What are the antecedents and decedents of mindful innovation? Do more mindful organizations experience superior innovation outcomes? What factors lead to greater or lesser mindfulness at the level of a population?
Technology Destiny (Fichman 2003)	Technology destiny is the ultimate disposition of a technology at the point it is no longer considered to be something new among most members of	The destinies that await IT innovations vary dramatically (from universal adoption to total abandonment); models based on assumptions of	Do the determinants of innovation adoption or assimilation vary depending on destiny? Are any organizational factors associated with a

	its target adoption community	normative rationality become problematic in the context of innovations with a poor destiny.	greater propensity to adopt innovations with a positive destiny, and avoid those with a negative destiny?
Quality of Innovation (Haner 2002)	The quality of innovation is the extent to which an organization has adopted the “right” innovation, at the “right” time and in the “right” way.	The quality of innovation, rather than the mere quantity of innovation, is the more important determinant of beneficial innovation outcomes.	What characterizes organizations that exhibit greater quality of innovation? How does quality of innovation relate to innovation outcomes?
Performance Impacts (Dos Santos and Peffers 1995; Hitt <i>et al.</i> 2001)	Performance impacts capture the effect an innovation has on business process measures (e.g., inventory turns), firm level measures (e.g., productivity and accounting profit based), and market-based measures (e.g., capitalization, Tobin’s Q).	Performance impacts from IT innovation will vary depending on when and how innovations are adopted and on the presence or absence of complementary competitive conditions, organizational resources, and managerial actions.	What is the relationship between IT innovation and organizational performance? Under what contextual conditions are the performance impacts the greatest? How do mindfulness, quality of innovation and technology destiny affect the returns to innovation?

One limitation of the dominant approach is that it does not allow for complex interactions among the factors that go beyond simple linear additive (or multiplicative) effects. In particular, there may be theoretical contexts where what matters is the holistic *configuration* of factors that are present or absent. A brief example illustrates the point. In the context of large-scale enterprise package implementations, it is fairly well established that rapid, big-bang style implementations are especially prone to disaster. Thus, in a standard innovation model formulation, one would posit a negative correlation between implementation success and the big-bang method that would exist regardless of the presence or absence of other factors. Yet, Austin *et al.* (1999) document an ERP implementation at Cisco that was quite successful despite using a rapid, big-bang approach. A closer examination of this case reveals that it contained a distinctive configuration of factors, including: (1) a decision by top management to assign eighty of the “best and brightest” employees in the company to work full time on the project, (2) the participation of a highly experienced integrator (KPMG), and (3) the presence of a particularly hungry vendor (Oracle) that hoped to use this implementation as a showcase for future clients. Under the configurational approach, then, the logic explaining this case would be: rapid pace + big bang + best and brightest dedicated full time to the effort + highly experienced integrator + hungry vendor with something to prove => major ERP project success. This one case does not, of course, prove that this particular configuration ensures success, or that all the factors in the configuration are necessary, or that other quite different configurations might not work just as well. To reach such conclusions requires that one examine several cases in a systematic way. The *qualitative comparative analysis* method, to be described next, provides this sort of systematic approach.

Qualitative Comparative Analysis

Interest in a configurational approach to understanding complex phenomenon originated

in the field of comparative sociology with the work of Charles Ragin (Ragin, 1987, 1999). Ragin sought to develop a case study method that had quantitative rigor, yet would treat each case holistically and preserve its full causal complexity: "The basic idea is that a phenomenon or a change emerges from the intersection of appropriate pre-conditions—the right ingredients for change. In the absence of any one of the essential ingredients, the phenomenon—or the change—does not emerge. This conjunctural or combinatorial nature is a key feature of causal complexity" (Ragin, p. 25). The resulting method came to be known as *qualitative comparative analysis*, or QCA. This method uses the logic of Boolean algebra to determine the most parsimonious sets of inter-related conditions that explain the outcomes observed among a given set of case examples. In particular, the method identifies *necessary* and/or *collectively sufficient* conditions to produce an outcome.

QCA has been used to examine such phenomenon as labor management practices in southern textile mills (Coverdill and Finlay, 1995), labor organizing drives (Boswell and Brown, 1999), and barriers to government policy diffusion in the United States (Kitchener et al., 2002). The method differs from standard comparative case study logic both by the formal structure it imposes and by its ability to handle larger numbers of cases. With a traditional comparative case study approach, the research will compare all cases to each other in the search for commonalities among them in terms of causal factors or contextual conditions leading to an outcome. As the number of cases grows beyond the single digits, this process becomes increasingly unmanageable, because the number of pair wise comparisons grows geometrically (Ragin, 1987, p. 49). By contrast, the complexity of QCA does not increase rapidly with the number of cases—although as explained below complexity does rapidly increase with number of included causal factors.

QCA allows that the same factor might promote an outcome, inhibit it, or have no effect, depending on the presence or absence of other factors. Returning to the ERP example above, we might imagine that if a firm is to put eighty of its best people on a project, then that project had better be accomplished quickly lest the operational business deteriorate. Thus, it might be that the factor *devote the best people full time* works well when combined with a *rapid pace* of implementation, but could lead to a poor outcome if the pace is very slow. QCA is a method that allows researchers to isolate this sort of causal complexity.

Following is a brief sketch of how QCA works. First, the researcher defines the theoretical scope based on some outcome of interest (e.g. extreme success on very large enterprise implementations) and then samples cases as exhaustively as possible from a relevant universe of cases (e.g., very large ERP implementations). Then the researcher analyzes each case to determine, through some combination of deduction and induction, which factors are most salient in producing the outcome of interest. Next, the researcher must assign dichotomized values to these factors for each case (representing the presence or absence of the factor). Following this, a "truth table" is constructed that enumerates all possible combinations of the explanatory variables, and for each combination, what outcome (if any) was observed. Then the researcher analyzes the truth table (with the help of software packages developed for this purpose) to determine the most parsimonious sets of conditions (i.e., presence or absence of each factor) that explain the outcome of interest. Due to its focus on identifying necessary and collectively sufficient conditions to produce some particular outcome, QCA can be seen as a rigorous method for testing the logic of a process model, rather

than the more-is-better logic of a variance model (Markus and Robey, 1988).

QCA: Methodological Challenges

While QCA has not been used in IT innovation research before, it holds promise in addressing a key limitation of variance research; namely, given the very large number of potentially important variables affecting innovation, how is a manager to know which are *essential* to attend to in any given context? And how is a manager to know which mechanisms and strategies go together especially well? Yet the benefit of QCA's rigor comes at the price of some methodological challenges. While many of these challenges can be managed with good research design and proper matching of the technique to the problem (see Table 2), they do limit the scope of potential application. Thus, the configurational approach is unlikely to emerge as a serious rival to the variance approaches that dominate among quantitative researchers, but rather is more likely to serve as a new tool in the case researcher's arsenal. Indeed, some case method researchers have observed that QCA (despite its challenges) helps to overcome subtle biases, and provokes examinations of possibilities that would otherwise have been overlooked (Kitchener *et al.*, 2002).

Two issues arising from Table 2 warrant additional discussion. The first is what to do about configurations for which no outcome was observed. The usual approach is to code unobserved combinations as "doesn't matter." When this is done, the QCA algorithm assigns a hypothetical outcome for such missing cases that results in the most parsimonious explanation for observed cases. Other approaches are to assume missing configurations result in a non-event or to use one's judgment to assign the most theoretically plausible outcome that would have been observed for each missing configuration. Clearly, as the observed configurations for the truth table become increasingly sparse, the study's conclusions become increasingly speculative. There are no hard-and-fast rules for what portion of the truth table must be observed, although most past examples have had about half of the table covered. Whether coverage of the truth table is adequate will also depend on the plausibility of "missing" cases occurring in practice. Configurations that can not happen in practice should be discounted in evaluating the coverage of the truth table. A related issue is whether all known cases of a phenomenon or a sampling of cases were studied. If all known cases were studied but a configuration was not observed, this lends credence to the conclusion that a configuration is missing because it is infeasible.

The second issue is the problem of sampling bias. Unlike large scale statistical studies, where inclusion or exclusion of a single case is unlikely to affect results, with QCA a single case can have a dramatic effect on conclusions. Thus, the further a researcher gets from full enumeration of the relevant cases, the greater the possibility that different results might have been observed with a different sample.

The issues raised in Table 2 suggest that QCA will work best when examining phenomenon that are large-scale, occur infrequently, and for which extreme outcomes (major success, major failure) are of particular theoretical or practical interest. The choice of large scale and infrequent phenomenon, aside from being in keeping with the QCA's origins in the study of notable macro-social events, makes it easier to employ an enumeration rather than sampling strategy for case selection. Besides being desirable in its own right, enumeration lessens concerns about unobserved values. The focus on extreme outcomes is suggested because it is here—rather than the "mushy middle"—

that the configurational logic, rather than the standard more-is-better logic of variance models, is most likely to hold.

Table 2. QCA Issues and Strategies

Issue	Explanation	Strategies
All variables must be dichotomized in the standard application of QCA	When Boolean logic is used each condition must be either present or absent	<ul style="list-style-type: none"> For conditions that are not naturally dichotomous the researcher can look for large natural breaks in the sample of cases
The size of the “truth table” grows intratable with more than 4 or 5 factors (a 5 factor model has 32 possible configurations)	While it is not necessary to have observations for all configurations, prior QCA studies have had observations for about half of the possible configurations. Research that covers much less than half of the table will be more prone to objections that the results are highly speculative.	<ul style="list-style-type: none"> When theorizing, focus on the most essential conditions to achieve an outcome. Limit use of the method to more unusual or extreme outcomes (major success, major failure) because such outcome are more likely to have distinctive configurations of causes (rather than result from the “more X leads to more Y” logic of variance models”. Narrow the theoretical scope to “control” for certain variables. In the Cisco case, one could take “rapid pace” and “big-bang” out of the truth table by confining theoretical scope to explaining outcomes on rapid big-bang implementations .
The method in its standard form does not allow contradictions	If two cases have the same exact configuration of factors but different outcomes the method in its standard form can not resolve this anomaly	<ul style="list-style-type: none"> Adding an additional variable can often resolve a contradiction, although then the variable must be captured for all cases, which doubles the size of the truth table. Some more advanced statistical procedures can be used to assign single value to combinations has have contradictions (Ragin 1987, ch. 7) Researcher can avoid trying using QCA to explain small differences in the “mushy middle” of potential outcomes, as one is more likely to encounter contradictions here.
The method works best when the entire universe of cases can be identified	If the research only samples from within a larger set, this raises the specter of a selection bias. Selection bias is a stronger threat to QCA than for large scale statistical studies.	<ul style="list-style-type: none"> Focus on larger scale phenomenon that are well publicized and have a smaller universe instances.

Examples of innovation phenomenon that might fit the suggested criteria of being large-scale and infrequent include massive enterprise system implementations, instances of some dominant IT standard being overturned, or attempts to deploy new national IT infrastructures. Massive enterprise system implementations are often announced in the

press (thus increasing the ease of identification) and have occurred in limited numbers (thus supporting an enumeration strategy in case selection). Instances of major challenges to a dominant standard occur infrequently; and on an intuitive level, it seems likely that a complex confluence of factors would explain those rare instances where a standard is successfully overturned. Likewise, attempts to institute major new national IT infrastructures are comparatively rare, tend to be well-publicized, and are less likely to be explained by the more-is-better logic of variance models. Some examples from the relevant universe could include positive cases such as Minitel in France, DOCOMO in Japan, and smartcards in Hong Kong; and negative case such as ISDN, use of web phones, and smartcards in the US.

Social Contagion Models of Adoption

A key assumption of most research falling within the dominant paradigm is that organizations each make independent assessments of the innovation and decide whether and when to adopt based on the inherent merits of the technology (e.g., potential to improve the organization versus the costs to adopt). Yet, there are at least four mechanisms that can lead to adoption decisions being interrelated across organizations. First, when a technology is subject to positive network externalities or other forms of increasing returns to adoption (Arthur, 1996), managers must consider the actions of other adopters when evaluating an innovation's merit because, by definition, the value of using the technology is contingent on the size of the eventual adoption network. Second, when a technology integrates a community of organizations under one umbrella, such as with supply chain management or other interorganizational systems, the adoption decision of any one firm will depend on the firm's role and position in the collaborating community (Hart and Saunders, 1998). The research streams associated with these two mechanisms are described in Fichman (2000), and will not be considered further here. The third mechanism to be considered in this section concerns interrelated adoption decisions that arise from social contagion. The fourth mechanism, to be discussed in the next section, is management fashion.

Most of the theorizing within the social contagion stream is founded on neo-institutional theory (DiMaggio and Powell, 1983; Rowan, 1977) or social learning theory (Bandura, 1977). The former is concerned with institutional forces that lead firms residing in an *organizational field* to increasingly resemble each other, resulting in *institutional isomorphism*. Researchers consider three forces that cause institutional isomorphism: (1) *coercive* forces (arising from societal expectations or from organizations on which the focal firm depends), (2) *normative* forces (arising from professionalization), (3) and *mimetic* forces (arising from the tendency to imitate peers perceived to be successful under conditions of uncertainty) (DiMaggio and Powell, 1983). Social learning theory overlaps with the mimetic force in positing that people engage in a process of social learning by examining the actions of similar peers.

When the forces of social contagion are strongly present, the adoption decision will be contingent not just on a firm's own independent assessment of the innovation's merits (per the dominant paradigm) or on the raw number of prior adoptions (as per work on network externalities), but will also depend on just *who those prior adopters were* (Strang and Tuma, 1993). Researchers can take two approaches to incorporate the effects of social contagion into models of IT innovation. One is to capture static managerial perceptions, as was done by Teo *et al.* (2003), who used perceptions of the three institutional forces as predictors of intention to adopt inter-organizational systems.

With this approach, one gets a single snapshot of managers' reported perceptions of influence, but there is no attempt to link a particular firm's adoption decision to specific prior adoptions by other firms. In fact, the link to the firm's own adoption may be tenuous, depending on how the research is designed. If the outcome variable is adoption timing, then there is a potential temporal ordering problem as *current* perceptions are assumed to cause the firm's own *previous* adoption. This temporal ordering problem looms larger here because, unlike many other kinds of less volatile explanatory factors, institutional forces are expected to be highly dynamic over time as contagion builds or ebbs away. Alternatively, if perceptions are linked to adoption intention, as with the Teo *et al.* study (2003) then one must assume that the strength of intention is closely linked to subsequent adoption timing or extent of use. While much evidence exists to support a strong link from intention to usage in more closely controlled studies of individual acceptance of personal use technologies (Venkatesh *et al.*, 2003), comparatively little evidence exists one way or another to link managerial intention to firm-level adoptions of complex organizational technologies.³ As one final pitfall of this approach, managerial self-reports of subjective perceptions (i.e., to capture institutional influences and adoption intention) can be more subject to bias, especially in single-respondent studies.

A second more direct and rigorous approach is to model contagion dynamically based on the pattern of actual prior adoptions. To support this approach, Strang *et al.* have developed a model (described below) based on an enhanced event-history analysis technique that they term the *heterogeneous diffusion model* (Greve *et al.*, 1995; Strang and Tuma, 1993),

The Heterogeneous Diffusion Model

The heterogeneous diffusion model posits that an organization's likelihood of adoption in any given period is determined by three elements: (1) a vector of factors that determine a firm's inherent *propensity* to adopt in the absence of contagion effects, (2) a vector that determines *susceptibility to contagion* from prior adopters, and (3) a vector that determines a firm's *proximity* (physical distance or organizational similarity) to prior adopters.

Greve's (1995) study of the abandonment of the easy-listening radio format over a ten-year period provides a very nice illustration of how to model *propensity*, *susceptibility*, and *proximity* in an innovation context. (Since abandoning one format means adopting another, more promising format, e.g., adult contemporary, this can be viewed as positive innovation despite what the label would imply.) Greve modeled *propensity* to abandon based on factors similar to those that have been used in prior innovation research (e.g., size of market, density of competition). So, this element of the model can be viewed as capturing the standard logic of the dominant paradigm, where firms with the "Right Stuff" are expected to adopt earlier. Conversely, *susceptibility* to prior abandonments was modeled based on the size and prominence of the prior abandoners. The basic logic here is that such adopters are, in effect, more socially contagious than smaller, less prominent adopters. Finally, Greve modeled *proximity* based on the degree of corporate contacts and market contacts between firms, and also whether they were direct

³ This is because there have been only a few studies of IT adoption intention at the organizational level, and those studies have typically been cross-sectional and did not capture subsequent usage data.

competitors. The logic is that firms with more contacts and in more direct competition are metaphorically closer together, and thus social contagion can spread between them more easily. Greve found strong support for all three effects, i.e., propensity, susceptibility, and proximity.

Methodological Challenges

While the heterogeneous diffusion model represents an attractive approach to capturing the forces of social contagion, the method imposes some methodological challenges. First, the full universe of firms residing in the organizational field should be identified, as well as the timing of the innovation event of interest (e.g., adoption, implementation, abandonment) for all firms in this universe. Second, the research must develop a sound theoretical rationale for locating each factor in the most appropriate of the three vectors (Greve *et al.*, 1995).⁴ Then, data must be captured for each factor. In cases where the value of factors will vary over time, the researcher must capture the value of that factor for each time period in the study. Once that has been accomplished, standard statistical packages can be used to estimate the resulting models.

In summary, the social contagion thread offers well-established theories (neo-institutional, social learning) for why the dominant paradigm will often provide an incomplete picture of the forces driving innovation, and also offers a novel and powerful method (heterogeneous diffusion model) for studying social contagion effects.

Management Fashion

Another emerging stream that—like the social contagion research just described—questions the assumption of rationalistic, independent assessments is based on the theory of management fashion (Abrahamson, 1996). Unlike the dominant paradigm, where innovations are assumed to be technically efficient (i.e., they provide economic benefits to adopters) this thread has been motivated more by the question of what might provoke the diffusion of technically *inefficient* innovations (Abrahamson, 1991). In addition, while the dominant perspective is more concerned with understanding the forces that produce innovations that become *institutionalized*, this perspective is concerned with innovations that experience only *transient* adoption.

As argued by Abrahamson and others, the forces of fad and fashion can have a powerful effect on the innovation diffusion process (Abrahamson, 1996; Abrahamson and Fairchild, 1999; Abrahamson and Rosenkopf, 1997). In their view, *knowledge entrepreneurs*—academics, consultants, business gurus, technology vendors—have an interest in generating demand for innovative ideas, and will seek to do so by actively promoting the idea as being on the leading edge of managerial practice. The resulting wave of media attention and innovation-related discourse produces a spurt of managerial interest and organizational adoptions—even in the absence of any systematic evidence of the innovation’s efficacy. Then, to the extent that adoption becomes more widespread, other forms of institutional pressure—from business partners, boards of directors, and shareholders—are often brought to bear on perceived laggards, thus producing a self-reinforcing adoption bandwagon. As a result, certain

⁴ Tan and Fichman (2002), who used the heterogeneous diffusion model in an exploratory study of the determinants of the adoption of web-based transactional banking by US banks and thrifts, chose to locate the same factors in multiple vectors, a tactic not observed in other applications of the method

innovations may garner managerial attention and organizational adoption out of all proportion to the ultimate benefits flowing from their actual use.

While prior research in this stream has examined fashionable managerial techniques rather than fashionable IT-related ideas and tools, the above sketch of the role of fad and fashion seems to hold for many IT innovations, which appear to be particularly prone to overly enthusiastic waves of discourse, or “hype” (Swanson and Ramiller, 1997). Thus, IT innovation appears to be a promising domain in which to extend the application of management fashion concepts and techniques.

Management Fashion Definitions

Definitions for what constitute management fashions differ. Abrahamson and Fairchild (1999, p. 709) define management fashion as “relatively transitory collective beliefs, disseminated by the discourse of management-knowledge entrepreneurs, that a management technique resides at the forefront of rational management progress.” This definition has four key elements. First, fashions are transient rather than enduring; this is perhaps the most essential feature of fashions as commonly understood. A second element is that the fashion itself is not some real world thing, but rather, collective beliefs about a thing. This element may be problematic, as I will discuss shortly, but does echo the featured role of a similar concept, *collective tastes*, in more general theories of fashion (see especially Blumer (1969)). The third element is a focus on management fashion setters as the mechanism by which fashions are disseminated. This echoes the important role of social elites in the spread of fashions more generally. The fourth feature is that the technique is viewed as beneficial and progressive.

Two facets of the above definition are noteworthy. First, the definition is agnostic about the merits of the underlying technique itself; that is, despite the common connotation that fashions are somehow lacking in substance, this aspect does not enter into the definition. In fact, Abrahamson and Fairchild state explicitly that this definition “does not imply that a fashion either is or is not dysfunctional” (1999, p. 709).⁵ Second, according to Abrahamson’s definition, a fashion is not constituted by the technique itself, but rather, by what people believe about it. This seems to lose a key element of our everyday understanding of fashions. For example, it would not seem quite right to say, by analogy, that mini skirts or hula hoops themselves were not fashions, only what people thought about them. This focus on belief also leads to some imprecise and potentially misleading phraseology, such as when Abrahamson discusses particular fashions being “adopted” by organizations, or the failure of fashions to become “institutionalized” across the broader community. We certainly want to discuss such matters in the context of particular fashions, but it does not make much literal sense to say that it is the collective beliefs, not the underlying technique, that are being adopted or institutionalized.⁶

Carson et al. (2000) provide an alternative approach that defines management fashions as “*interventions that*:

⁵ It is possible Abrahamson thinks this lack of substance is implied by the transitory nature of the collective beliefs that the technique is on the forefront. However, unless progress itself were to come to a halt, any particular technique can reside only temporarily at the “forefront”, thus all innovations, not just fashions, would qualify as transitory in this particular sense.

⁶ It could be argued that what Abrahamson really means in such cases is the adoption or institutionalization of the technique-for-which-there-was-an-associated-fashion.

- Are subject to social contagion because they are novel and perceived to be progressive, or preferable to preexisting fashions,
- Are, or are perceived to be, innovative, rational, and functional,
- Are aimed at encouraging better organizational performance either materially or symbolically, through image enhancement,
- Are motivated by a desire either to remedy some existing operational deficiency or to prospectively capitalize on opportunities for improvement, and
- Are considered to be of transitory value because, despite a ‘post latency’ period of acceptance, no systematic and comprehensive research legitimizing their prolonged utility or generalizability emerges.”

This definition differs markedly from Abrahamson by defining a fashion as the technique or intervention itself rather than collective beliefs about the technique. It also differs by including some indication that an intervention lacks verified substance—in this case because the intervention is viewed as having transitory value and because of an absence of research legitimizing its utility or generalizability. Yet this definition is not without limitations. For one thing, it is rare for interventions, even those that become universally institutionalized and taken for granted as beneficial, to be studied post-hoc in systematic, rigorous ways to prove their utility. So the fifth criterion does not really offer the intended degree of discrimination. Also, it would be preferable to have a conceptual definition that can be communicated in one sentence.

It may be difficult to develop a single conceptual definition of management fashion that adequately captures all of the following key elements: that some real world intervention exists that, together with collective discourse and beliefs, constitutes the fashion; the necessarily transient status of any intervention as a fashion; and that to be a fashion in the common understanding of the word, an intervention (or at least the interest surrounding it) must be somehow lacking in substance. Thus, I will suggest that we treat the concept of management fashion as having three distinct connotations, each with its own precise labeling and definition. Toward this end, I suggest that the term management fashion be understood to mean either a *fashion wave*, a *fashionable intervention*, or a *transient fashion*, with the specific connotation being explicitly delineated by the user of the term. I suggest that Abrahamson and Fairchild’s definition could be redeployed as-is to define a *fashion wave* as *relatively transitory collective beliefs, disseminated by the discourse of management-knowledge entrepreneurs, that a management technique is at the forefront of rational management progress*. I suggest that a *fashionable intervention* be defined, based on an adaptation of the prior definition, as *an intervention for which there are relatively transitory collective beliefs, disseminated by the discourse of management-knowledge entrepreneurs, that it resides at the forefront of rational management progress*. (An alternative definition for *fashionable intervention* could be constructed using the first of four of Carson *et al.*’s criteria.) Finally, I suggest that a *transient fashion* be defined as *a once-fashionable intervention that fails to become broadly institutionalized among its intended adopter population*. This connotation would be most similar to Carson *et al.*’s meaning of the term, and could be operationalized by substituting lack of institutionalization for Carson *et al.*’s fifth criteria for defining a management fashion. The lack of institutionalization can be viewed either as an aspect of an innovation that lacks substance (i.e., it lacks substance because it did not live up to expectations and had marginal long-term impacts), or as an indicator that the innovation lacked merit (i.e., because it lacked substance, it went on to wholesale abandonment despite initial popularity).

Management Fashion Research Methods

Like the social contagion thread described in the prior section, prior research on management fashion views institutional forces as a key driver of adoption. However, unlike the social contagion thread, fashion research has focused on *management discourse*, rather than the weight of prior adoptions, as the mechanism by which institutional forces are brought to bear. As a result of this distinctive focus on discourse, fashion research employs correspondingly distinctive empirical methods based on bibliographic search techniques. These techniques have been used to examine the following general research questions: (1) why do *fashion waves* arise when they do? (2) what characterizes the typical patterns of discourse *intensity* surrounding *transient fashions*, (3) what patterns exist in the *content* of this discourse? (4) what intervention characteristics or macro-level environmental factors might correlate with patterns of discourse intensity?⁷

Despite the fact that research questions (2), (3), and (4) could, in principle, be investigated for any fashionable intervention, not just the subset of transient fashions, prior fashion research has, in fact, focused exclusively on transient fashions. This is not surprising, since the question of what accounts for transience is central to the whole stream. However, as will be discussed later, this exclusive focus represents a limitation of prior research—but also an opportunity for IT researchers. Unlike the general management domain, where apparently very few fashionable interventions become institutionalized, a fair number of fashionable IT interventions have become institutionalized.

Two recent studies on management fashion provide exemplars of research addressing research questions (1) through (4) above, and also illustrate the diversity of potential methods that can be employed. Abrahamson and Fairchild (1999) conducted an intensive case study of a single fashion, quality circles. They found that a new fashion is triggered when the downswing of a prior fashion occupying the same niche occurs—and when the performance gap the fashion purportedly addresses becomes more salient due to a combination of environmental forces intensifying the gap and increased discourse highlighting the importance of the gap. In addition, they demonstrated that during the upswing in discourse intensity, the content of discourse tends to be more emotional, enthusiastic, and unreasoned, while on the downswing it tends to be more unemotional, reasoned, and qualified.

As a second exemplar of empirical research on management fashions, Carson *et al.* (2000) examined 16 management fashions that emerged over a five decade span (management by objectives, T-groups, TQM, reengineering, etc.). They found a general trend toward shorter and more intense cycles of discourse. Of particular interest to IT innovation researchers, they also found correlations between characteristics of the fashions themselves (e.g. radicalness, difficulty of implementation, broadness of potential application) and characteristics of the discourse intensity lifecycles surrounding the fashions (year of introduction, time to reach the peak of discourse intensity, intensity of the discourse peak, and duration of the entire lifecycle). This sort of analysis is similar in spirit to a long line of diffusion research matching innovation characteristics to aspects of the diffusion lifecycle (i.e., rate and extent of diffusion).

⁷In articulating these research questions I have substituted the more specific terms suggested earlier for the generic term “fashion”:

Interestingly, although these and other studies have considered the macro and historical conditions that enable a particular fashion to emerge, little research has been devoted to the question of what aspects of the innovations themselves might make them more or less fashion driven. However there are a few hints in the literature. For example, Abrahamson has argued that administrative innovations are more likely to be fashion driven than technical innovations, because their outputs are more difficult to observe (Abrahamson, 1991). In other words, ambiguity about effects will tend to make an innovation more fashion driven. Miller and Hartwick (2002) suggest the following as the hallmarks of fads: simple; prescriptive; falsely encouraging; one-size-fits-all; easy to cut-and-paste; in tune with the zeitgeist; novel, not radical; and legitimized by gurus and disciples. Benders (2001) offers a similar list in describing the characteristics of discourse associated with fashions: promises of, preferably substantial, performance enhancement; the threat of bankruptcy in the case of non-adoption; using well-known and successful users of the concept in question; stressing the concept's universal applicability; presenting the concept as an easily understandable commodity with a catchy title; presenting the concept as timely, innovative, and future-oriented; interpretive viability, i.e., leaving certain room for interpretation.

Future Directions for Research on IT Fashions

Although to date most published research has focused on “pure” management fashions, Abrahamson (1996, p. 258) argues that “management researchers should remain open to the possibility that not only management, but also marketing, finance, accounting, operations, as well as almost every area of technical endeavor, are open to the swings of *fashion*. Theorists should therefore attempt to develop theories of *fashion* in technical realms...” This suggests an opportunity for IT researchers to make a contribution to the emerging field of management fashion even while they enlarge the scope of IT innovation research.

Conceptual and Methodological Challenges

Research on IT fashion is subject to a few conceptual and methodological challenges. A first challenge is to clearly delineate what does and does not constitute an IT fashion. If we want to preserve a stronger linkage between IT fashions and the broader management fashion stream, it would be appropriate to limit IT fashion research to organizational IT innovations that have non-trivial effects in the management domain, i.e., they impact key organizational processes, structure, roles, incentives, etc. These could be innovations in IT management *per se*, such as those that relate to the software development process, or they could be those innovations embedded in software/hardware artifacts that have a significant managerial component, such as groupware or ERP. So, on one extreme, purely technical innovations (e.g., cable modems) might be of less interest from a management fashion perspective. On the other extreme would be innovations that are so management-oriented (e.g., knowledge management, business process reengineering, IT outsourcing) that one must think carefully about whether the intervention is sufficiently IT-specific to qualify as an “IT” managerial fashion. This latter concern would, of course, be most salient for research that wishes to draw distinctions between IT management fashions and “pure” management fashions. In fact, one interesting research opportunity would be to develop a measure of the degree of “IT-ness” of an innovation, and to develop plausible hypotheses for how the degree of “IT-ness” might affect the discourse intensity lifecycle.

A second challenge is that as a general rule, management fashions can be difficult to

study—or even to identify—in real time. If the primary interest is transient fashions, as has been the case in past fashion research, we must wait to see which fashionable interventions are in fact transient. However, even if that were not the case, the actual research methods that have been used to study fashions often rely on variables relating to the lifecycle of discourse (i.e., duration) that can only be gathered in retrospect. This raises a question about the potential for practical relevance of management fashion research in the IT domain. More specifically, if we can not identify which fashionable interventions are destined to become transient fashions, rather than enduring institutions, then the opportunity to impact managerial practice will be limited. This suggests attention in IT research to developing propositions discriminating these two ultimate outcomes, a topic I will take up below.

The third challenge is that rigorous bibliographic analyses are not easy. Researchers must identify which key word or words (if any) can be used to identify a valid and complete set of articles potentially about the innovation. Criteria must be established to decide which articles are really about the innovation, rather than just mentioning it in passing. Counts of articles must be normalized to account for growth in the total number of articles published over time. This is particularly an issue for IT, which has probably seen a greater rate of increase than other domains. Fortunately for IT researchers, the two aforementioned exemplars of empirical fashion research provide excellent templates for how to manage these challenges.

Distinctive Features of IT Managerial Fashions

Using the fashion lens to study managerial innovations that involve IT—rather than the prior focus on pure management innovations—can be best motivated by considering some ways in which IT innovations tend to differ from pure management innovations and the potential implications of those differences. Toward this end, I offer several preliminary ideas.

First, IT innovations usually have an identifiable market of tangible artifacts and/or services that can provide objective evidence about the innovation's uptake and persistence, and this opens new opportunities for linking the lifecycle of discourse surrounding IT innovations to the diffusion lifecycle of the innovation itself. While it can be difficult to establish with certainty whether a purely managerial innovation has in fact been transient without gathering primary data, there are often good market data available to make this assessment for IT innovations, which usually have a commercial hardware or software component.

Second, there have been quite a large number of fashionable IT innovations over the last three decades, and many of them (e.g., relational database management technologies, EDI, certain forms of ecommerce) seem to have become widely institutionalized. This gives IT researchers a unique opportunity to examine whether it is possible to distinguish in real time transient fashions from enduring institutions that were once fashionable. As argued above, this ability may be central to establishing the managerial relevance of fashion research in the IT domain.

Third, IT innovations, once deployed, tend to have heavy switching costs at both the organizational and community levels. This might be used to develop propositions relating the degree of switching costs to a fashionable intervention's final status as either transient or enduring. That is, other things equal, one would expect innovations with the

highest switching costs to be more enduring. This may well explain, for example, the persistence of the once highly fashionable but subsequently much maligned ERP movement.

Fourth, IT innovations vary widely in terms of the relative emphasis between administrative and technical components. Abrahamson (1991) has posited that the fashion perspective should have the greatest explanatory power for administrative innovations because they have unclear outputs. This suggests that a collection of IT innovations could be used to investigate whether this is true, i.e., is the link between fashionability of discourse and transience stronger for IT innovations that have the heaviest administrative components?

Fifth, IT innovations are often subject to economically-based self-reinforcing adoption dynamics (related to learning and positive network externalities), thus providing an alternative to the fashion-based explanation for observed volatility of managerial interest and related discourse and adoption.

In conclusion, IT innovation researchers interested in management fashions have several avenues available to them. One would be to attempt to replicate key findings found in prior research in the IT domain. A second, more promising route is to take the distinctive elements of IT innovations (such as those identified above) and use them to develop and test a specific theory of IT fashion. A third avenue is to take concepts related to managerial fashion and integrate them with other emerging streams, such as those related to technology destiny or innovation mindfulness, as discussed below.

A Parting Word on the Macro-level Effects of Fashions

Although transient management fashions—almost by definition—tend to produce little lasting value in themselves, it would be a mistake to view the larger fashion phenomenon as necessarily dysfunctional or counter productive, for a few reasons. First of all, fashionable management discourse can serve as an important mobilizing force in the business community surrounding not just transient fashions, but also what will become enduring institutions (Abrahamson, 1991; Swanson and Ramiller, 1997). It could well be that the resources lost to transient fashions are a price that a community should be willing to pay in order to discover those fashionable innovations worth retaining. It is perhaps instructive that the United States seems especially prone to swings in management fashion, but also has one of the most productive, innovation-driven economies in the world. Also, some transient fashions may have complementary or catalytic effects that enable or otherwise facilitate the adoption of other beneficial innovations.⁸ For example, first generation PDAs—a highly fashionable innovation circa 1990 that failed miserably in the market—did lay a foundation for the highly successful second generation PDAs introduced about a decade later.

⁸ I thank Bob Zmud for offering this observation.

Reconceptualizing Some Linkages between Independent & Dependent Variables

In the previous section I considered some streams of research that suggest either new kinds of innovation antecedents (i.e., social contagion, management fashion) or a different logic for how antecedents may act in combination to affect innovation (i.e., innovation configurations). In this section I consider two streams that question the fundamental more-is-better logic that underlies the linkages between antecedents and outcomes in the dominant paradigm (see Figure 2). The first stream focuses on the nature of the innovation itself and posits that a technology's *destiny* will moderate the relationship between traditional antecedents and some kinds of traditional innovation outcome variables. The second stream focuses on organizations, and posits that *mindfulness* will serve as a key moderator or mediating variable in models of innovation.

Technology Destiny

Despite the pro-innovation bias associated with the dominant paradigm, not all innovations are beneficial. Abrahamson (1991) draws a distinction between technically efficient and technically inefficient innovations. He calls an innovation technically efficient if it would economically benefit most members of the target population to adopt and deploy it. Conversely, an innovation is technically inefficient if it would not benefit most members of the target population to use it.

However, a key limitation of this notion of technical efficiency is that for most IT innovations, technical efficiency is not a static feature that exists outside a community's response to the innovation. Rather, the typical IT innovation only *becomes efficient* through a recursive process where each round of adoption contributes to positive network externalities, triggers more rapid maturation of the technology, promotes additional investments in essential complementary technologies, and increases public knowledge about the innovation and how best to apply it. Thus, I will argue that a technology's ultimate fate, or *destiny*, may provide a better insight into a technology's lasting benefits than some static concept of technical efficiency.

Elsewhere, I have defined destiny as the *ultimate disposition of a technology at the point it is no longer considered to be something new among most members of its target adoption community* (Fichman, 2003). On one extreme are technologies with the most favorable ultimate dispositions. These become universally adopted and deployed within the targeted adoption community (e.g., relational DBMS, email, and GUIs). On the other extreme are technologies that, despite initial promise and rapid rates of adoption, are only sparsely deployed by adopters or even abandoned altogether (e.g., expert systems for knowledge work, CASE, first generation PDAs). Of course, there are many points lying between the two extremes of pervasive institutionalization and complete abandonment.

While much prior research has examined factors that may affect a technology's destiny (e.g., relative advantage, complexity, compatibility) no prior research has considered how a technology's unfolding destiny might in turn affect the determinants of organizational innovation. Yet the implications of technology destiny here could be far reaching. Under the dominant paradigm, firms that fit the innovator profile are expected to be more innovative because innovation is assumed to be good, and such firms have greater needs and abilities pertaining to this good thing. In the case of technologies with

a favorable destiny, this logic holds. But, in the case of innovations with a poor destiny, this logic may break down because sustained innovation with such technologies is much less likely to be a good thing. This line of argument suggests that technology destiny may be an important moderating variable in models of organizational innovation with IT.

I will explain in more detail how destiny could act as a moderating variable shortly, but first let us briefly consider how the concept could be measured. Several possibilities come to mind, such as: (1) the proportion of the intended adopting population that eventually institutionalizes the innovation, (2) the cumulative expenditures for products related to the innovation, or (3) the aggregated net economic benefits flowing from ongoing use. Perceptual measures (from managers or experts) on these or other indicators could represent another potential approach.

I now consider the effects of destiny on traditional models of organizational innovation with IT, starting with models of adoption, then moving on to models of implementation among adopters.

Technology Destiny and the Innovator Profile

To understand the effect of destiny on models of innovation, it is important to distinguish pre-adoption and post-adoption behaviors. All firms, even those that fit the generic innovator profile, will tend to suffer from limited information about technology destiny pre-adoption. As a result, the standard innovation model should hold when predicting adoption and pre-adoption behaviors (such as adoption timing) even for technologies with a poor destiny. This is because while potential adopters will naturally differ in their assessments about destiny, a technology's *actual* destiny will be so difficult to foresee during the time when most early adoptions are occurring that it is unlikely to have a large affect on the predictors of adoption timing or propensity to adopt. Thus, we should still expect organizations that fit the innovator profile to be more prone to be an early adopter of a technology even when it is destined to fail, for the simple reason that its destiny is still largely unforeseeable to potential adopters at the time that most early adoptions are occurring. This is especially true when we confine our attention, as most innovation research does, to those emerging technologies that appear quite promising.

However, when we look at post-adoption behaviors, we are now examining a special group of firms that are beginning to get better information on a technology's likely destiny. These organizations are gaining direct experience with a technology, and so managers will have a better opportunity to assess whether it is performing as expected or whether benefits have been exaggerated. Also, if implementation is going poorly they might conclude that other similar organizations are also having difficulties and therefore that costs and risks have been underrated. Thus, the process of implementation itself can afford organizations special insights about a technology's ultimate destiny that are difficult to obtain as a passive observer. In addition, since implementation often unfolds over a period of years, the passage of time will increase external information about a technology's unfolding destiny. We may expect that adopting firms will have a greater interest in and appreciation for this external information because they have more at stake, and because they have greater knowledge and experience and thus greater absorptive capacity related to the technology than non-adopters (Cohen and Levinthal, 1990).

When a technology continues along a trajectory toward a favorable destiny, we can

expect that firms that were justifiably more aggressive in their adoption timing due to their fit with the innovator profile will continue to be more aggressive in their rate of implementation. This is because implementation is still desirable, and especially desirable for them due to their strong fit. Thus, I posit that the innovator profile should still predict speed/extent of implementation among adopters under the condition of a favorable destiny.

On the contrary, when destiny takes a turn for the worse, it is no longer clear that sustained implementation is desirable. While it might be argued that those firms with greater innovation-related needs and abilities—and thus better fit the innovator profile—could, if they wanted, achieve a more rapid and thorough implementation than those that do not fit the profile, these firms (along with everyone else) should be less likely to want to. Thus, we should expect to see a dampening of the usual positive effect of innovator profile on implementation under the condition of poor unfolding destiny. It might even be argued that firms that fit the innovator profile would be quicker to perceive a negative turn in destiny, though this is speculative and not necessary to the argument.

As result, the extent to which the economic-rationalistic logic of the dominant paradigm explains the rate or extent of implementation should vary depending on a technology's destiny. In particular, we should expect the innovator profile to diminish in predictive power for technologies with a poor destiny. By contrast, models predicting adoption timing should be comparatively immune to the effects of destiny because of the difficulty that passive observers have in assessing destiny early in the overall diffusion cycle.

Technology Destiny and the Escalator Profile

If the variables associated with the innovator profile become less influential for technologies with a poor destiny, this naturally leads one to wonder whether any variables become more influential for such technologies. Here, the literature on escalation of commitment to failing courses of action may provide some clues. This literature has found, among other things, that escalation is more likely when: (1) current project managers are responsible for initiating the project; (2) project managers' level of risk perception and risk propensity is high, (3) project managers hold asymmetric information about project status, and (4) perceived sunk costs associated with the project are higher (Bowen, 1987; Brockner, 1992; Keil *et al.*, 2000; Schmidt and Calantone, 2002). While persisting in the implementation of an innovation with a poor destiny is not always the kind of "failing course of action" considered in the escalation literature, it will more often be so than when a technology has a favorable destiny. Technologies with poor destinies are more likely to have been overrated or misunderstood compared to those with a favorable destiny. Also, persisting with the implementation of a "stranded" technology carries the opportunity cost of not joining the more robust network for the winning technology. Thus, potential escalation situations will be more prevalent among the adopters of technologies with poor destinies, and this in turn suggests that organizations possessing characteristics associated with escalation may be more likely to persist in the implementation of those technologies.

Methodological Challenges

A key methodological challenge is that technology destiny can only be determined with certainty in retrospect, so research that includes destiny as a variable may have to wait until later in the diffusion cycle to be completed. This is not necessarily a debilitating limitation, because, as a practical matter, enough time has often elapsed by the time a

study is completed and being prepared for publication to discern a technology's destiny. Also researchers have the option to design a study with multiple objectives—some that require destiny to be known and others that do not require a classification of destiny and can be published immediately.

Alternatively, rather than waiting for destiny to be known with certainty, there may be ways to get an early window on a technology's likely destiny based on expert ratings, initial diffusion patterns (such as the presence of a large assimilation gap (Fichman and Kemerer, 1999)), or characteristics of discourse surrounding the innovation (are more "fashionable" innovations less likely to be broadly deployed?).

As a parting word on the destiny concept, I will grant that the very idea that an emergent property measured *a posteriori* at the level of an innovation could "affect" innovative behaviors in real time at the level of an organization seems paradoxical. The resolution of this paradox is that destiny at the level of a technology does not act as a *causal factor* in the traditional sense. Rather, it is a *correlate* of a complex mix of time-varying, context-specific, micro-level factors that do affect behaviors in real time at the firm level. Among these are: the adopter's time-varying perceptions of a technology's unfolding destiny; the level of implementation barriers actually experienced by adopters; and the extent to which adopters' expectations have been disappointed. Thus, when I argue that destiny serves as a moderator, this is meant in the statistical sense that the predictions in the dominant paradigm will hold more strongly for innovations with a positive destiny, rather than in the traditional causal sense. Nevertheless, the destiny concept does give an overall conceptual framework for examining presumed micro-level causes such as those just described.

Innovation Mindfulness

The dominant innovation paradigm takes a normatively rational perspective on the determinants of innovative behaviors in organizations. That is, managers are assumed to approach the decision process well-armed with information about the innovation and how well it fits with their organizational context and objectives—or at least they are expected to behave as if they were so armed. Yet, even a casual observer will see divergences from this rational ideal wherein individual organizations, or even whole populations make decisions that seem divorced from normative rationality. In fact, two of the streams already highlighted (management fashion, social contagion) are concerned specifically with understanding how apparently inefficient innovations can gain broad acceptance, at least for a time.

If we can not trust that the rational ideal will always hold, this suggests the need to open the "black box" of organizational decision making so we can examine instances where innovation does and does not conform to the rational ideal, and thereby identify the antecedents and consequences of such conformance (or lack of conformance). To support this task, researchers have begun to consider the concept of *mindfulness* (Swanson and Ramiller, 2004) and the role it plays in determining organizational innovation, particularly in the context of technologies subject to bandwagon dynamics. The concept of mindfulness suggests a more complex relationship between traditional innovation antecedents and outcomes than the ones reflected in the dominant paradigm.

Fiol and O'Connor (2003) draw on the work of Langer (1989, 1997) to define the mindfulness of an individual as a *watchful and vigilant state of mind* that is distinguished

by three characteristics: the continuous creation of new categories, openness to new information, and an implicit awareness of more than one perspective. They see mindfulness as a key to understanding the “*micro level decision context that influences whether organizational leaders will make discriminating choices that fit an organization's unique circumstances in the face of such bandwagons, or whether they will simply follow the pack*” (Fiol and O'Connor, 2003). Toward this end, they develop a framework whereby greater mindfulness among decision makers changes the way in which mechanisms for environment scanning and information processing are used. In particular, they argue that mindful managers will have more *expanded* scanning and more *context relevant* interpretations, and this will lead to more discriminating decisions in the face of bandwagons. The implication is that such organizations will be more likely to avoid bandwagon behaviors that add little or no value to the firm. Thus, while the traditional model would be concerned with predicting who would be the first to jump on an innovation bandwagon (regardless of whether the underlying technology has merit), research on mindfulness seeks to explain how some firms successfully resist bandwagons that in the end do not have merit.

In the IT domain, Swanson and Ramiller (2004) develop a comprehensive examination of the role of mindfulness in IT innovation. Building primarily on Weick *et al.* (1999), they define an organization as mindful when it *attends to innovation with reasoning grounded in its own facts and specifics*. They argue that “attention to organizational specifics is crucial in supporting sound judgments about *whether* adopting a particular innovation is a good thing to do, *when* committing to the innovation is likely best to take place, and *how* implementation and assimilation can best be pursued” (emphasis in original) (Swanson and Ramiller, 2004, p. 4). Thus, like Fiol and O'Connor, Swanson and Ramiller see mindfulness as a key antecedent of discriminating choices. However, they are more concerned with mindfulness as an organizational-level property rather than as a property of individual managers. In addition, they are interested in the interplay between mindfulness and the evolving discourse that surrounds the innovation in the larger institutional field (Swanson and Ramiller, 1997). Their work enlarges the scope of mindfulness as it pertains to organizational innovation to consider how it impacts adoption and diffusion at a community level. Accordingly, they have developed propositions linking the prevalence of mindfulness to characteristics of innovations (e.g., radicalness), characteristics of the diffusion process (e.g., rate of adoption, size of assimilation gaps), characteristics of discourse surrounding the innovation (e.g., fashionability), and characteristics of firms (e.g., recent performance, IT sophistication). They also articulate several other opportunities for incorporating mindfulness into research on innovation processes and related managerial decision making.

Measurement Challenges

In a moment I will summarize some opportunities to incorporate mindfulness into models of IT innovation, but first let's consider some measurement challenges related to the concept. In Fiol and O'Connor's formulation, mindfulness is an individual level cognitive construct, dealing as it does with things that exist in managers' minds (concept *creation*, *openness* to information, and *awareness* of multiple perspectives). Cognitive constructs are difficult to capture, however the authors do give some guidance on how this might be accomplished. In particular, they suggest that operational evidence of mindfulness could be derived from certain managerial behaviors and tendencies, such as to “make distinctions and thus create/refine categories not common to their competitors” as well as from “openness to and interest in new information as well as an awareness of

multiple possibilities not considered by competitors" (Fiol and O'Connor, 2003, p. 67). The authors also suggest measures for three processes that lead to greater mindfulness (i.e., greater reluctance to simplify, greater commitment to resilience, and greater preoccupation with both success and failure), and for the proposed outcomes of mindfulness (i.e., scanning, interpretation, and decision making processes).

Swanson and Ramiller are concerned with mindfulness as an organizational property, which is "grounded in, although not reducible to, the minds of participating individuals" (Swanson and Ramiller, 2004, p. 7). Although they give no specific guidance on potential operational definitions, presumably scales could be developed to capture the extent to which a firm considers situational specifics, such as "the innovation's ramifications for operational efficiency; the organization's preparedness for change involved; the quality and availability of complementary resources needed; implications for various common and conflicting interests, both internally and in inter-firm relationships; and the effects of adoption on the firms legitimacy with outside constituencies" (Swanson and Ramiller, 2004, p. 4).

Incorporating Mindfulness into Models of IT Innovation

Several complimentary avenues of investigation are suggested by the mindfulness concept. First, the propositions offered in the two pioneering works just described (Fiol and O'Connor, 2003; Swanson and Ramiller, 2004) have yet to receive an empirical test. These propositions could hold appeal to scholars working within a variety of research styles (micro or macro, variance or process, field study or survey).

In addition, the mindfulness concept suggests some avenues for investigation that, although not explicitly posited, are consistent with the spirit of these works. I would begin by noting that there is no indication that more mindful organizations will do more innovation overall. In fact, Fiol and O'Connor (2003) are most concerned with how mindfulness could predispose an organization to be less likely to follow the pack in joining social bandwagons surrounding innovations of dubious merit. This suggests that, in some contexts, more mindful organizations would be expected to do *less* innovation. Or, to link this to the social contagion stream, such firms should have a lower coefficient of *susceptibility*. On the other hand, it can be argued that more mindful organizations will be less likely to follow the pack in rejecting beneficial (but for whatever reason, unpopular) innovations, and thus could, in some contexts, do *more* innovation. This means that whether an innovation is beneficial or not could moderate the link between mindfulness and the propensity to innovate. Furthermore, in keeping with the view that mindful organizations will make better decisions throughout the innovation process, one could posit that mindfulness serves as an antecedent of the *quality of innovation* (to be discussed below). Alternatively, it could be posited that mindfulness will moderate the link between the *quantity of innovation* and resulting *performance impacts* (also discussed below), the rationale being that for any given amount of innovation quantity, organizations that are more mindful will have better results.

Finally, one might examine how the links between mindfulness and innovation might vary across innovation stages. Swanson and Ramiller (2004) argue that mindful organizations should be more prone to engage in early evaluations of new technology, which implies a positive link between mindfulness and this traditional measure of innovation. A corresponding implication, however, is that such organizations will be less likely to fully implement any given innovation that has been initiated. The reason for this

discrepancy is that initiation is necessary for an organization to gather the information required to ascertain whether an innovation is relevant given their own facts and specifics, but is also the mechanism that produces information needed to make discriminating choices regarding which innovations should actually be implemented from among those that have been initiated. (Interestingly, a similar pattern of more promiscuous initiation but more choosy implementation is suggested by the real options perspective on technology adoption (Fichman, 2004)).

The concept of mindfulness holds much potential to form the basis of a robust alternative stream of research on IT innovation. It directly addresses the unifying theme of IT innovation research, which is to help managers understand how they can do a better job of discerning whether, when, and how to innovate with IT. Also, owing to the high degree of uncertainty that often attends IT innovation, and the prevalence of bandwagon dynamics, IT innovations represent an ideal domain from which to select focal innovations with a good balance of both mindful and mindless responses. Finally, the perspective suggests a large variety of interesting linkages involving both traditional innovation variables and also new variables suggested by the alternative perspectives considered here.

Reconceptualizing the Dependent Variable

Previous sections have presented several new conceptual directions that identify—based on the key ideas of innovation configurations, social contagion, management fashion, technology destiny, and innovation mindfulness—different kinds of innovation antecedents and modeling approaches. The central argument has been that in many innovation contexts, we must go beyond the traditional economic-rationalistic perspective to get a fuller picture of what causes different sorts of innovative behaviors.

I now shift to the domain of innovation outcomes (depicted in the right side of Figure 2) and consider two opportunities to go beyond the traditional focus on the quantity of innovation (i.e., frequency, earliness, extent). In particular, I will examine *quality of IT innovation* and *performance impacts* as alternative candidates for the ultimate innovation outcome variable.

Quality of IT Innovation

The idea that some innovation decisions (about when, whether, and how to adopt) are better than others is a common theme, both within the dominant innovation paradigm and beyond it. In fact, the focus on identifying the profile of a likely or appropriate adopter, which characterizes much innovation research, carries with it the implication that there just might be something wrong with firms that fit the profile but fail to innovate (i.e., have they lost an opportunity?), or choose to innovative despite a lack of fit (i.e., can they really expect to benefit?).

Research within the dominant paradigm has been content to assume a positive relationship between the degree of innovation and beneficial outcomes, as per the *pro-innovation bias* already mentioned. However as this essay has argued so far, there are situations where this assumption may not hold. The streams on managerial fashion and social contagion suggest that innovations that lack merit may attract a large following, at least for a time. Furthermore, even a potentially beneficial innovation can have a poor

destiny if it loses a standards war or gets superceded by a different technology. And finally, even when an innovation holds benefits for most adopters, it may not benefit the subset that adopts mindlessly, i.e., when it is a poor fit to their organizational facts and specifics.

Thus it seems a little surprising that no attempt has been made to define some notion of the *quality of IT innovation* that captures the extent to which an organization has adopted the “right” innovation, at the “right” time and in the “right” way.⁹ In fact, a direct link from mindfulness to quality of innovation so conceived is implied by the passage quoted above linking mindfulness to sound judgments about whether, when, and how to innovate (Swanson and Ramiller, 2004, p. 4). Yet, if we are to examine linkages between quality of innovation and mindfulness or any other variables, we will need an operational measure of the concept.

Measures for Quality of Innovation

There are several options for measuring quality of innovation. One approach would be to take performance impacts as a proxy, the logic being that if a beneficial result has been achieved, then the quality of the decisions and processes surrounding the innovation effort must have been reasonably good (i.e., the proof is in the pudding). And, in fact, one article that could be found on the subject of innovation quality does include quantifiable performance impacts (e.g., return on investment, efficiency, productivity improvement) among several potential operational measures of the concept (Haner, 2002).

The use of innovation impacts as a proxy for the quality of innovation may be appropriate in some cases, but also introduces challenges, especially in the context of IT. First of all, IT innovations often produce benefits that are difficult to quantify, such as increases in managerial flexibility or contributions to the core competencies or absorptive capacities of the firm (Sambamurthy *et al.*, 2003). Furthermore, even when impacts are quantifiable in principle, it is rare for individual firms to gather the data required to accomplish such measurement in a way that would allow researchers to directly link particular innovations to particular outcomes. Owing to this rarity, investigations that employ such measures are usually limited to case studies of single firms, or a small group of firms. This perhaps explains why researchers in this domain more often use publicly available firm-level performance data. Some firm-level measures, such as Tobin’s Q, have the advantage of capturing intangible benefits of innovation that should, sooner or later, show up at the firm level (Bharadwaj *et al.*, 1999). However, such measures can only be used for innovations that are substantial enough to have enterprise-level impacts. A last limitation of performance impacts as a proxy for innovation quality is that this approach does not recognize that a firm may implement an innovation at the *wrong* time or in the *wrong* way, and still get *some* benefits, even while the majority of the potential benefits from the opportunity have been squandered.

The implication of the above is that while there may well be situations where it is useful to treat performance impacts as a proxy for innovation quality, it may also be fruitful to view innovation quality and innovation performance impacts as distinct, but related

⁹ While innovation quality is essentially a measure of the efficacy of the innovation process, it does not pertain only to process innovations. A firm can also be more adept at introducing the “right” product innovations at the “right” time and in the “right” way.

concepts. In fact, it might be most appropriate to view innovation quality as a variable that mediates (or perhaps moderates) the impacts of other innovation variables (e.g., extent of use) on performance impacts.

Measures for innovation quality that go beyond performance impacts will require some ingenuity on the part of researchers to operationalize. The following are offered as potential approaches.

Innovator Profile-Innovation Behavior Fit

Much prior research on IT innovation seeks to identify the profile of organizations that *should be* an innovator from a normative, economic-rationalistic point of view. This suggests that *innovator profile/innovation behavior fit*—a variable that captures the extent of agreement between *expected* levels of innovation (according to normative rationality) and *actual* levels of innovation—might serve as one measure of quality of innovation (especially the extent to which a firm has adopted the “right” innovation at the “right” time). The actual operational measure could be based on the error term (i.e., the difference between the fitted and actual innovation scores) produced in a regression model in which innovator profile predicts the extent of innovation. More specifically, one could calculate the absolute value of residuals and then rescale so that the smallest residuals are assigned the highest values.

This proposed measure, while relatively easy to capture, raises a concern that other kinds of error (e.g., measurement error, model specification error) may account for the vast majority of the statistical error term, rather than the kind of managerial error that signals poor quality of innovation. That is, organizations that innovate despite a *measured* lack of fit may not be making poor decisions, but rather, may possess compensating characteristics that have gone *unmeasured*. Nevertheless, one must remember that an interest in quality of innovation as an alternative to quantity of innovation arises in the first place from skepticism about whether managers always behave according to the dictates of normative rationality throughout the innovation process. Thus, in circumstances where managerial error is expected to be more prevalent, one may expect the portion of the residuals that are attributable to such error to be correspondingly larger, and hence will represent a better negative indicator of innovation quality.

Winners and Losers

The technology destiny concept, which draws a broad distinction between innovations that eventually succeed versus those that fail at the community level, suggests another avenue for developing a measure of the quality of innovation. Rather than using a simple count of innovations adopted, as with the traditional quantity oriented approach, one could develop a different kind of aggregate innovation score that takes into account the destiny of the innovations in question. In particular, a measure could be developed that gives positive weight to the propensity to adopt and implement *winners* (technologies that ultimately succeed) but a negative or neutral weight to the propensity to adopt and implement *losers* (technologies that ultimately fail). The rationale for this measure is that a key aspect of adopting the *right* technology is choosing those innovations destined for broad institutionalization at the community level, since as already argued, community adoption is the means by which most IT innovations become beneficial. The main limitation of this measure is that it will require gathering adoption and implementation data on several technologies (both winners and losers), and it can

only be used in studies that take a retrospective look at innovations for which technology destiny is already known. This will complicate the collection of variables that may be expected to predict this measure of innovation quality, due to the elapse of time since such variables would have had their presumed effects.

Quality of Use

A third avenue for measuring innovation is suggested by the *quality of use* construct developed by Boudreau and Seligman (2003). They define *quality of use* as a person's *ability to correctly exploit the appropriate capabilities of software in the most relevant circumstances*. They suggest that this should be viewed as an individual-level construct measured based on the perceptions of technology users themselves. However, it may be possible to develop an analogous organizational level measure, either by aggregating individual ratings, or by asking key informants to assess the construct at the level of the organization.

Conclusions on the Quality of Innovation

A general notion of quality of innovation underlies (at least implicitly) much prior work on innovation, yet no attempt has been made to operationalize the concept. In this section I have offered three candidate approaches to such operationalization, and have also suggested the use of performance impacts as a proxy measure. I now turn to a detailed examination of the role of performance impacts in IT innovation research.

Innovation Performance Impacts

In the end, organizations adopt IT innovations to improve some aspect of their performance. Yet, there are two ways in which the actual performance impacts of IT innovation can diverge from the normative ideal. One is that the innovation can fail to produce expected benefits, or indeed, any benefits at all. A second is that the innovation could produce some benefits, but not enough to recover the costs of implementation. These possibilities suggest more attention going forward to studies that examine performance impacts as a key outcome of IT innovation.

A sizeable stream of research has examined the more general domain of the business value of IT investments (see Melville *et al.* (2004) for a review). Early studies failed to find the expected link between enormous increases in IT capital investments on the one hand and productivity improvement on the other, leading to the so-called "productivity paradox" (Brynjolfsson, 1993). However, subsequent studies established that, in general, investments in IT capital do produce net efficiency benefits, although this varies depending on other factors such as management practices, and organizational and industry structure (Bresnahan *et al.*, 2002). Also there is no assurance the investing firm, rather than customers or competitors, will capture the value of those efficiency improvements (Hitt and Brynjolfsson, 1996).

Within this broader stream, a growing number of studies have examined individual applications of IT that were either explicitly characterized as innovative, or could have been. These have included studies of automated teller machines (Dos Santos and Peffers, 1995), enterprise resource planning, (Hitt *et al.*, 2001; McAfee, 2002), electronic data interchange (Mukhopadhyay *et al.*, 1995) and hospital information systems (Devaraj and Kohli, 2003; Kohli and Devaraj, 2003). Yet most of these studies, even though they concerned adoption of innovative technologies, were not intended to

address the fundamental question of *whether, when, and how* firms should innovate with IT, and thus fall outside the domain of IT innovation research as defined here. For example, while Mukhopadhyay *et al.*'s (1995) examination of EDI at Chrysler illustrates how to measure the value of a particular implementation, it was not intended to answer the question whether implementation occurred at the right time or in the right way, and therefore gives little insight into how to manage IT innovation.

To bring performance impacts of innovative IT under the innovation research umbrella requires that researchers link innovation-related variables to performance impacts. So, the simple fact that adoption of some innovation has lead to business value would not be particularly germane from an innovation research perspective. Rather, innovation researchers would be concerned with what *innovation-related conditions* have lead to *differences in the degree of business value produced*. Figure 3 depicts three generic assumptions for how innovation-related conditions lead to performance impacts.

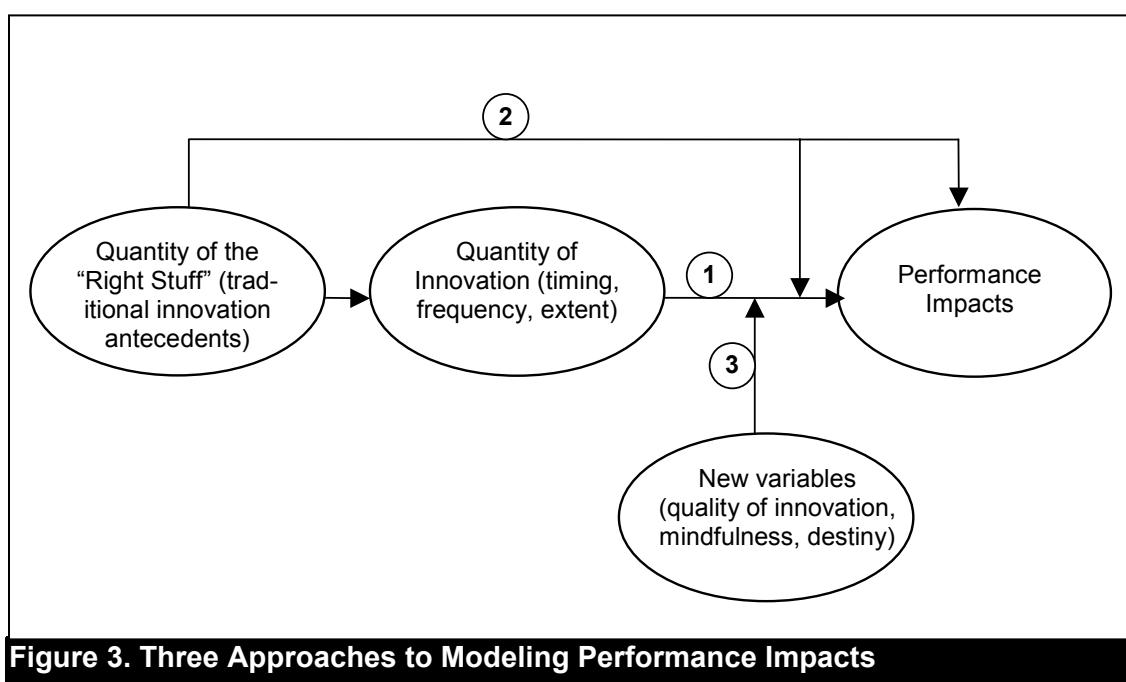


Figure 3. Three Approaches to Modeling Performance Impacts

Assumption 1: Innovation Directly Causes Performance

The first approach (depicted as Arrow 1 in Figure 3), is to assume that the quantity of innovation directly determines performance impacts. Of those few studies of performance impacts that fall within the scope of IT innovation research as defined here, most have employed this approach and have found a positive link between quantity of innovation and performance impacts. For example, Dos Santos and Peffers (1995) show that banks who were earlier adopters of ATMs increased their market share and net income relative to later adopters. This research links a traditional innovation outcome variable—adoption timing—to performance outcomes, and thus provides insight into the question of *when* to innovate with IT. In a more recent study, Hitt *et al.* (2001) examined the performance of ERP adopters versus non-adopters, and also examined whether performance was affected by *how* firms implemented ERP. Specifically, they found that organizations with a greater breadth of use achieved greater benefits—thus linking a common measure of innovation quantity to performance. These two studies employed

objective measures of performance. Some other studies with a similar intent to link quantity of innovation directly to performance have used perceptual performance measures (Ramamurthy *et al.*, 1999; Subramanian and Nilakanta, 1996). Both studies developed two-stage models where a number of traditional innovation antecedents were assumed to increase quantity of innovation (e.g., number of adoptions, adoption timing), which, in turn, increased performance. As discussed below, perceptual measures of performance, while comparatively easy to capture, can be prone to some rather serious methodological limitations in the context of innovation research.

Assumption 2: Traditional Innovation Antecedents Have Direct or Moderating Effects on Performance

A second generic approach to modeling performance impacts is to assume that traditional innovation antecedents are not fully mediated by innovation quantity-oriented outcome variables, but rather, that they have direct or moderating effects over and above their indirect effects working through innovation quantity (see Arrow 2 in Figure 3). As just one example, it might be argued that absorptive capacity not only increases the amount of innovation itself, but also facilitates the translation of any given amount of innovation into positive performance impacts.

One interesting approach for testing direct linkages between innovation antecedents and performance is to analyze the organizational learning curves associated with new technology. Kemerer (1992) first proposed applying learning curves to study the benefits associated with CASE, an innovation in software process technology. Learning curve models allow researchers to estimate how performance improves with cumulative experience, and how other variables affect this rate of improvement. So, one could hypothesize that certain characteristics predispose an organization to progress along the learning curve more rapidly (e.g., some of the same sorts of capability-oriented characteristics that predict early adoption), and then gather data to test for these relationships. Because cumulative experience is controlled for, this means that one measure of innovation quantity—extent of use—is already built into the model. While no studies of IT innovation have used the learning curve model, a recent study of the adoption of a novel surgical technique provides an excellent illustration of how to link work on learning curves and organizational innovation (Edmondson *et al.*, 2003; Pisano *et al.*, 2001).

Assumption 3: Novel Variables Have Moderating Effects on Performance

A third generic approach is to posit that some of the new constructs that exist outside the bounds of the dominant paradigm will serve to moderate the relationship between innovation quantity and performance. One alternative here would be to examine whether the innovator profile/innovation behavior fit (which might be viewed as a measure of the quality of innovation) plays a moderating role. If the implicit assumption that organizations that better fit the innovator profile (i.e., they have the "Right Stuff") should be expected to have higher returns due to this fit is actually true, then it follows that organizations that innovate *despite a lack of fit* should expect lesser returns.

Relatedly, we could hypothesize that organizations that exhibit greater innovation mindfulness will have more beneficial impacts for any given level of innovation. Because mindful organizations make innovation decisions that are more strongly grounded in organizational facts and specifics, this implies that there will be stronger complementarities between the innovation and other organizational aspects (needs,

capabilities, strategies). Examination of this hypothesis provides an opportunity for interdisciplinary research joining economics and behavioral science.

As a third alternative within this generic approach, it seems clear that the destiny of an adopted innovation would often moderate the level of benefit achieved from a given level of adoption. In the case of networked technologies, this follows virtually by definition because greater acceptance of the innovation (an element of a positive destiny) improves its utility.

Measurement Challenges

As with the stream on IT business value more generally, the measurement of performance impacts in the innovation context poses some significant challenges. The nature of these challenges varies depending on which kind of measure is used. Below I consider both perceptual and objective measures of performance.

Perceptual Measures

Perceptual measures facilitate the construction of more comprehensive multi-stage models, such as those already mentioned (Ramamurthy *et al.*, 1999; Subramanian and Nilakanta, 1996). However, such measures may be especially prone to certain methodological problems when investigating direct innovation-performance linkages due to: (1) the pro-innovation bias, and (2) the possibility respondents are involved with or responsible for the innovative effort itself.

The pro-innovation bias refers to the assumption that innovation is "good." The linking of one "good" variable (innovation) to another "good" variable (improved performance) should tend to magnify concerns about social desirability bias and hypothesis guessing, especially in single respondent studies. Those people most prone to social desirability bias will tend to over-report both innovation and performance, thus leading to exaggerated estimates of relationships between the two. Likewise, the ease of guessing the expected relationship between innovation and performance should lead to exaggerated estimates.

Further magnifying the potential for bias is the common practice of using respondents that were responsible for bringing an innovation into the organization and/or ensuring its future success. This problem will be most severe when researchers use convenience samples taken from professional groups or conference attendee lists that are directly associated with the innovation in question. But even when the respondents are taken from the ranks of IT managers in a random sample of organizations, it can be assumed that many of those managers would have had some responsibility for bringing the innovation into the organization. When respondents that are responsible for an intervention are asked to assess performance in areas that the innovation is expected to effect, this increases the risk that self-justification bias will enter into their performance judgments, leading to over-reporting of performance impacts.

Thus, it would appear that when perceptual performance measures are used it will be especially crucial to either use separate respondents to judge performance antecedents and performance, or to employ research designs where biases arising from single respondents are less of a threat. Regarding the latter, relationships that are more subtle (such as moderating effects) should be less prone to the hypothesis guessing.

Objective Measures

Objective measures of performance, which have dominated in the larger stream on the business value of IT, fall into four broad categories: process-level measures, firm-level productivity measures, firm-level accounting measures, and firm-level market measures. A brief summary of these measures is provided below.

Process-level performance measures are those that capture the effects of an innovation on the efficiency or effectiveness of the specific organizational processes they support. So, for example, McAfee (2002) documents the effects of ERP adoption on lead time and on-time delivery. The advantage of such measures is that they can isolate the direct tangible effects of innovations in use. The disadvantages are that they omit intangible and indirect effects (both positive and negative). Also, such measures are only infrequently captured by firms (thus confining most studies to small convenience samples), and do not account for the costs of implementation (which leaves unanswered whether benefits have exceeded the costs).

The **firm-level productivity** benefits of IT investments are usually investigated using the production function specification from the theory of the firm. This framework posits that output (e.g., sales) is a function of how a firm combines the fundamental inputs of labor and capital. These broad inputs can be further divided to examine more specific research questions. For example, Brynjolfsson and Hitt (1996) separate out IT capital and IT labor to examine the specific contributions of those variables to output. Although it appears the production function framework has yet to be used in an innovation context, one could, in principle, distinguish innovative IT capital from IT capital just as prior IT business value studies have separated out IT capital from all capital.

Firm-level accounting measures of performance include profitability-related measures (e.g., return on assets, return on sales, and return on equity) and cost measures (e.g., SGA expenses/sales). In the general IT business value literature, Bharadwaj (2000) uses a matched pair analysis to show that IT leaders outperform on several profitability measures and some cost measures. In the domain of IT innovation, Hitt *et al.* (2001) link ERP adoption to several profitability variables and other performance measures (e.g., labor productivity).

Firm-level accounting measures have the advantage of tapping directly into well-accepted notions of firm performance, but also have several disadvantages. These measures relate to the past rather than the future; they are not adjusted for risk; they are subject to temporary shifts in tax laws or accounting conventions; and they do not account for performance lags (Bharadwaj *et al.*, 1999). As a result, accounting profitability represents a somewhat noisy measure of the true economic returns from IT innovation. Also, they only can be used for innovations substantial enough to have measurable firm-level effects.

The fourth and final category of performance measures consider the market response to innovation. **Market-based measures** have the advantage of being forward looking, risk adjusted and less sensitive to accounting rule changes (see Bharadwaj *et al.* (1999) for a detailed discussion). However, as with accounting measures, these can only be used for innovations large enough to have firm-level impacts. In the IT business value literature, many studies have examined the market response to IT investment announcements, the idea being that if the market collectively judges that the investment

will be beneficial on net, then they will upwardly adjust the assessment of the company's value, thus leading to abnormal returns. An alternative market-based approach is to examine a firm's Tobin's Q ratio, which is defined as the ratio of a firm's market valuation to the cost of replacing all of its assets. This measure is especially promising for innovation value studies, in that it captures the extent to which a firm's value is based on intangibles, and an innovation often produces a majority of its benefits in the domain of intangibles.

Summary words

Researchers interested in the business value of IT are increasingly coming around to the idea that the central question is not whether IT investments pay off, but rather "*under what conditions* do investments in IT pay off"? (Dehning et al., 2003, p. 638). A similar statement can be made concerning the performance impacts of investments in innovative IT. This sensitivity to contextual conditions could lead to increasing interest among business value researchers in examining specific information technologies, and particularly those that are just emerging. Meanwhile, IT innovation researchers have developed a rich cumulative tradition of examining the contextual factors that affect innovation, and potentially, the performance impacts of innovation. This suggests excellent opportunity for collaboration across these two traditions.

Summary and Conclusions

This essay has presented seven emerging perspectives that go beyond the dominant paradigm's focus on explaining the quantity of IT innovation with economic-rationalistic arguments. These perspectives can be summarized by their central research questions, articulated as follows:

- Innovation Configurations: Which holistic combinations of factors explain IT innovation outcomes on large-scale deployment efforts?
- Social Contagion: When are the forces of social contagion the strongest and what variables carry the contagion effect?
- Management Fashion: What triggers the emergence of IT fashions and determines whether a fashion will be transient or become an enduring institution?
- Innovation Mindfulness: What characterizes more mindful organizations, and how do their innovation outcomes differ from less mindful organizations?
- Technology Destiny: How do the determinants of the rate/extent of innovation implementation vary depending on a technology's destiny?
- Quality of Innovation: What characterizes organizations that exhibit greater quality of innovation, and how does this quality relate to innovation impacts?
- Performance Impacts: Under what contextual conditions are the performance impacts of innovative IT the greatest?

Three trends favor the new perspectives advocated here. First, the 1990s ushered in a new emphasis on enterprise-scale IT innovation initiatives, including corporate reengineering, ERP, customer relationship management, knowledge management, and

general efforts to “digitize” the firm using Internet technologies. The scale and complexity of these initiatives is especially compatible with a configurational view of causal factors, and also opens up the opportunity to study innovation performance impacts at the firm level, where a richer array of publicly-available data resides.

Another trend that gained traction in the 1990s is skepticism about whether the massive increases in corporate investment in IT have produced positive payoffs. This skepticism strikes at the heart of the pro innovation bias on which so much of the dominant paradigm depends. If it can not be safely assumed that innovation is beneficial for all adopters at all times, this increases interest in explicitly examining innovation value by incorporating quality of innovation or performance impacts into our research. It also increases interest in perspectives—such as social contagion, management fashion, and innovation mindfulness—that do not rely on the assumption that innovation is beneficial, and, in fact, posit mechanisms for how innovations that are not beneficial can nevertheless become widely adopted.

A third trend is an increase in the pervasiveness of bandwagon phenomenon, and perhaps not coincidentally, a quickening in the pace in the rise and fall of managerial fashions (Carson *et al.*, 2000). The underlying drivers for this trend could include the increasing importance of technical standards, an increasing emphasis on interconnected IT systems (both within and across firms), and a general increase in environmental uncertainty due to accelerating rate of change. Also, the increasing prominence of a new class of business periodicals with an unabashed championing of innovation (e.g., *Business 2.0*, *MIT Technology Review*, *Fast Company*, *Wired*) may be seen as resulting from and contributing to the rise of bandwagons. As one more example of this trend, consider Gartner Group’s success in building a consulting practice surrounding their “Hype Cycle” concept, which portrays hype-driven bandwagons as the norm, rather than the exception, for emerging IT. Most of the new perspectives advocated here increase in salience in the presence of bandwagons and “hype.” Social contagion and management fashion both contribute to and feed off of bandwagon dynamics. Pressure to innovate mindlessly by following the crowd is greatest in the presence of bandwagons. And finally, bandwagons tend to result in the sort of extreme outcomes (pervasive adoption or wholesale abandonment) that make technology destiny a more compelling consideration.

The new perspectives articulated here not only hold promise for reinvigorating innovation concepts, but also suggest a variety of new research methods that take us beyond the cross-sectional sample surveys and traditional regression or covariance-based statistical models that have dominated in prior IT innovation research. The QCA method provides a radical new approach to modeling the causal complexity of large-scale innovation phenomenon based on intensive, holistic examination of individual cases. The heterogeneous diffusion model allows explicit specification of both traditional and contagion-carrying variables in a single statistical model that, unlike conventional regression models, allows explanatory factors to vary over time. Emerging bibliographic analysis techniques permit the examination of innovation-related discourse, an element that most researchers would acknowledge as important, but that can not be captured with traditional survey methods. In the domain of performance impacts, the learning curve and production function models represent powerful tools that have yet receive attention in the IT innovation context.

While there is still much good work to be done working within the dominant paradigm, it

is going to be increasingly difficult to make ground-breaking contributions. Research that goes beyond the dominant paradigm holds more promise to tell us things about the IT innovation phenomenon that we do not already know. Researchers will be using novel constructs (e.g., susceptibility to contagion, fashionability of discourse, destiny, mindfulness, quality of innovation, performance impacts) in combination with novel methods (e.g., QCA, heterogeneous diffusion models, discourse analysis models, learning curve models, production function) to examine more subtle theoretical relationships. While the challenges in doing so will be quite high, the alternative of staying too firmly ensconced within the dominant paradigm risks an ebbing away of interest in what we as a field have to say on the subject of whether, when, and how to innovate with IT.

References

- Abrahamson, E. (1991) "Managerial fads and fashions: The diffusion and rejection of innovations," *Academy of Management Review* (16)3, pp. 586-612.
- Abrahamson, E. (1996) "Management Fashion," *Academy of Management Review* (21)1, pp. 254-285.
- Abrahamson, E., and G. Fairchild (1999) "Management fashion: Lifecycles, triggers, and collective learning processes," *Administrative Science Quarterly* (44)4, pp. 708-740.
- Abrahamson, E., and L. Rosenkopf (1997) "Social network effects on the extent of innovation diffusion: A computer simulation," *Organization Science* (8)3, pp. 289-309.
- Armstrong, C.P., and V. Sambamurthy (1999) "Information technology assimilation in firms: The influence of senior leadership and IT infrastructures," *Information Systems Research* (10)4, pp. 304-327.
- Arthur, W.B. (1996) "Increasing returns and the new world of business," *Harvard Business Review* (74)4, pp. 101-109.
- Austin, R.D., R.L. Nolan, and M. Cottleer (1999) "Cisco Systems, Inc.: Implementing ERP," Harvard Business School Press, Case #9-699-022, Boston.
- Bandura, A. (1977) *Social Learning Theory* Prentice-Hall, Englewood Cliffs, NJ.
- Benders, J., and K. van Veen (2001) "What's in a fashion? interpretative viability and management fashions," *Organization* (8)1, pp. 33-53.
- Bharadwaj, A.S. (2000) "A resource-based perspective on information technology capability and firm performance: An empirical investigation," *MIS Quarterly* (24)1, pp. 169-196.
- Bharadwaj, A.S., S.G. Bharadwaj, and B.R. Konsynski (1999) "Information technology effects on firm performance as measured by Tobin's q," *Management Science* (45)7, pp. 1008-1024.
- Bijker, W. (1995) *Of Bicycles, Bakelite, and Bulbs: Towards a Theory of Sociotechnical Change* MIT Press, Cambridge, MA.
- Blumer, H. (1969) "Fashion: From class differentiation to social selection," *Sociological Quarterly* (10, pp. 275-291.
- Boswell, T., and C. Brown (1999) "The scope of general theory - Methods for linking deductive and inductive comparative history," *Sociological Methods & Research* (28)2, pp. 154-185.
- Boudreau, M.-C., and L. Seligman (2003) "A learning-based model of quality of use: Insights from a case study of ERP implementation," Diffusion Interest Group in Information Technology Workshop, Seattle.

- Bowen, M.G. (1987) "The escalation phenomenon reconsidered: Decision dilemmas or decision errors?," *Academy of Management Review* (12)1, pp. 52-66.
- Brancheau, J.C., and J.C. Wetherbe (1990) "The adoption of spreadsheet software: Testing innovation diffusion theory in the context of end-user computing," *Information Systems Research* (1)2, pp. 115-143.
- Bresnahan, T., E. Brynjolfsson, and L.M. Hitt (2002) "Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence," *Quarterly Journal of Economics* (117)1, pp. 339-376.
- Brockner, J. (1992) "The escalation of commitment to a failing course of action: Toward theoretical progress," *Academy of Management Review* (17)1, pp. 39-61.
- Brynjolfsson, E. (1993) "The productivity paradox of information technology," *Communications of the ACM* (36)12, pp. 66-77.
- Brynjolfsson, E., and L. Hitt (1996) "Paradox lost? Firm-level evidence on the returns to information systems spending," *Management Science* (42)4, pp. 541-558.
- Carson, P.P., P.A. Lanier, K.D. Carson, et al. (2000) "Clearing a path through the management fashion jungle: Some preliminary trailblazing," *Academy of Management Journal* (43)6, pp. 1143-1158.
- Clark, J., and K. Guy (1998) "Innovation and competitiveness: A review," *Technology Analysis & Strategic Management* (10)3, pp. 363-395.
- Cohen, W.M., and D.A. Levinthal (1990) "Absorptive capacity: A new perspective on learning and innovation," *Administrative Science Quarterly* (35)1, pp. 128-152.
- Cool, K.O., I. Diericks, and G. Szulanski (1997) "Diffusion of innovations within organizations: Electronic switching in the Bell System, 1971-1982," *Organization Science* (8)5, pp. 543-559.
- Cooper, R.B., and R.W. Zmud (1990) "Information technology implementation research: A technological diffusion approach," *Management Science* (36)2, pp. 123-139.
- Coverdill, J.E., and W. Finlay (1995) "Understanding mills via Mill-type methods: An application of qualitative comparative analysis to a study of labor management in southern textile manufacturing," *Qualitative Sociology* (18)4, pp. 457-478.
- Damanpour, F. (1991) "Organizational innovation: A meta-analysis of effects of determinants and moderators," *Academy of Management Journal* (34)3, pp. 555-590.
- Damsgaard, J., and K. Lyytinen (2001) "The role of intermediating institutions in the diffusion of electronic data interchange (EDI): How industry associations intervened in Denmark, Finland, and Hong Kong," *Information Society* (17)3, pp. 195-211.
- Dehning, B., V.J. Richardson, and R.W. Zmud (2003) "The value relevance of announcements of transformational information technology investments," *MIS Quarterly* (27)4, pp. 637-656.
- DeSanctis, G., and M.S. Poole (1994) "Capturing the complexity in advanced technology use: Adaptive structuration theory," *Organization Science* (5)2, pp. 121-147.
- Devaraj, S., and R. Kohli (2003) "Performance impacts of information technology: Is actual usage the missing link?," *Management Science* (49)3, pp. 273-289.
- DiMaggio, P.J., and W.W. Powell (1983) "The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields," *American Sociological Review* (48), pp. 147-160.
- Dos Santos, B.L., and K. Peffers (1995) "Rewards to investors in innovative information technology applications: First movers and early followers in ATMs," *Organization Science* (6)3, pp. 241-259.
- Downs, G.W., and L.B. Mohr (1976) "Conceptual issues in the study of innovation," *Administrative Science Quarterly* (21)4, pp. 700-714.

- Edmondson, A.C., A.B. Winslow, R.M.J. Bohmer, et al. (2003) "Learning how and learning what: Effects of tacit and codified knowledge on performance improvement following technology adoption," *Decision Sciences* (34)2, pp. 197-223.
- Eveland, J.D., and L.G. Tornatzky (1990) "The Deployment of Technology," in: *The Processes of Technological Innovation*, L.G. Tornatzky and M. Fleischer (eds.), Lexington Books, Lexington, Massachusetts, pp. 117-148.
- Fichman, R.G. (2000) "The diffusion and assimilation of information technology innovations," in: *Framing the Domains of IT Management: Projecting the Future Through the Past*, R.B. Zmud (ed.), Pinnaflex Educational Resources, Inc., Cincinnati, OH.
- Fichman, R.G. (2001) "The role of aggregation in the measurement of IT-related organizational innovation," *MIS Quarterly* (25)4, pp. 201-429.
- Fichman, R.G. (2003) "How does technology destiny affect organizational innovation?," Academy of Management Conference, Seattle, WA.
- Fichman, R.G. (2004) "Real options and IT platform adoption: Implications for theory and practice," *Information Systems Research* (15)2, pp. 132-154.
- Fichman, R.G., and C.F. Kemerer (1997) "The assimilation of software process innovations: An organizational learning perspective," *Management Science* (43)10, pp. 1345-1363.
- Fichman, R.G., and C.F. Kemerer (1999) "The illusory diffusion of innovation: An examination of assimilation gaps," *Information Systems Research* (10)3, pp. 255-275.
- Fiol, C.M., and E.J. O'Connor (2003) "Waking Up! Mindfulness in the Face of Bandwagons," *Academy of Management Review* (28)1, pp. 54-70.
- Gallivan, M. (2001) "Organizational adoption and assimilation of complex technological innovations: Development and application of a new framework," *Data Base* (32)3, pp. 51-85.
- Greve, H.R. (1995) "Jumping ship: The diffusion of strategy abandonment," *Administrative Science Quarterly* (40)3, pp. 444-473.
- Greve, H.R., D. Strang, and N.B. Tuma (1995) "Specification and estimation of heterogeneous diffusion models," *Sociological Methodology* (25), pp. 377-420.
- Grover, V., K. Fiedler, and J. Teng (1997) "Empirical evidence on Swanson's tri-core model of information systems innovation," *Information Systems Research* (8)3, pp. 273-287.
- Hamel, G. (1998) "The challenge today: Changing the rules of the game," *Business Strategy Review* (9)2, pp. 19-26.
- Haner, U.-E. (2002) "Innovation quality--A conceptual framework," *International Journal of Production Economics* (80), pp. 31-37.
- Hannan, M.T., and J. Freeman (1977) "The population ecology of organizations," *American Journal of Sociology* (82)929-964.
- Hart, P.J., and C.S. Saunders (1998) "Emerging electronic partnerships: Antecedents and dimensions of EDI use from the supplier's perspective," *Journal of Management Information Systems* (14)4, pp. 87-111.
- Hitt, L.M., and E. Brynjolfsson (1996) "Productivity, business profitability and consumer surplus: Three different measures of information technology value," *MIS Quarterly* (20)2, pp. 121-142.
- Hitt, L.M., D.J. Wu, and X. Zhou (2001) "Investment in enterprise resource planning: Business impact and productivity measures," *Journal of Management Information Systems* (19)1, pp. 71-98.
- Iacobou, C.L., I. Benbasat, and A.S. Dexter (1995) "Electronic data interchange and

- small organizations: Adoption and impact of technology," *MIS Quarterly* (19)4, pp. 465-485.
- Ives, B., J. Valacich, R.T. Watson, et al. (2003) "What every business student needs to know about information systems," *Communications of the AIS* (9) pp. 467-477.
- Keil, M., J. Mann, and R. Arun (2000) "Why software projects escalate: An empirical analysis and test of four theoretical models," *MIS Quarterly* (24)4, pp. 631-664.
- Kemerer, C.F. (1992) "How the learning curve affects CASE tool adoption," *IEEE Software* (9)3, pp. 23-28.
- Kimberly, J.R. (ed.) (1981) *Managerial innovation*. Oxford University Press, New York, 1981.
- Kitchener, M., M. Beynon, and C. Harrington (2002) "Qualitative comparative analysis and public services research: Lessons from an early application," *Public Management Review* (4)4, pp. 485-204.
- Kohli, R., and S. Devaraj (2003) "Measuring information technology payoff: A meta-analysis of structural variables in firm-level empirical research," *Information Systems Research* (14)2, pp. 127-145.
- Langer, E.J. (1989) "Minding matters: The consequences of mindlessness-mindfulness," in: *Advances in Experimental Social Psychology*, B. L. (ed.), Academic Press, pp. 137-173.
- Langer, E.J. (1997) *The Power of Mindful Learning* Addison-Wesley, Reading, MA.
- Markus, M.L., and D. Robey (1988) "Information technology and organizational change: Causal structure in theory and research," *Management Science* (34)5, pp. 583-598.
- McAfee, A. (2002) "The impact of enterprise information technology adoption on operational performance: An empirical investigation," *Production and Operations Management* (11)1, pp. 33-53.
- Melville, N., K. Kraemer, and V. Gurbaxani (2004) "Review: Information technology and organizational performance: An integrative model of IT business value," *MIS Quarterly* (28)283-322.
- Meyer, A.D., and J.B. Goes (1988) "Organizational assimilation of innovations: A multilevel contextual analysis," *Academy of Management Journal* (31)4, pp. 897-923.
- Miller, D., and J. Hartwick (2002) "Opinion - Spotting management fads," *Harvard Business Review* (80)10, pp. 26-27.
- Mukhopadhyay, T., S. Kekre, and S. Kalathur (1995) "Business value of information technology: A study of electronic data interchange," *MIS Quarterly* (19)2, pp. 137-156.
- Nadler, D.A., and M.L. Tushman (1999) "The organization of the future: Strategic imperatives and core competencies for the 21st century," *Organizational Dynamics* (28)1, pp. 45-60.
- Nutt, P.C. (2002) "Making Strategic Choices," *Journal of Management Studies* (39)1, pp. 67-96.
- Orlikowski, W.J. (1992) "The duality of technology: Rethinking the concept of technology in organizations," *Organization Science* (3)3, pp. 398-427.
- Pisano, G.P., R.M.J. Bohmer, and A.C. Edmondson (2001) "Organizational differences in rates of learning: Evidence from the adoption of minimally invasive cardiac surgery," *Management Science* (47)6, pp. 752-768.
- Premkumar, G., K. Ramamurthy, and S. Nilakanta (1994) "Implementation of electronic data interchange: An innovation diffusion perspective," *Journal of Management Information Systems* (11)2, pp. 157-186.
- Prescott, M.B., and S.A. Conger (1995) "Information technology innovations: A

- classification by IT locus of impact and research approach," *Database* (26)2&3, pp. 20-41.
- Purvis, R.L., V. Sambamurthy, and R.W. Zmud (2001) "The assimilation of knowledge platforms in organizations: An empirical investigation," *Organization Science* (12)2, pp. 117-135.
- Ragin, C.C. (1987) *The Comparative Method: Moving Beyond Qualitative and Quantitative Strategies* University of California Press, Berkeley, CA.
- Ragin, C.C. (1999) "Using qualitative comparative analysis to study causal complexity," *Health Services Research* (34)5, pp. 1225-1239.
- Rai, A., and D.S. Bajwa (1997) "An empirical investigation into factors relating to the adoption of executive information systems: An analysis of EIS for collaboration and decision support," *Decision Sciences* (28)4, pp. 939-974.
- Ramamurthy, K., G. Premkumar, and M.R. Crum (1999) "Organizational and interorganizational determinants of EDI diffusion and organizational performance: A causal model," *Journal of Organizational Computing and Electronic Commerce* (9)4, pp. 253-285.
- Rogers, E.M. (1995) *Diffusion of Innovations*, (4th edition ed.) The Free Press, New York.
- Rowan, B. (1977) "Organizational structure and the institutional environment: The case of public schools," *Administrative Science Quarterly* (27)2, pp. 259-279.
- Sambamurthy, V., A. Bharadwaj, and V. Grover (2003) "Shaping agility through digital options: Reconceptualizing the role of information technology in contemporary firms," *MIS Quarterly* (27)2, pp. 237-263.
- Schmidt, J.B., and R.J. Calantone (2002) "Escalation of commitment during new product development," *Journal of the Academy of Marketing Science* (30)2, pp. 103-118.
- Solow, R.M. (1957) "Technical change and the aggregate production function," *Review of Economics and Statistics* (39), pp. 312-320.
- Strang, D., and N.B. Tuma (1993) "Spatial and temporal heterogeneity in diffusion," *American Journal of Sociology* (99)3, p 614.
- Subramanian, A., and S. Nilakanta (1996) "Organizational innovativeness: Exploring the relationship between organizational determinants of innovation, types of innovations, and measures of organizational performance," *Omega-International Journal of Management Science* (24)6, pp. 631-647.
- Swanson, E.B. (1994) "Information systems innovations among organizations," *Management Science* (40)9, pp. 1069-1092.
- Swanson, E.B., and N. Ramiller (2004) "Innovating mindfully with information technology," Working Paper, The Anderson School at UCLA, Los Angeles.
- Swanson, E.B., and N.C. Ramiller (1997) "The organizing vision in information systems innovation," *Organization Science* (8)5, pp. 458-474.
- Tan, S.L., and M. Fichman (2002) "Adoption of web-based transactional banking: Efficiency-choice and neo-institutional perspectives," International Conference on Information Systems, Barcelona.
- Teo, H.H., K.K. Wei, and I. Benbasat (2003) "Predicting intention to adopt interorganizational linkages: An institutional perspective," *MIS Quarterly* (27)1, p 31.
- Tornatzky, L.G., and K. Klein, J. (1982) "Innovation characteristics and innovation adoption-implementation: A meta-analysis of findings," *IEEE Transactions on Engineering Management* (EM-29)1, pp. 28-45.
- Venkatesh, V., M.G. Morris, G.B. Davis, et al. (2003) "User acceptance of information technology: Toward a unified view," *MIS Quarterly* (27)3, pp. 425-478.
- Weick, K.E., K.M. Sutcliffe, and D. Obstfeld (1999) "Organizing for high reliability:

- Processes of collective mindfulness," *Research in Organizational Behavior*, Vol. 21, 1999 (21), pp. 81-123.
- Wolfe, R.A. (1994) "Organizational innovation: Review, critique and suggested research directions," *Journal of Management Studies* (31)3, pp. 405-431.

About the authors

Robert Fichman is an associate professor of IT in the Department of Operations, Information, and Strategic Management in the Carroll School of Management at Boston College. His research interests include IT diffusion, management of innovation, and real options. He has published broadly on these subjects in CACM, IEEE Computer, Information Systems Research, Management Science, MIS Quarterly, Sloan Management Review and other scholarly journals. Dr. Fichman holds BS and MS degrees in Industrial & Operations Engineering from the University of Michigan and a Ph.D. in Management from the MIT Sloan School of Management.

Copyright © 2003 by the **Association for Information Systems**. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and full citation on the first page. Copyright for components of this work owned by others than the Association for Information Systems must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or fee. Request permission to publish from: AIS Administrative Office, PO Box 2712 Atlanta, GA, 30301-2712, Attn: Reprints, or via e-mail from ais@aisnet.org.



Journal of the Association for Information Systems

ISSN: 1536-9323

EDITOR
Sirkka L. Jarvenpaa
University of Texas at Austin

JAIS SENIOR EDITORS

Soon Ang Nanyang Technological University	Izak Benbasat University of British Columbia	Matthias Jarke Technical University of Aachen
Kalle Lyytinen Case Western Reserve University	Tridas Mukhopadhyay Carnegie Mellon University	Robert Zmud University of Oklahoma

JAIS EDITORIAL BOARD

Ritu Agarwal University of Maryland	Paul Alpar University of Marburg	Anandhi S. Bharadwaj Emory University	Yolande E. Chan Queen's University
Alok R. Chaturvedi Purdue University	Roger H.L. Chiang University of Cincinnati	Wynne Chin University of Houston	Ellen Christiaanse University of Amsterdam
Alan Dennis Indiana University	Amitava Dutta George Mason University	Robert Fichman Boston College	Henrique Freitas Universidade Federal do Rio Grande do Sul
Guy G. Gable Queensland University of Technology	Rudy Hirschheim Louisiana State University	Juhani Iivari University of Oulu	Matthew R. Jones University of Cambridge
Elena Karahanna University of Georgia	Robert J. Kauffman University of Minnesota	Prabhudev Konana University of Texas at Austin	Kai H. Lim City University of Hong Kong
Claudia Loebbecke University of Cologne	Mats Lundeberg Stockholm School of Economics	Stuart E. Madnick Massachusetts Institute of Technology	Ann Majchrzak University of Southern California
Ryutaro Manabe Bunkyo University	Anne Massey Indiana University	Eric Monteiro Norwegian University of Science and Technology	B. Jeffrey Parsons Memorial University of Newfoundland
Nava Pliskin Ben-Gurion University of the Negev	Jan Pries-Heje Copenhagen Business School	Arun Rai Georgia State University	Sudha Ram University of Arizona
Suzanne Rivard Ecole des Hautes Etudes Commerciales	Rajiv Sabherwal University of Missouri – St. Louis	Christopher Sauer Oxford University	Peretz Shoval Ben-Gurion University
Sandra A. Slaughter Carnegie Mellon University	Christina Soh Nanyang Technological University	Ananth Srinivasan University of Auckland	Kar Yan Tam Hong Kong University of Science and Technology
Bernard C.Y. Tan National University of Singapore	Dov Te'eni Bar-Ilan University	Yair Wand University of British Columbia	Richard T. Watson University of Georgia
Gillian Yeo Nanyang Business School	Youngjin Yoo Case Western Reserve University		

ADMINISTRATIVE PERSONNEL

Eph McLean AIS, Executive Director Georgia State University	Samantha Spears Subscriptions Manager Georgia State University	Reagan Ramsower Publisher, JAIS Baylor University
---	--	---