

When What is Useful is Not Necessarily True: The Underappreciated Conceptual Scheme

T. Grandon Gill
University of South Florida, Tampa, FL, USA

grandon@usf.edu

Abstract

A major emphasis of research in informing science and its contributing disciplines is the development of theory. As it is most commonly used, the term theory refers to our attempt to describe—in symbolic terms—some underlying truth. Through developing such descriptions, we seek to improve our understanding of phenomena that occur in the natural or social world and, perhaps, enhance our ability to make predictions of such phenomena.

Theory creation and validation, however, are not the only possible goals for research. It is also possible for research to focus on establishing conceptual schemes. These schemes are models that can be used to think about the same types of phenomena described by theory, but are not assumed to be representations of truth. Instead, conceptual schemes are evaluated based upon their usefulness to a client.

This research essay proposes that, as the complexity of the environments we seek to understand and control grows, the goals of “truth” and “usefulness” tend to diverge. Where this occurs, an obsession with truth can lead to research impotence: we are motivated to validate truths that are already widely accepted, our theory becomes too convoluted to be applied or communicated, and we are prone to becoming infatuated with the nobility of our quest—while others outside of our tight circle cease to care about our activities. Where our goal is usefulness, in contrast, we are far less likely to be seduced by our research. A good conceptual scheme: a) is interesting, meaning it conveys something novel to the client, b) is simple enough to be communicated effectively, and c) recognizes its own limitations. This essay therefore proposes that we widen our research objectives to include the development of conceptual schemes and establish standards for the rigorous evaluation of such schemes. Evaluation is critical because conceptual schemes are actually intended for use; a bad one can do far more damage than bad theory that lies fallow and unapplied.

Keywords: theory, conceptual scheme, complexity, informing science, research methods.

Introduction

Material published as part of this publication, either on-line or in print, is copyrighted by the Informing Science Institute. Permission to make digital or paper copy of part or all of these works for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage AND that copies 1) bear this notice in full and 2) give the full citation on the first page. It is permissible to abstract these works so long as credit is given. To copy in all other cases or to republish or to post on a server or to redistribute to lists requires specific permission and payment of a fee. Contact HPublisher@InformingScience.org to request redistribution permission.

Chip and Dan Heath (2010) recently published a book called *Switch*, dealing with the challenges of making personal and organizational changes. Throughout the book, they referred to the process of a rider guiding an elephant—an image they, in turn, credited to Haidt’s (2006) *The Happiness Hypothesis*.

The authors likened the elephant-rider image to the internal processes through which rationality (the rider) and emo-

Conceptual Schemes

tional drives (the elephant) interact within the human mind. They use the model—an example of a *conceptual scheme* in the parlance of this essay—to make a number of points, including:

1. While the rider may perceive himself to be in control, the elephant provides the force through which goals are accomplished.
2. When the rider fights the elephant, he can maintain control over the short run but he tires more quickly than the elephant, meaning the elephant will always win in a long term struggle.
3. The best way for the rider and elephant to work together is for the rider to guide the elephant down a path that meets the needs of both.

What I found particularly intriguing about this image were two key points. The first was how much I could recall after putting the book down, largely as a consequence of the authors' repeated use of the image. Indeed, I listed the previous three points without even bothering to look back at the book, now sitting on my shelf. Naturally, the authors—who are top researchers in their own right—reinforced the lessons with specific research results; Haidt's (2006) book, which I immediately ordered based upon the authors' recommendation, provided many additional research findings grounded in solid science. Nevertheless, I must admit to finding the elephant-rider lessons coming to mind far more readily than the other examples.

The second point I found intriguing was somewhat more troubling. I tried to imagine what might happen if I were to submit a paper based upon elephants and riders or—more to the point—my own favorite analogies to a top research journal in my chosen field of MIS. I fear that the outcome would not be pretty. First, I am reasonably confident that I would be informed that what I had written was not “research”. I base this belief upon the vitriolic comments that popular business trade books typically garner within the academic community. Second, I suspect that I would be informed that the goal of the journal is to create and validate knowledge, not to publish works whose principal goal is to explore better ways to communicate that knowledge effectively. Finally, I am convinced that the relaxed narrative tone that I would employ in conveying such a message would provoke rebuke. I would be reminded that no serious researcher would risk trivializing his or her findings by presenting them so informally.

If the reader finds these predicted responses credible overall—and I gleefully admit that I have been surprised, from time-to-time, by unexpectedly positive reactions to some of my more adventurous publication efforts—it should be clear that such attitudes pose a particularly serious problem to the informing science transdiscipline. Given the huge range of contributing disciplines represented in informing science—MIS, education, computer science, communications, library science, information science, psychology, cognitive informatics, philosophy, and so forth—it is likely that if each of us sticks to the narrow strictures for conducting and presenting research that have evolved in each of our respective areas, none of us will understand what anyone else is talking about. In fact, I propose that the greatest challenge informing science faces is how we can best communicate with each other. For this reason, research whose principal goal is to find ways to communicate existing knowledge effectively should be among our top priorities. It should not be viewed as a “nice to have” add-on to our knowledge creation and validation activities.

In this essay, I contrast the development of theory—the search for what is true—with the development of conceptual schemes—the search for what is useful. I begin by examining an unanswerable question: What is theory? I then consider the nature of conceptual schemes, identifying the ways in which they are similar and different from theory. I then explore how the complexity of the phenomena being studied can impact the relative merits of theories versus conceptual schemes. I conclude by encouraging the development of more robust conceptual schemes within informing science and other disciplines that deal with complex phenomena.

What is Theory?

The question of what constitutes theory could easily fill an entire paper—one that would lead to a predictable conclusion: “It depends upon whom you ask.” The implication of this state of definitional disarray is that anyone using the term must specify how he or she intends to interpret it. Before meeting this obligation, I will review some extremes of the considerable range of definitions that have been proposed. After that, I will provide some examples of theory—as I choose to employ the term—and consider how such theory can be established and defended.

Definitions of Theory

For our purposes, I use the term *model* to describe any system of observations, assumptions, relationships, and principles held by an individual or collection of individuals. What constitutes *theory*, in turn, is some subset of all possible models.

Under the most stringent definitions, a scientific theory is a set of relationships strongly supported by observations and well established principles. As an example, in its defense of evolutionary theory, the *American Association for the Advancement of Science* (AAAS, 2010) asserts the following:

A scientific theory is a well-substantiated explanation of some aspect of the natural world, based on a body of facts that have been repeatedly confirmed through observation and experiment. Such fact-supported theories are not “guesses” but reliable accounts of the real world.

According to this definition, relatively little research in the social sciences would pass for scientific theory.

Table 1: A taxonomy of theory types in IS research (Gregor, 2006, p. 620)

Theory type	Distinguishing Attributes
I. Analysis	<i>Says “what is”.</i> The theory does not extend beyond analysis and description. No causal relationships among phenomena are specified and no predictions are made.
II. Explanation	<i>Says “what is”, “how”, “why”, “when”, “where”.</i> The theory provides explanations but does not aim to predict with any precision. There are no testable propositions.
III. Prediction	<i>Says “what is” and “what will be”.</i> The theory provides predictions and has testable propositions but does not have well-developed justificatory causal explanations.
IV. Explanation and prediction (EP)	<i>Says “what is”, “how”, “why”, “when”, “where” and “what will be”.</i> Provides predictions and has both testable propositions and causal explanations.
V. Design and action	<i>Says “how to do something”.</i> The theory gives explicit prescriptions (e.g., methods, techniques, principles of form and function) for constructing an artifact.

Nearing the opposite extreme, five types of theory have been identified in the MIS field (Gregor, 2006). Under this heavily cited taxonomy, presented in Table 1, analysis theory (Type I) requires neither that causal relationships be specified nor predictions be made; furthermore, such theory is specifically presumed to be non-testable.

The Gregor (2006) taxonomy classifies theory based upon the types of questions that the research is designed to answer. Another classification scheme is based upon the constituent elements. For example, Roethlisberger (1977) asserts that scientific theory mainly consists of general propositions, empirical propositions, and elementary concepts that, in turn, are used to build deductive systems and statements of relationships.

A third approach to defining theory—and the one that I will use—involves the researcher’s ultimate motivation or goal. McKelvey (2002, p. 1) states this succinctly as “Good theory is truthful explanation”—implying that the goal of theory is, or should be, uncovering “truth”. This sentiment is echoed by Hevner, March, Park, and Ram (2004), who differentiate traditional MIS behavioral research from design science by characterizing the goal of the former as being “truth”, while the latter strives for “effectiveness”.

Creating and Testing Theory

One advantage of differentiating theory from other types of models by virtue of its intended relationship to truth is that doing so establishes a clear boundary for what is, and what is not, to be considered theory. Under the “truth definition”, as long as one believes the components of the model being used are individually true—meaning that they accurately reflect the phenomena being modeled at an appropriate level of abstraction or approximation—then the collected components constitute a theory. Where the relationship to truth does not exist or fails, the components remain a model but cease to be a theory.

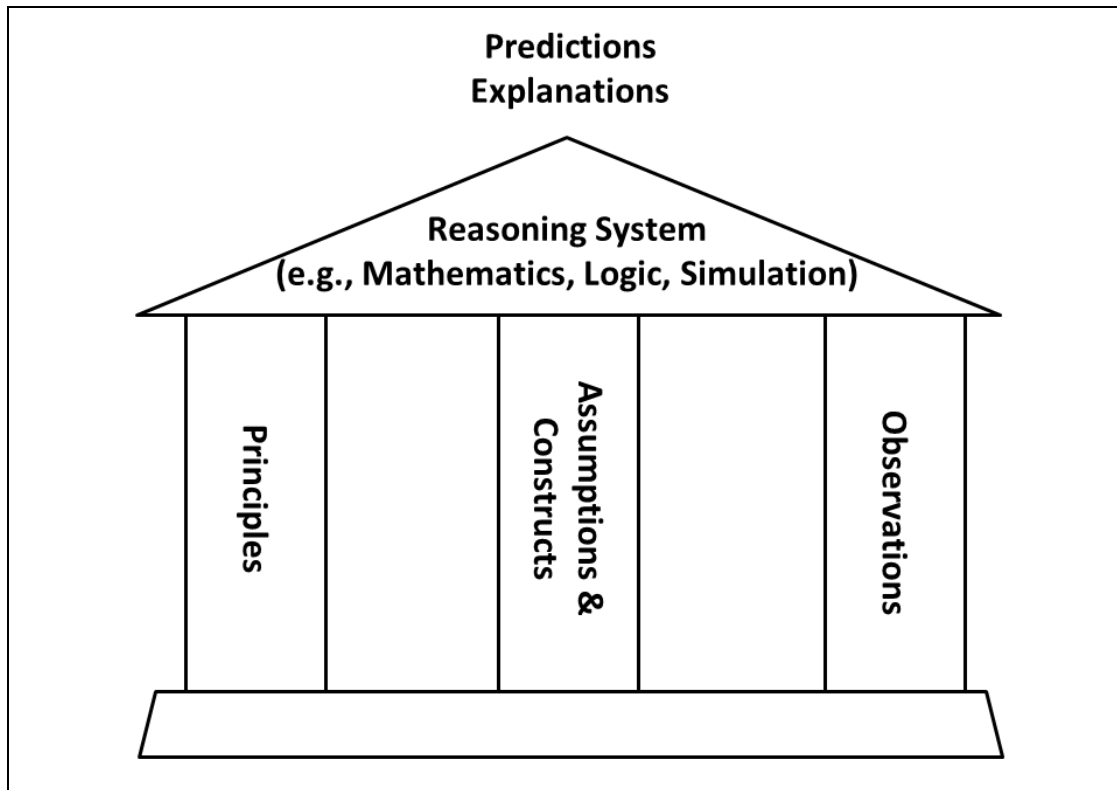


Figure 1: Construction and components of a theory

With apologies to the many philosophers who have written treatises on the subject, Figure 1 presents a simplified view of the components of a theory. The inputs of a theory include principles and laws, assumptions and constructs, and observations. Principles and laws consist of relationships or assertions of fact that are accepted without question. Assumptions and constructs are key

elements needed to construct the theory that are potentially subject to test. Observations are what we use to link the theory to “truth”. Using these inputs as our raw materials, we reach conclusions by applying a reasoning system—such as logic, mathematics or simulation—that may either strive to explain existing observations or predict future behavior/observations.

As previously stated, as long as a model continues to provide satisfactory explanations and/or predictions *and* its underlying assumptions remain unchallenged, we may continue to call it a theory. Much research is devoted to testing theory. Under the definition that I am using here this process is tantamount to asking the question: Does the model and its component parts have the relationship to truth necessary to be called a theory? To illustrate the process, let us quickly look at two well-known examples: the existence of ether and evolution.

The Existence of Ether

Almost everyone who has taken a college-level physics course has been lectured on the origins of Einstein’s special theory of relativity. It started with a vigorous debate in the 17th and 18th centuries over the nature of light. The corpuscular proponents argued that light was best viewed as a stream of particles. The wave theory proponents, on the other hand, argued that light was better thought of as a wave, like sound. The wave advocates could point to a variety of characteristics of light—such as its constant speed and the way it behaved when encountering objects such as a prism or an aperture—that were consistent with wave behaviors. The corpuscular proponents, on the other hand, argued that waves need a medium. If outer space was truly empty, how could light then travel through it?

By the beginning of the 19th century, the wave paradigm had come to dominate science. To address the problem of the lack of a medium, they proposed that a massless substance known as ether (also spelled aether) served to carry light waves. Based upon this assumption, it was predicted that ether drift should be observed (Kuhn, 1970, p. 73). The problem was that no experiment had been able to detect such drift. The most famous experiment along these lines was the Michelson-Morley experiment, which split a beam of light into two paths, reflected them, and measured the interference between them. Given that the two paths were perpendicular, the velocity of light on one path was expected to be different from that of the other path (since the rate of travel through the ether would necessarily be different in the two directions, assuming that ether was fixed in space). Unfortunately for wave theory, the velocities appeared to be identical.

To address this discrepancy, it was proposed that matter “drags” ether with it—the way a thin layer of water adheres to an object travelling through it. Various estimates, such as the Fresnel drag coefficient, provided mathematical estimates of this drag consistent with the observed behavior. The Lorentz transformations provided another approach, arguing that lengths were compressed and time slowed as objects traveled through the ether.

In the early 20th century, Albert Einstein proposed another explanation, special relativity. In the relativistic model, the invariance of the speed of light from all points of observation was taken as a starting point. Einstein was able to derive the Lorentz transformations using this approach, which did not rely on the existence of ether. Subsequent experiments universally confirmed Einstein’s predictions, causing the ether paradigm to be discarded. In this example, therefore, we see two stages of theory evolution:

1. The assumptions behind the ether theory are revised as a consequence of observations (e.g., the incorporation of ether drag or Lorentz transformation).
2. The ether theory is abandoned when it becomes evident that a simpler model offers comparable or improved predictions without the need for assuming the existence of a substance whose presence cannot be detected.

Evolution

Darwin's Theory of Evolution provides an interesting counterpoint to the previous ether example. Whereas physics offered considerable opportunity for prediction and experimentation, evolutionary theory was focused on explanation. Even today, scientists rarely attempt to "predict" what the next great species will look like.

At the time he began his research, the prevailing Western view was that plant and animal species had been placed on the planet by God and were immutable. Through meticulous observation of many species around the globe, Darwin continually found evidence of traits that were uniquely adapted to local conditions. Ultimately, that led to his theory of survival of the fittest—whereby beneficial traits in plant and animal species tend to be propagated to future generations at higher rates than less beneficial traits, leading to a gradual change in the species.

What was particularly interesting about his theory was the critical—and, at the time, heroic—assumptions he made in order to give the theory substance. Among these (Hayden, 2009):

1. That a mechanism existed whereby new traits could spontaneously emerge and be passed on to offspring. During Darwin's time, the study of genetics was relatively unknown. While millennia of animal husbandry had demonstrated that traits could, to some extent, be passed to offspring, this observation was viewed as a special case demanding human intervention (an example of "intelligent design" in today's parlance).
2. That intermediate examples of species representing the transition from one set of traits to another could be found. While fossils and fossil hunting had become popular in Darwin's time, the assumption was that these represented animals that had become extinct (perhaps during the great flood mentioned in the Bible) and had no connection to today's species.
3. That an explanation—other than the intervention of a higher power—could be found for how very similar species (such as the ostrich and the rhea) could appear in entirely different locations separated by impenetrable ocean barriers.

The need for all of these assumptions represented a potential threat to the theory. The robustness of evolution stems from the fact that justifications for each assumption have been found. For example, the study of genetics offers mechanisms for the emergence and transmission of traits, a vastly expanded fossil record confirms the existence of intermediate states in the transition from one life form to another, and the discovery of plate tectonics explains how species could have migrated from one location to another across land bridges that no longer exist (Hayden, 2009). Thus, by confirming the assumptions of the theory, the linkage between the theory and truth has grown ever stronger. To quote a biologist at the University of California (Hayden, 2009, p. 48):

Darwin didn't know 99 percent of what we know... but the 1 percent he did know was the most important part.

Conceptual Schemes

A direct correspondence to truth is not the only goal we might have in mind for a model. In the introduction to this essay, for example, I presented the elephant/rider model employed by Heath & Heath (2010) and Haidt (2006). That model, I proposed, could be useful although any direct connection to truth would be speculative at best. I therefore proposed a different term, *conceptual scheme*, be used as an alternative to the term "theory" for such models.

The specific term conceptual scheme was suggested by Fritz Roethlisberger's book *The Elusive Phenomena* (1977), which explored the evolution of his work with Mayo and others in the field of organizational behavior. He proposed the model of knowledge evolution presented in Figure 2.

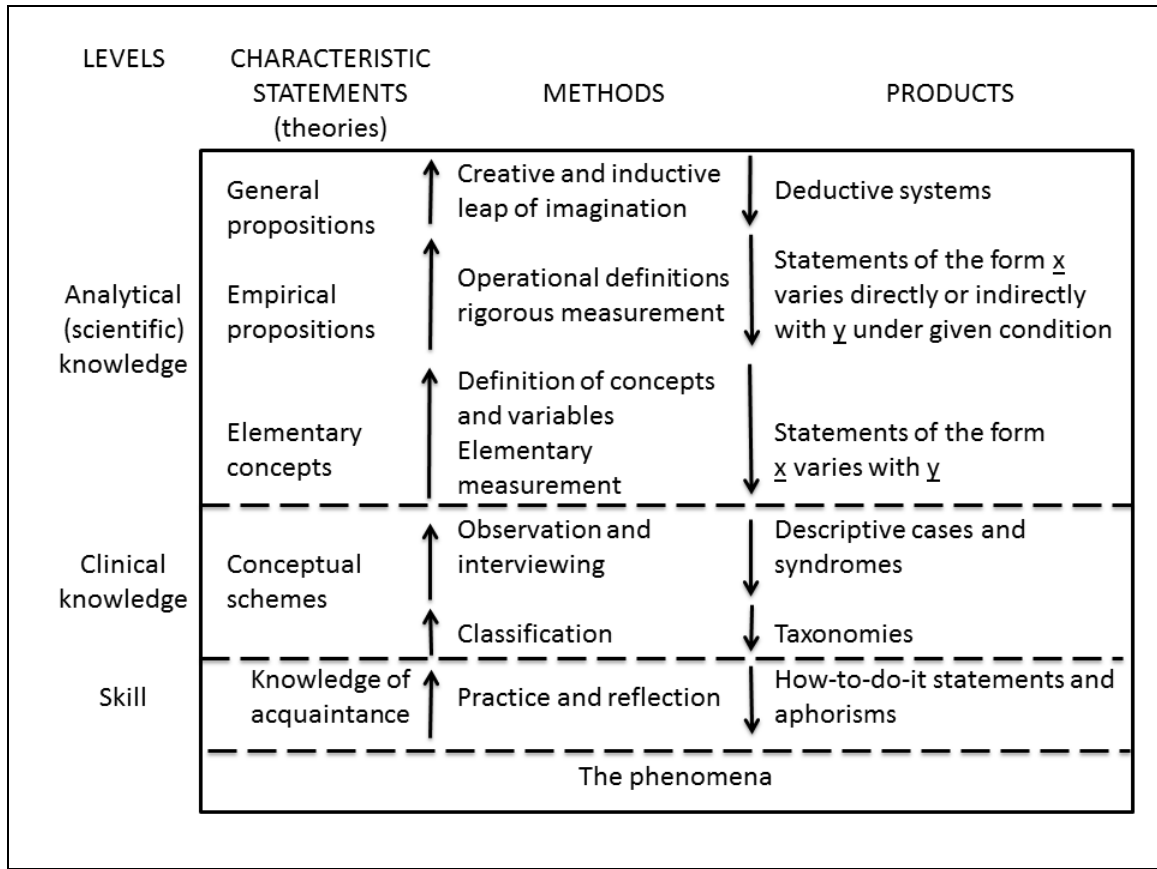


Figure 2: Evolution of knowledge (Roethlisberger, 1977, p. 393)

Within this model, the term “conceptual schemes” is used to describe clinical knowledge—the type of knowledge that is applied in practice. Davidson (1973-1974) describes the conceptual scheme as follows:

Conceptual schemes, we are told, are ways of organizing experience; they are systems of categories that give form to the data of sensation; they are points of view from which individuals, cultures, or periods survey the passing scene. There may be no translating from one scheme to another, in which case the beliefs, desires, hopes and bits of knowledge that characterize one person have no true counterparts for the subscriber to another scheme. Reality itself is relative to a scheme: what counts as real in one system may not in another. (p. 5)

Here, once again, we find the notion that such schemes may be relevant in a particular context, but not necessarily true in the objective sense. Davidson argues that conceptual schemes are necessarily translatable, and that a single underlying truth may exist. Others argue for conceptual relativism (e.g., Aune, 1987), suggesting that different schemes may reflect different truths that each may be valid in their own way.

For our purposes, I will focus on the clinical applicability of conceptual schemes emphasized in Figure 2. In other words, a conceptual scheme is a model that can be usefully applied. This definition suggests the relationship to theory presented in Figure 3.

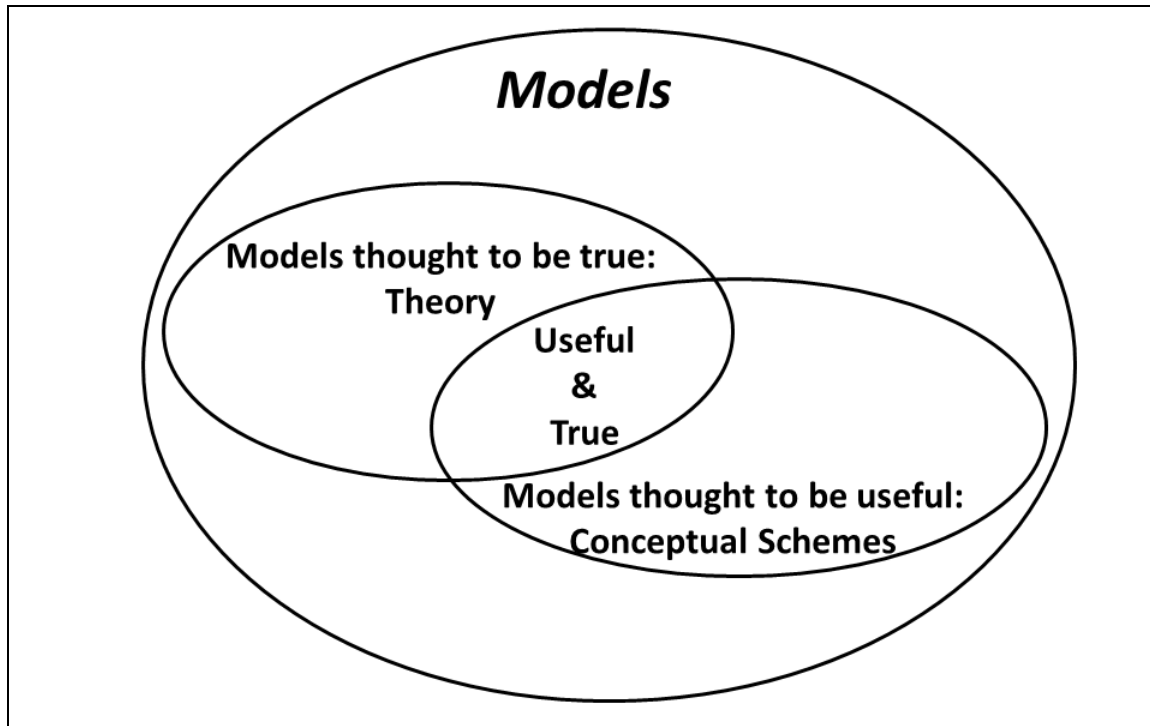


Figure 3: Models, theories and conceptual schemes

The most significant aspect of Figure 3 is that usefulness is not presumed to be a subset of truth. Why such a truth subsumes use relationship is dubious is illustrated by a well-worn gag of stand-up comics and situation comedies where an individual asks his or her spouse:

Does this outfit make me look fat?

Here, the useful response—a strong negative with no hesitation—may diverge considerably from the truthful one; indeed, in the context of comedy it *always* does.

One may reasonably argue that one-liners have little to do with theory. They do, however, have a great deal to do with communicating effectively. If truth acts as a barrier to communication, then in order to achieve usefulness an indirect path to the truth may be needed as part of the informing process. Such a path would, more or less as a matter of definition, require the development of a series of conceptual schemes specifically designed to resonate with the client.

Relating the truth-usefulness divergence to a more scientific context, we often find that models whose existence begin as attempts to build theory continue on as conceptual schemes once their underlying truth has been refuted. A good example of this is Maslow's (1943) Hierarchy of Needs. Still widely taught in management and marketing, this model proposes that human needs form a hierarchy, ranging from hygiene needs (e.g., food, clothing, shelter, sex) at the bottom to self-actualization at the top. Individuals, it is proposed, are not motivated to satisfy needs at higher levels until the needs at lower levels have been satisfied. Since the time that the model was proposed as theory, however, many criticisms have been raised with respect to its validity. Among these are (a) findings that the proposed hierarchy is biased towards the motivational structure of certain cultures and (b) many examples where individuals bypass or ignore some lower levels in the process of striving for higher-level goals to which they are committed. Nevertheless, the model is probably right more often than it is wrong and has the added benefit of being easy to communicate. Thus, it continues to be taught and applied.

Usefulness and Informing

Before turning to some examples of conceptual schemes, we need to think about what being “useful” actually means. From an informing system standpoint, the first issue we need to address is the identity of the client. The usefulness of a particular model, for example, is likely to be very different for a student client versus an experienced practitioner. To be consistent with Figure 2, I will emphasize practitioner clients in this essay, since they would most commonly be in a position to apply clinical knowledge and the findings of advanced research.

Whatever the nature of the client, a particular model will be useful only to the extent that it is:

1. *Relevant*: It addresses a problem or decision area faced by the client. I will refer to this quality as *relevance*, and will return to it when specifically considering the question of what we should be researching.
2. *Acceptable*: It can be communicated effectively to the client, impacting the client’s own internal models without being rejected or seriously distorted, and
3. *Potentially Actionable*: It subsequently impacts the client’s activities or has the potential to impact client activities should certain circumstances be encountered.

Stated as compactly as possible, as I define useful it exists at the intersection of the three conditions, i.e.,

$$\text{Useful} \equiv \text{Relevant} \cap \text{Acceptable} \cap \text{Potentially Actionable}$$

Including “potential” in the last condition is important since there are some types of knowledge that prepare us for future circumstances that may never occur. For example, while I have never performed the Heimlich maneuver to dislodge an obstruction in the windpipe of a choking victim, I have been taught to apply it and I judge it to be useful knowledge because I *might* need to apply it at some point in the future. The potential value is increased because, should that need occur, the knowledge must already be in place since there will be no time to acquire it at that point. To frame it in decision science terms, the knowledge has an expected value of usefulness. That expected value of usefulness might be computed by taking the usefulness of actually applying the knowledge and discounting it by the likelihood that the knowledge will be applied.

Chip and Dan Heath (2007), mentioned in the introduction, wrote an earlier book, *Made to Stick*, that does a very nice job summarizing the key elements of a communication that make it likely to resonate with a client—what they refer to as *stickiness*. The stickiest messages, they argue, have the six properties: Simple, Unexpected, Concrete, Credible, Emotional, and Stories (the SUCCESS framework). In order to meet the second condition for model usefulness, some level of stickiness is likely to be required. Otherwise, the message is likely to be ignored or forgotten.

Obviously, whether or not a theory can be made simple will depend on the phenomenon being studied; I will discuss that topic in much greater detail later in the essay. With respect to the remaining criteria, the “unexpected” and “credible” characteristics of stickiness often exist in an uneasy relationship. Imagine that we are attempting to communicate a simple proposition to a client. That client will, quite naturally, have a prior belief with respect to the proposition. He or she may believe the proposition to be false, to be true or—the intermediate case—have no opinion regarding it. As illustrated by Figure 4, the expected value of the informing process will vary based upon the client’s initial perception. Where the client already believes the proposition to be true, there will be no change in the client’s model; the message is not useful because it is not unexpected. At the other extreme, where the client strongly believes the proposition to be false, a simple message will have a very low likelihood of changing the belief; here we encounter a failure of credibility. It is in the region of no opinion (uncertainty) where the greatest likelihood of valuable informing will take place. This notion of an intermediate (inverted U) peak in informing

effectiveness is well supported by the literature. It is implied in Davis' (1971) widely cited article "That's Interesting". The importance of prior beliefs in achieving effective informing is also well established in the diffusion (e.g., Rogers, 2003), management (e.g., Dane, 2010) and education (e.g., Bain, 2004) literatures.

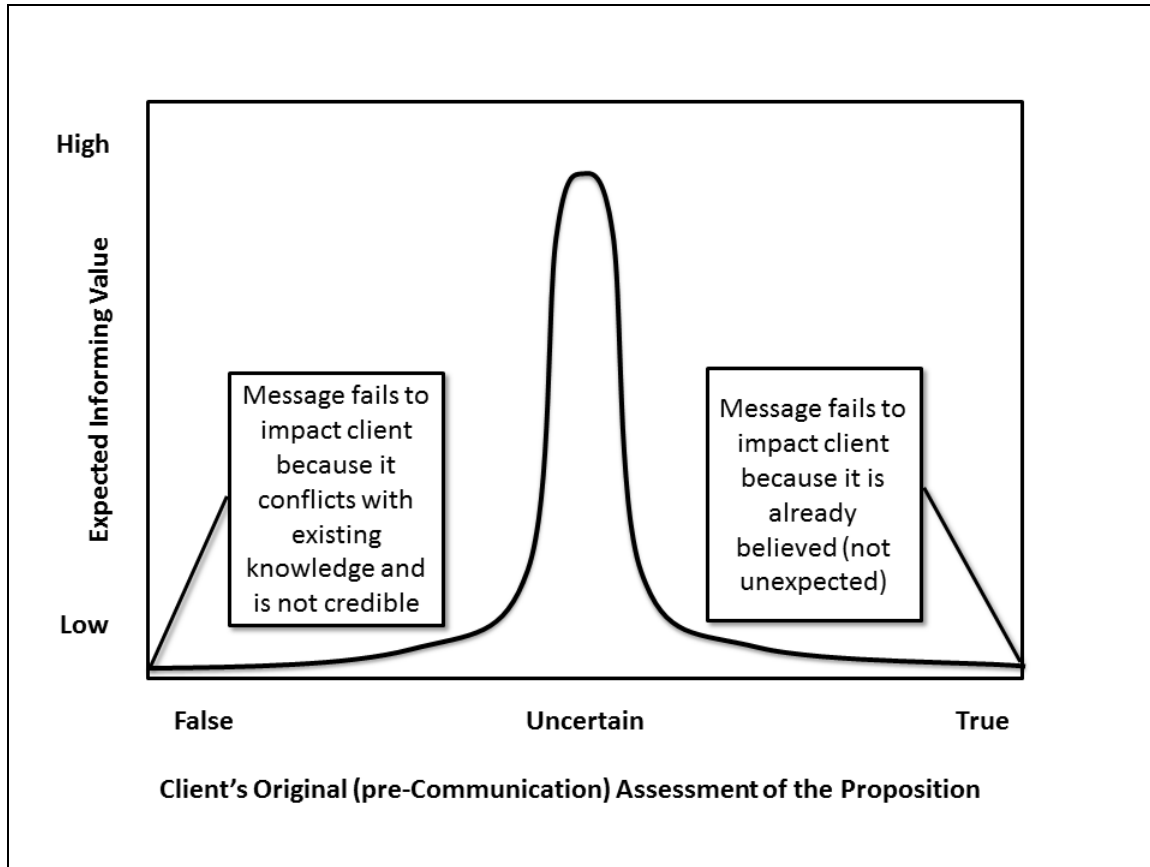


Figure 4: Expected message value as a function of prior beliefs

Conceptual schemes have a considerable edge over theory with respect to the remaining three SUCCESs factors. In order to enhance its credibility, theory—particularly in the social sciences—tends to be:

- Presented in abstract terms rather than describing a specific situation (i.e., it is not concrete)
- Justified as objectively as possible (i.e., emotion is frowned upon)
- Framed in terms of principles, constructs, and relationships rather than as a narrative (i.e., it avoids storytelling).

There is nothing, however, that prevents a conceptual scheme from being concrete, from incorporating emotions, or from being presented in the form of a story. Indeed, if you look at religious beliefs—conceptual schemes that are applied by billions of people to guide their daily lives—we see principles almost universally conveyed in the form of concrete, emotion-laden stories. This application of a particular religion's tenets takes place irrespective of what non-believers happen to think of their relationship to truth. Indeed, faith is often taken to be a matter of accepting that which is not proven to be true, even in the face of conflicting evidence.

Conceptual Schemes in Economics

Nowhere is the distinction that I make between conceptual schemes and theory more apparent than in the fields of economics and its close business cousin, finance. In these fields, the models used are nearly universally described as theory. But do they meet the mapping-to-truth test I propose in this paper?

There is ample evidence that economists routinely ignore inconvenient observations that would make their models impossibly intractable. Discussed in greater detail elsewhere (e.g., Gill, 2010), numerous examples of this practice can be found. Among these:

- Axioms of consistent preference, both with respect to risk and with respect to time, are virtually never confirmed by experiments that use human subjects (e.g., Ariely, 2009; Kahnemann, Slovic, & Tversky, 1982; Thaler, 2000). Because these axioms serve as the foundation of most economic and financial mathematical models, their failure to be confirmed by observation calls into question the truth of the entire enterprise.
- Assumptions regarding the normality of risk distributions are central to many financial models, such as portfolio theory. A series of recent market failures (e.g., credit markets, real estate markets) have demonstrated that we are prone to misunderstanding the nature and magnitude of portfolio risks, vastly overestimating our ability to predict them (Taleb, 2007).
- Economic models appear to have very little predictive power. In the U.S., for example, attempts to predict unemployment changes in light of fiscal and monetary interventions have been almost laughably inadequate. The problem of prediction is well established in economics (Tetlock, 1999) and other fields where complexity is high (Shanteau, 1992).
- Multiple competing models exist, many of which appear to be more heavily influenced by ideology than by truth. For example, in attempting to explain the events of the U.S. Great Depression of the 1930s, a right-leaning economist will (predictably) assert that the recovery was impeded by the fiscal interventions of the Roosevelt Administration while a liberal-leaning economist will (equally predictably) assert that recovery was impeded by the reluctance of the U.S. Congress and, to some extent, the Roosevelt Administration itself to go far enough in its policy of fiscal intervention.

These examples—and many more could be presented—should cause any reasonable scholar to dismiss the objective truth of economic “theory”. But just because a model maps poorly to the truth does not make it a conceptual scheme as I define it. It must also be useful.

With respect to usefulness, I would propose that the situation with respect to economics is radically different. In both government and finance, decisions with a large potential to impact investments and the economy *must* be made. Postponing such decisions until the underlying truth of a model has been substantiated is not an option; *the decision not to act is, in itself, a decision to follow a particular course of action*. Thus, we need models to help us make decisions. In using them, *we necessarily accept the fact that these models are not completely true* (i.e., they do not meet the definition of theory that I am using). Nevertheless, they are useful.

What is particularly interesting about the economics and finance fields is the degree to which the models proposed by academics—as well as the academics themselves—have become integrated into practice. Keynesian economics has academic roots, as do most economic models. Administrations not only adopt these models (often on ideological grounds), they also hire prominent academics to help interpret and implement recommendations based on these models. The same is true in finance—the one area of business academia where academic institutions frequently find themselves competing with the private sector for doctoral students and faculty. Economists and

financial specialists have also established a large infrastructure for acquiring the data needed for their models, such as the U.S. Bureau of Labor Statistics and its U.N. counterparts. Academic expertise is often the prerequisite for making a sensible interpretation of the data that has been gathered. In other words, the area of business-related research that has diffused to practice more effectively than any other is the one that has focused its efforts on the creation of conceptual schemes (even if researchers choose to label them as “theories”).

The Rugged Landscape Conceptual Scheme

My current favorite conceptual scheme is that of the rugged landscape, originally proposed by evolutionary biologist Stuart Kauffman (1993). The model is a way of thinking about how an entity’s characteristics determine its fitness. I find it particularly useful because it highlights the limitations of theory as complexity grows, and I will therefore apply it extensively throughout the remainder of this essay. There are three different forms of the scheme that I use to visualize rugged landscape relationships: one as proposed by Kauffman, one viewing the relationship in terms of an actual landscape, and one involving a cookbook. I now summarize each of these briefly.

The NK landscape

The NK fitness landscape model, proposed by Kauffman (1993), provides a tool for visualizing how interactions between variables can impact a dependent variable, referred to as fitness. Imagine, for example, that we have 20 binary (having values of 0 or 1) attributes, x_1 through x_{20} , that impact the survivability of an organism. On one extreme—the ordered landscape—each of the 20 impacts would be independent of each other. That would lead to an expression of the form:

$$\text{Fitness} = c_0 + c_1x_1 + c_2x_2 + \dots + c_{20}x_{20}$$

This is essentially the familiar regression equation that is routinely employed—in its various mutations, such as structural equation modeling—in much of the research conducted in the MIS field and throughout the social sciences. If this equation truly reflects the environment being modeled, then we have a very compact theory-of-fitness: with only 21 values (c_0 through c_{20}) we can capture the entire relationship existing between the attributes and fitness. In Kauffman’s model, where N is the number of attributes, and K is the number of interactions between attributes, this relationship would be described as an $20,0$ landscape. Such a landscape will necessarily have one, and only one, peak value—that combination for which the most advantageous value of x_i is chosen depending on whether the coefficient c_i is positive or negative. As illustrated by the left side of Figure 5, fitness is determined by a stack of individual influences.

At the other extreme, illustrated by the right hand side of Figure 5, we have Kauffman’s chaotic landscape. In this environment, we cannot determine the contribution to fitness of any attribute x_i without knowing all 19 other attribute values. In other words, each variable interacts with 19 other variables, leading to a $20,19$ landscape. Kauffman simulates this situation by assigning a random value to each of the 2^{20} possible attribute combinations. A complete theory-of-fitness for this landscape therefore require 2^{20} (about a million) values. Quite obviously, this is not a compact theory; in fact, it would be nearly impossible to communicate. Even the number of local fitness peaks—estimated to be $2^N/(N+1)$ —is very large (about 50,000, in this example).

What Kauffman argues is that living systems tend to evolve towards intermediate K values. Ordered systems with a single peak tend to be efficient but brittle in the face of environmental change, since all entities tend to migrate towards the single peak over time. Chaotic systems, on the other hand, tend to trap many organisms on sub-optimal local peaks but, in the face of major environmental change, are more robust as a consequence of the diversity of peaks where entities reside. This balance is described as the *complexity* that exists on the boundary of order and chaos.

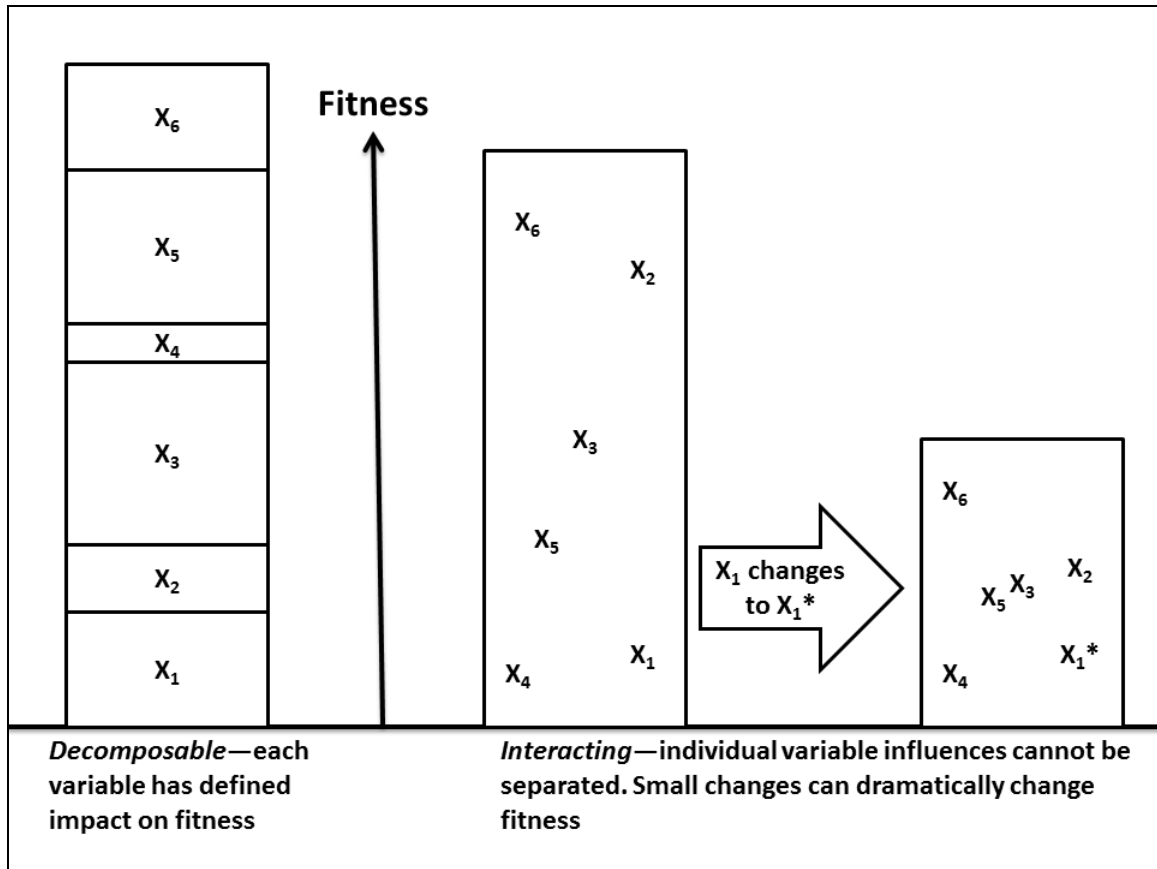


Figure 5: Decomposable ($K=0$) vs. fully interacting ($K=N-1$) relationships between attributes and fitness

The key ideas that the NK fitness landscape model helps cement in my mind are the following:

1. As interactions between variables grow, there is an explosion in the size of the theory that will be required to accurately and completely describe a landscape.
2. Entities will migrate to peaks over that time and the local behavior around those peaks can be very different.
3. As environments become more dynamic, the influence of interaction effects on fitness should grow, since diversity is the key to survival under turbulent conditions.

The rugged terrain

Another way of viewing the rugged landscape model—more akin to the elephant-rider analogy—is to think in terms of actual terrains. Imagine that a certain number of local peaks exist. The presence and number of such peaks does not tell the whole story. Some landscapes may be like sand dunes. On such landscapes, changes in fitness that result from changing values are gradual. At any arbitrary position in the landscape, the best estimate for adjacent positions is the same as the position you are on. Such a landscape is *smooth*.

At the other extreme, landscapes such as Kauffman's chaotic ($N,N-1$) landscape are very different in character. For these landscapes, knowing the value of one position tells you absolutely nothing about the values of adjacent positions. No matter where you are on the landscape, your best estimate for adjacent values is the average for the landscape as a whole. These *rugged* landscapes

suggest a visual image more like that of Utah’s Bryce Canyon, where hoodoos suddenly jut out of the ground giving the terrain an almost digital flavor, as illustrated in Figure 6.

The significance of the smooth-rugged continuum becomes apparent when you imagine searching for peaks blindfolded or in a heavy fog. The smooth terrain responds reasonably well to incremental search. Take a step. If the step moves you up, take another step. If it does not, step back to your original position and step in another direction. Once you’ve tried stepping in all possible directions and always returned to your original position, you are on a peak.




Figure 6: Smooth landscape (left) and rugged landscape (right).

Dunes image from Microsoft Office ClipArt. Bryce Canyon photo (right) taken by the author.

On the rugged terrain, however, such a search procedure is very risky since any step could take you off a cliff—even small changes in state can produce huge changes in fitness. In such landscapes, the least risky strategy may be to stay put if your fitness is above average. Even better, you can try to identify someone who is very much like you (since even small differences can exert a huge impact on fitness): someone who also appears to be at a higher level. Then imitate that person. This phenomenon of accepting influence most readily from individuals very much like you is referred to as *homophily*. It is widely observed in informing systems (Rogers, 2003) and makes complete sense on rugged landscapes where the fitness of particular states is not readily apparent.

The rugged terrain conceptual scheme offers some particularly useful insights with respect to searching for higher fitness:

1. As the terrain becomes more rugged, the challenges of searching for fitness become greater. Perhaps not coincidentally, the unofficial motto for Bryce Canyon is “helluva a place to lose a cow”, the words supposedly uttered by Ebenezer Bryce upon viewing the landscape for the first time. Because of the risks of low fitness being encountered during search, our motivation to satisfice can be expected to grow with ruggedness.
2. A rugged landscape also gives us the incentive to pay close attention to others like ourselves carefully, so as to determine if it makes sense to imitate them.

The cookbook

The final way in which I view the rugged landscape conceptual scheme is in terms of a cookbook. In many ways, a collection of recipes can be viewed as the end-product of an evolutionary fitness process. Each recipe, presumably, represents a local peak. After all, would an author publish a recipe knowing that an incremental change—such as adding more oregano—would improve its fitness? The range of recipes included illustrates how interactions between ingredients lead to diverse peaks. No one seeks to create an “optimal” dish, which is to say the single recipe (peak) that is the best choice for every possible occasion.

Three other important issues can be highlighted by thinking in cookbook terms:

1. The lack of decomposability between individual characteristics and fitness. While some ingredients—such as spices—may be individually distinguishable, asking questions such as “What percentage of a cake’s deliciousness is attributable to the eggs in the recipe?” is quickly seen to be nonsensical. For this reason, we do not see techniques such as multiple regression or structural equation modeling used to determine each ingredient’s impact on recipe fitness.
2. The potential for differing perceptions of fitness is readily apparent. As a consequence of my individual tastes, were I to rank each recipe in a particular cookbook according to my preferences, that ranking would likely differ substantially from yours.
3. The relationship of attributes to fitness may change over time. After three days of turkey leftovers, my own assessment of the fitness of turkey recipes is likely to be altered for several weeks.

The second issue—inability to agree on what constitutes fitness—can play an important role in conceptual schemes. In an intriguing op-ed piece for the *Wall Street Journal*, Jonathan Haidt (2010), of elephant-rider fame, proposed that part of the debate between economists of the left and the right has to do with alternative views regarding what constitutes “fairness”. On the one hand, he argues that the left is most likely to view fairness in terms of remedying inequality; e.g., it is unfair that one person is rich while another is not. The right, on the other hand, views it more from the perspective of karma (e.g., what goes around comes around); e.g., it is unfair when individuals who have worked and accumulated wealth are penalized for their industriousness through confiscatory taxes. Constructed through these two lenses, the shape of the fitness landscape could be radically different. It is hardly surprising that the two groups end up advocating radically different policies to achieve their own conception of fitness.

Ruggedness and truth vs. usefulness

My reason for presenting the rugged landscape conceptual scheme in such detail is that it carries important implications for the relationship between truth and usefulness. Ruggedness is, for the most part, a natural byproduct of complex adaptive systems—which is to say environments consisting of many diverse, interrelated elements that evolve over time. Using the various conceptual schemes I have just presented, the following conclusions are virtually inescapable:

1. Theory in rugged environments will need to be much larger than that required for decomposable environments. Recall the 21 vs. 2^{20} pieces of information required to describe 20,0 vs. 20,19 environments using the NK scale. **Implication:** *Truth becomes more difficult to communicate.*
2. As environments grow more rugged, we are increasingly motivated to imitate those similar to us (homophily). Recall how the search for peaks changes depending upon whether you are walking in the dunes vs. Bryce Canyon. **Implication:** *Truth becomes less important than finding a good entity to imitate.*
3. Rules applicable to specific peaks in rugged environments will vary considerably from peak-to-peak. Recall how the ingredients and instructions for different recipes in a cookbook vary. **Implication:** *The nature of truth becomes increasingly local.*
4. Environments that are rugged tend to emerge where criteria for fitness are subject to change. Recall how diversity tends to survive better in dynamic environments and how changing tastes are likely to influence our preferences for different recipes. **Implication:** *Truth becomes more fleeting.*

Collectively, these factors all suggest that as ruggedness (i.e., complexity) grows, the area of overlap between truth and usefulness—originally illustrated in Figure 3—is likely to become smaller. This is illustrated in Figure 7.

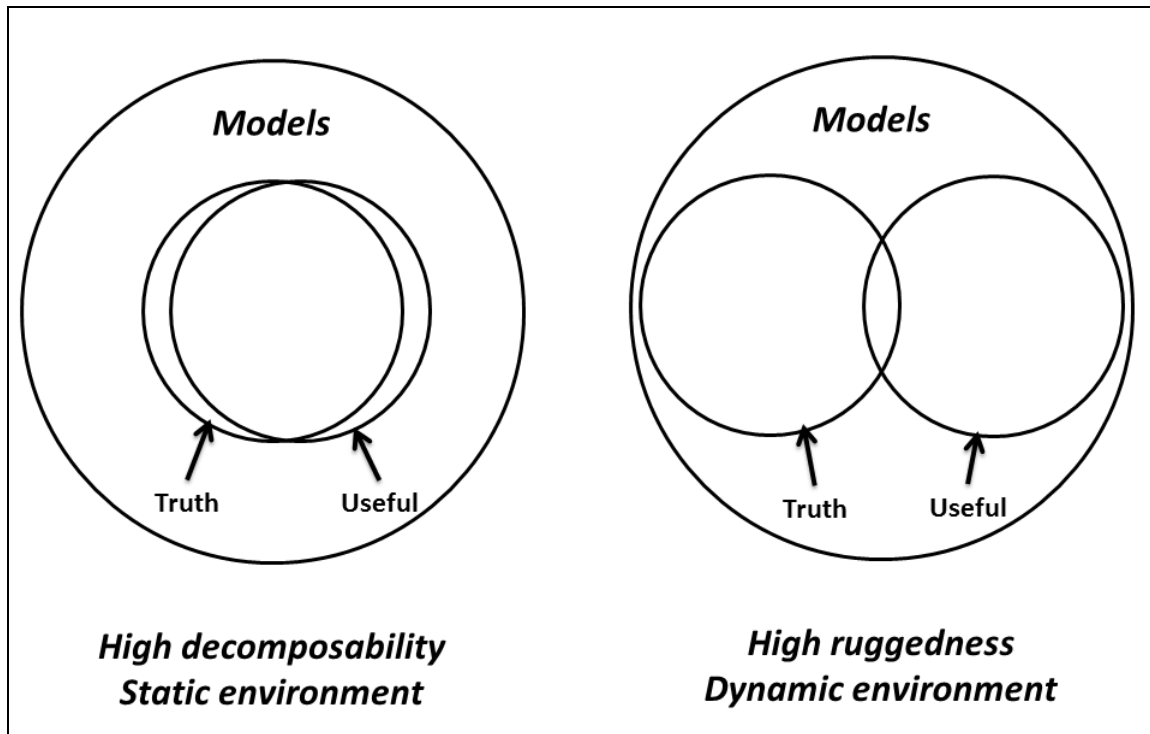


Figure 7: Change in overlap between truth (theory) and conceptual schemes (usefulness) as ruggedness increases

When I use the term useful in Figure 7, I specifically include the model’s potential to inform, which means that the client’s mental model needs to be changed as a consequence—part of the earlier “acceptable” condition. There are a number of barriers to informing that parallel the challenges that ruggedness presents to theory:

1. As theory grows larger and more involved, it becomes increasingly difficult to achieve resonance with a client. Simplicity, for example, is the first item in Heath & Heath’s (2007) model for “sticky” communications.
2. Complex knowledge tends to diffuse almost exclusively through face-to-face channels. Only a very small percentage of any population will accept an innovation based upon communications alone. Such changes are mainly the result of imitating self-similar peers (Rogers, 2003).
3. Generalizations will nearly always contradict some client experience on a landscape where the client is focused on a peak that is even slightly different from the peaks upon which the theory has been based. That contradiction will damage the credibility of the theory—another Heath & Heath (2007) stickiness factor.
4. Credibility of the theory will be further damaged where the time taken to develop it results in some or all of it becoming obsolete as a consequence of the continuing evolution of the system being modeled.

To repeat my earlier assertion, if a model is not constructed in a manner that allows it to be communicated to those clients who can put it into practice, it is not useful. This assertion seems self-evident no matter how true *and useful* a model *might* be if it were available to the client.

To summarize, I am proposing that in environments governed by static principles and simple interactions—such as the superposition principle that governs forces and waves in physics—it is very likely that a theory that articulates truth will be useful in the long run. In dynamic environments where interactions between elements are significant, any governing theory is likely to be so large that it will be nearly impossible to communicate and will take a long time to develop. The long development time makes it likely that the systems we are describing will have changed materially by the time we have fully developed our theory and rigorously tested it. Fields such as MIS, where the nature of the technology artifact often plays an important role in our theory, are particularly vulnerable to this problem. For example, well before we had developed solid theory about centralized mainframe systems with clients who knew little about computers we found ourselves confronted by a world of distributed microcomputers in the hands of far more knowledgeable client users.

The implication of Figure 7 is that our research goals may benefit from taking into account the inherent ruggedness of the domain we are studying. We now consider that issue more fully.

Theory vs. Conceptual Schemes in Research

In an ideal world, of course, we would demand that our research be aimed at uncovering the underlying truths of our world *and* that it should be useful as well. In the previous section, however, I argued that the intersection between truth and usefulness grows smaller and smaller as the landscapes we study become more rugged. Moreover, it has further been argued that informing landscapes and business landscapes are likely to be quite rugged (Gill, 2010), since they necessarily involve: a) establishing a fit between diverse senders and clients, a diversity accelerated by today's movement towards globalization, b) channels that are continually being altered by technology, c) content to be communicated that is rapidly evolving, and d) clients who are continually adapting to changing conditions. Thus, the question becomes one of priorities: do we worry about creating theory (truth) first, or should we focus on developing (useful) conceptual schemes?

Since the value of achieving both truth *and* usefulness seems fairly self-evident, in this section I will focus mainly on the drawbacks of becoming too single-minded in the quest for either.

The Drawbacks of Theory Research

I have already argued that as ruggedness increases, we can expect valid theory to become larger in size, more complicated in substance, and more transient in its validity. I would suggest, however, that this is not the only potential impact. In addition, I see a single-minded obsession with truth as being likely to lead to: 1) research that is specifically directed towards those hypotheses least likely to be useful, 2) being crowded out of the channels most effective for communications with practice, and 3) a level of self-absorption that is likely to act as a barrier to knowledge diffusion. I will examine these issues with a mix of the ruggedness conceptual scheme and examples from my own discipline.

Hypotheses unlikely to be useful

Earlier, I proposed that a model is useful a client only to the extent that it addresses a topic that relates to a client problem/decision (relevance), it is in a form the client can absorb (changes client mental models), and ultimately results in client action or potential action (actual or probabilistic effect). Unfortunately, the term “relevance” appears to have as many definitions as “theory” and as much variability between them. The definition that I am using is one that is particularly

focused on problems of interest. An MIS executive turned academic proposed, for example, that business research would be relevant if it addressed problems including (Weinberg, 2001, p. 99):

- Improved customer service
- Increased market share
- Reduced costs
- Market differentiation

These examples are consistent with my definition. What is also significant about this list is that all the items on it appear to be variables that could easily be described as fitness-related. In fact, some research I have been doing with a doctoral student suggests that the majority of empirical MIS research involves establishing or testing a mapping between a set of characteristics and one or more fitness-related variables. This is consistent with the earlier results of Orlikowski & Baroudi (1991) that found most (over 96%) MIS empirical research was conducted from a positivist perspective, assuming that fixed relationships between characteristics exist. For that reason, the rugged landscape conceptual scheme seems to be a good starting point.

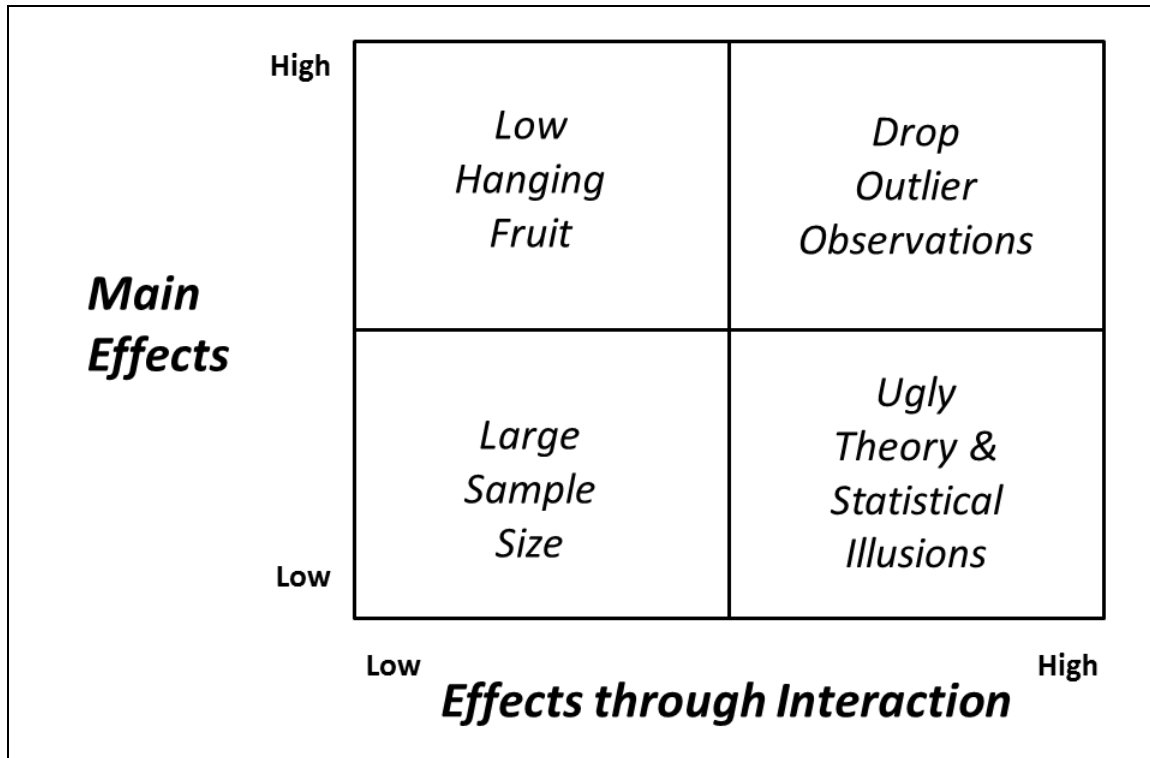


Figure 8: Types of variables in fitness landscape

Let us specifically look at the case of explanatory variables (i.e., statements of relationships in Figure 2) in a fitness landscape. As shown in Figure 8, each variable can exert effects on fitness in two ways: as a *main* (decomposable) effect or in *interaction* with other variables. Furthermore, we may also suppose that each of these effects may range from low (not particularly material to fitness) to high (important determinant of fitness). We also need to make two other assumptions:

1. That each entity on the landscape we are studying will try to migrate towards higher fitness, meaning that (eventually) each will reach a local peak.
2. That a researcher will, as a general rule, prefer results that support a positive relationship as opposed to a non-relationship.

Both of these assumptions are supported; either by “theory” (e.g., utility maximization) or by empirical research on publications (e.g., Hubbard & Armstrong, 1992).

Low Hanging Fruit. If we examine the four quadrants labeled in Figure 8, the upper left represents the case where a variable exerts a large influence on fitness that is independent of other variables (i.e., decomposable). This is the quadrant where the researcher first identifying the relationship can make an earth-shattering discovery—*provided that the relationship is not already familiar to practice*. That qualification, unfortunately, proves to be a great stumbling block to usefulness. We expect practitioners to be adapting continuously to improve their fitness. How likely is it that they will have missed such a straightforward and powerful relationship?

In my own field, MIS, the research literature is absolutely brimming with empirical tests of statements that—even if supported—are unlikely to add much to what practitioners already know. As an example, *in an article specifically focused on the relevance of MIS research*, Benbasat and Zmud (1999) identify the use of the theory of planned behavior (TPB) as *an exemplar of theory that was relevant and had been successfully applied in MIS research*. They described the theory as follows:

TPB posits that an individual’s use of a specific IT will be influenced by: (1) attitudes formed by the person’s beliefs about the expectations of outcomes associated with the IT use, (2) subjective norms associated with IT use, and, (3) perceived behavioral control formed by the person’s beliefs about the extent to which the person is capable of actually using the IT. (p. 10)

Attempting to translate this into slightly less convoluted terms, I see these factors as being roughly equivalent to: (1) Do I think the technology will do the job well? (2) Will other people think more (or less) of me for using the technology? (3) Will I be able to figure out how to make the technology work? The question I would then pose is this: How hard would it be to find a practitioner who does not think these factors matter—independent of their knowledge of MIS research? If the answer is *nearly impossible*, as I suspect it would be, then establishing the truth of these relationships through extensive empirical research is unlikely to produce significant changes to client mental models. Such research therefore fails to meet the criteria for usefulness that I have proposed. On the other hand, I agree fully with the authors that the topic is relevant, since the problem of figuring out what makes users accept or reject a particular technology is frequently of concern in organizations. The weakness in the research strategy we have pursued is that picking the low hanging fruit is unlikely to offer much value to individuals who are already experts. It does, however, greatly increase the likelihood of non-null findings.

From a research standpoint, another problem with low hanging fruit is that—by definition—there cannot be very many of them in a decomposable landscape, since the more there are, the lower the average impact of each on fitness. For example, the maximum average impact of 5 low hanging fruit is 20% of the total fitness range, while the average if there are 10 such variables is 10%. Thus researchers will quickly find themselves running out of high impact variables and will necessarily turn their investigations towards individual variables have much smaller impacts, often less than what might be considered clinically significant.

Large Sample Size. In the lower left-hand quadrant of the Figure 8 grid, we have variables with both low main and interaction effects. In their book *The Cult of Statistical Significance*, Ziliak and McCloskey (2008) particularly disparage this region. Their disapproval rests on the fact that we have come to accept significance testing as the gold standard for evidence, with $p < 5\%$ being the most common test. Thus, the goal of establishing such significance has trumped the question of whether or not the impact is significantly material. From a statistical standpoint, it is nearly always possible to achieve such significance if the sample is large enough. Thus, in this region we tend to detect relationships that are so small in their impact on fitness that they would not be

worth acting upon in practice. While the realm may be fertile grounds for the testing of theory—since there is essentially no limit to the number of minor impacts on fitness that can be detected—working in this region almost guarantees that research outcomes will not be useful since their impact tends to be too small to justify action. To compound the problem, operating in this region we are motivated towards activities leading to the identification or acquisition of large data samples, rather than towards solving problems of pressing interest to practice. Over the past decade, for example, MIS has developed a passion for studying online auctions. Researchers in the area have confided to me that the topic—one that seems to have little to do with MIS—is made particularly attractive by the large, high quality data sets that are readily available for analysis.

Drop Outlier Observations. The upper right quadrant of Figure 8 represents a region where characteristics exert a decomposable impact on fitness but also exert an impact as part of an interaction that may, in some cases, oppose the main effect. For example, in the NK fitness model, terms containing x_i might include:

$$c_i x_i \text{ and } f(x_1, x_2, \dots, x_i, \dots, x_K)$$

And, for some sets of values $\{x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_K\}$, the effect of changing x_i within $f(x_1, x_2, \dots, x_i, \dots, x_K)$ could oppose and exceed the main effect (c_i).

In cases such as these, strong main effect might dominate a given variable's impact across a sample but, from time-to-time, unusual observations would occur where the effect was absent or opposite. Continuing with the earlier theory of planned behavior example, it is generally assumed that the usefulness of an IT application will increase the user's disposition to use it. Occasionally, however, circumstances will arise where the user feels threatened by such a system's potential to eliminate—or, at least, deskill and devalue—his or her job; under such circumstances, the greater the usefulness of the system, the greater the threat and likelihood of user resistance (see Gill, 2010, pp. 348-349 for an example).

When the goal of research is to test of the truth of a model, the researcher may reasonably choose to treat occasional observations that do not follow a predicted pattern for an identifiable reason as outliers. In the interest of full disclosure, the researcher may—and probably should—describe the outlier observations and explain the reason they were omitted from the broader analysis. This procedure raises two problems in terms of the usefulness of results. First, by focusing only on generally observed main effects of material size, we return to the problem of low hanging fruit; such results are unlikely to be a surprise to practicing managers. Second, because high interaction effects are precisely what tend to produce diversity in a fitness landscape, focusing on main effects will tend to lead the researcher to ignore the less explored regions of the fitness landscape in their efforts to exhaustively explore the theory-friendly central peak. When discontinuities cause the environment to change dramatically, however, it is precisely these less explored peaks that are likely to contain the new dominant entities in the altered environment.

Ugly Theory and Statistical Illusions. The bottom right quadrant of Figure 8, where a variable's effect is large in interaction with other variables, but small individually, presents the most problematic quadrant. It is, however, also the quadrant that offers the greatest potential for useful discoveries. What makes it problematic is the large number of attribute combinations that must be studied. In the NK landscape model, for example, we saw that if N binary (0 or 1) attributes interact, there are 2^N combinations that must be examined. The theory that fully describes such a landscape will necessarily be huge; if the landscape is dynamic, the relationships will also be changing with time. This can only lead to what I describe as *ugly theory*.

The potential for usefulness offered by this region is also substantial—for precisely the same reason that the theory is ugly. As the number of combinations grows, the likelihood that even the most expert practitioners will have overlooked some high fitness peaks is great; indeed, the wide-

ly observed phenomenon of expert entrenchment (Dane, 2010) makes that likely. Thus, the researcher may be able to identify unexplored positions on the landscape.

Research in the bottom right quadrant of Figure 8 is further complicated when the entities being studied migrate to local fitness peaks, as they would be expected to do. In the pure case, the conceptual conundrum this poses is clear. If every entity is on a peak, then any “simple” recommendation made by a researcher will already be in place or, if implemented, will lead to a decline in fitness. That conclusion is purely a result of the definition of a peak. In either case, the informing value of the simple recommendation will likely be zero (recall Figure 4).

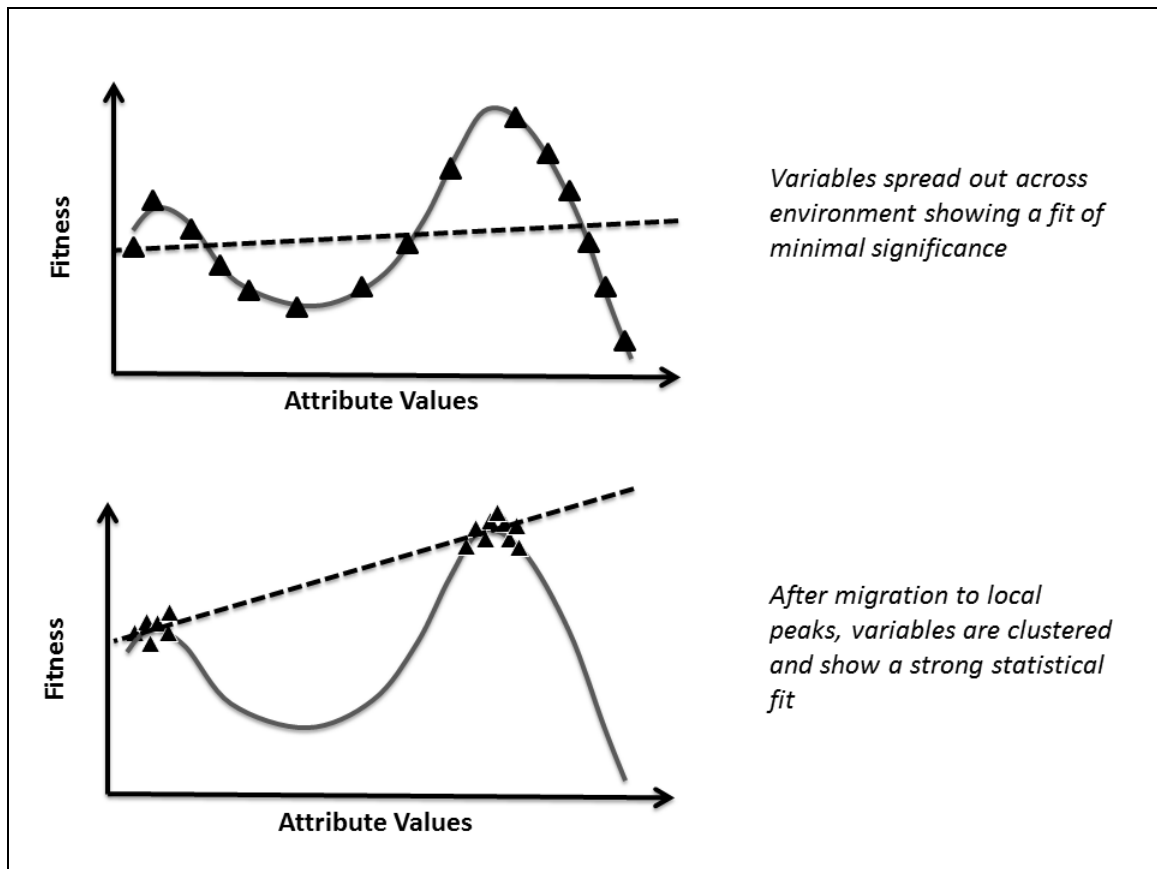


Figure 9: Statistical illusions that can result from observations becoming clustered around peaks

The high interaction/low main effect region is also prone to statistical illusions. Although a more detailed development of the problem involving simulated landscapes is presented elsewhere (Gill & Sincich, 2008), the basic nature of the illusions can be discerned by thinking in terms of the terrain conceptual scheme. Where entities are scattered randomly around a chaotic terrain, attempts to apply statistical tools such as multiple regression will produce only coincidental significances—just as we would expect (and hope) they would. When entities migrate to peaks, however, the situation changes. Tools such as multiple regression and structural equation modeling rely upon the variability between observations in order to estimate the reliability of their estimates. Imagining the multi-dimensional relationship collapsed into two dimensions, as shown in Figure 9, helps clarify why variables clustered on or around peaks are likely to exhibit significantly lower variability than widely dispersed observations. In essence, peak combinations of attributes are repeated many times after the migration to fitness has occurred. These are not readily detectable by observation because: a) there are likely to be many peaks (unlike the simplified 2-peak illustration presented in Figure 9), and b) decomposable subsets of variables—each seeking

peaks independently—are likely to exist in all but the pure chaotic landscape. As a result of the clustering, multivariate tools will be prone to return estimated significance values for coefficients suggestive of strong relationships when applied to observations drawn only from peaks. Moreover, since more entities tend to migrate to the highest peaks than to lower peaks (Kauffman, 1993), the sign of these relationships will often be consistent with attribute values associated with high fitness observations, making the relationships appear to be plausible. These significances, however, tell us little that would be useful in understanding the true underlying landscape.

Being crowded out of channels to practice

The second drawback of theory in a complex environment involves communications. A core principle of informing science is that there are limits on our ability to attend to the various channels through which we are informed. If we focus our research exclusively on establishing truth (i.e., theory), without concern for usefulness, as the landscapes we research become more rugged, our favored channels will tend to contain more and more content that is true but not useful to a particular practitioner (a consequence of Figure 7). As a result, we would expect practitioners to devote less attention to these channels.

In the MIS research area, there is some evidence for such devolution. In the early days of the discipline, the late 1970s, the practitioner-focused *Society for Information Management* (SIM) was heavily involved in the establishing and publicizing what became the discipline's premier journal, *MIS Quarterly* (*MISQ*). During those years, the percentage of articles that included at least one practitioner author was typically above 40% and SIM members were provided with a subscription to the journal as part of their membership. By the mid-1990s, however, the percentage of practitioner-authored articles dropped to under 20% and the *MISQ* subscription was unbundled from SIM membership, after which most members promptly dropped their subscription. During the past decade, practitioner contributions to *MISQ* have been almost non-existent—well under 5% on average and 0% most years (Gill & Bhattacharjee, 2009).

Contrast the academic journal channel to practice in MIS—which appears to be withering away—with another channel, that of third party MIS consultants. As a group, consultants are characterized as being less concerned with justification and more concerned with action (Benbasat & Zmud, 1999). Their services are widely used by companies, particularly where complex software technologies such as Enterprise Resource Planning (ERP) are being implemented (Ko, Kirsch, & King, 2005). The industry experienced “explosive growth” (Johnston, 2003) during the same period of the 1980s-1990s when the MIS academic journal channel appeared to be losing influence. While the consulting industry did experience a substantial setback after the “Internet Bubble” burst and Y2K concerns abated at the turn of the millennium, it has subsequently come roaring back and has found new drivers of demand, such as security (Johnston, 2003; Nolan & Bennis, 2003).

From an informing system standpoint, the consulting channel has a substantial advantage over any mass-media channel such as academic journals. Whereas the mass-media channel cannot take into account variations in pre-existing client models—and is therefore prone to conveying what is already known or what contradicts local experience (both of which interfere with the sender's credibility)—the consulting channel is inherently bi-directional in its communications, providing the consultant with the ability to assess the client's local knowledge prior to suggesting actions. This advantage naturally grows with the ruggedness of the landscape.

Tendency towards self-absorption

The last area that I will discuss is the potential of a rugged-landscape to produce communities of theory-focused researchers who communicate mainly with each other. An example of this phe-

nomenon is the development of disciplinary silos in academia, as well as the failure of applied disciplines, such as those that study business, to inform practice or other disciplines.

The rugged landscape model can help explain how such self-absorption could evolve. As the NK landscape conceptual scheme shows us, one of the key effects of ruggedness is a dramatic expansion in the likely size of any broadly applicable theory. To avoid addressing the entire landscape, it is natural for us to specialize. There is, however, a great benefit in having colleagues who specialize in the same region of the landscape, as we tend to judge the effectiveness of our research by metrics that measure its adoption by other researchers, such as citation counts. Targeting a group who shares our research values also means that the large set of assumptions upon which we necessarily base our research does not need to be justified every time we make a point. Moreover, working in silos prevents us from having to confront sets of assumptions that contradict our own. In economics, for example, *homo economicus*—the economic man assumed by most “theory”—knows everything and makes choices that are purely rational (Thaler, 2000). In MIS, on the other hand, we assume that managers can derive important insights from being told that they should “strategically focus on key improvement areas” while, at the same time noting that “non-focus metrics must still be managed” (Mitra, Sambamurthy, & Westerman, 2010, p. 44). It is hard to imagine how advice such as this could be useful unless we presume our clients to be novices or complete idiots. Precluded from the need to talk with economists as a consequence of our silos, neither discipline feels a pressing urge to address the fact that both extreme views are ludicrous.

The lack of concern with respect to informing practice and other disciplines within MIS has been amply documented elsewhere (e.g., Gill & Bhattacharjee, 2009). I cannot resist leaving the topic, however, without providing some concrete anecdotes of two aspects of this self-absorption that I have observed a number of times in conversations and correspondence with my professional colleagues: 1) the tendency to dismiss those who address the challenges faced by MIS practice from a different (i.e., non-theoretic) way than we do, 2) the tendency to abdicate responsibility for informing practice, either because: a) it is a problem that does not exist, or b) it is a matter for which we, as researchers, are not responsible.

With respect to the first issue, that of placing on low value on individuals who inform practice from a non-theoretic perspective, I am specifically thinking of the disregard with which many of us seem to hold IT consultants (see Baskerville & Myers, 2009, p. 652). A quote from an anonymous reviewer on a manuscript that I wrote demonstrates this attitude particularly well:

There are no particular pathological characteristics in the academic discipline of information systems that are not equally present in other academic disciplines. There may be, however, issues with the structure of IS practice. For example, the role filled by academic research in many disciplines, (e.g., medicine, sociology, psychology, etc.) is filled in IS by professional consultants. These are not just individual consultants, but big, wealthy consulting firms. The structure of IS practice includes an entire sub-industry dedicated to developing and delivering to the field knowledge in IT and its management. These firms depend on their own internally generated knowledge, which for want of proper research is perhaps a bad idea. Because academia lacks slick marketeers and constant face-time, it's not surprising that interest in academic research would fade away.

Competing head-to-head against the IT consulting industry is not the solution. Academia lacks the capital and the motive to play in the services market. However, it might be possible to compete internally to the IT consulting industry against their internally-generated knowledge. Academia indeed has better product to offer. It is important to recognize that such a shift is **not** a change in the academic discipline, it is a complex change in the social and cultural structures by which the IT consulting industry acquires its knowledge.

What this commentary suggests is that rather than ignoring practice, we were crowded out of our channels by better funded consultant competitors. Despite our lack of resources, however, our product is remains better than theirs. As an aside, if I were a Dean facing a chronic budget shortfall, upon hearing this argument I would immediately cut my MIS department's budget, since we claim to be able to maintain the quality of our efforts independent of the resources we are given.

With respect to the second issue, abdicating responsibility for informing practice, my most salient example came as a result of an email exchange I had with the Editor-in-Chief of one of the premier journals in the MIS field. In one email, I presented my argument that we must assume our research was having negligible impact on practice until otherwise proven. His reply included the following:

Why would we start with the null? That is good Popperian logic, but is contended by Cook and Campbell and most of the modern socio-psychological methodologists. They argue that we intrinsically want to confirm theories, not deny them. My theory is that we do impact practice. Since there is no good evidence against my theory, as I interpret it, then I will hold this view until there is compelling evidence against it.

Another line of argument—one that has been published in a number of articles—proposes that our failure to inform practice may be the result from shortcomings on the part of practice, rather than from any weaknesses in our informing activities. For example, consider the following from a commentary on Gill & Bhattacharjee's (2009) article on informing practice in MIS:

An unstated assumption in Gill and Bhattacharjee is that our informing problem lies in the kind of research being delivered into practice. But what if the problem is partly (or entirely) with IS practitioners, who fail to appreciate solid research results? ... Too many practitioners may not be sufficiently educated to comprehend the degree to which IS research addresses some of their most intractable issues. (Myers & Baskerville, 2009, p. 664)

As a point of contrast, consider the praise that the same authors heap upon the accomplishments of the discipline's premier journal:

MIS Quarterly is one of only two IS journals included in the *Financial Times* list of Top 40 academic journals in business. All the journals in this list are research, not practitioner, oriented. Achieving such a reputation did not happen because the evolving editorial objectives of *MIS Quarterly* took this journal down the wrong path. On the contrary, a focus on academically rigorous research has considerably enhanced the reputation of this journal and, by association, the field as a whole. *MIS Quarterly* now has an excellent, hard-won reputation, one of which all IS scholars should be proud. (Myers & Baskerville, 2009, p. 664)

Certainly such recognition should be a point of pride for those who measure achievement in terms of theory-building and testing. But what role does usefulness play when our community chooses to evaluate itself almost entirely on our ability to inform each other?

Even when we recognize the problem of practice being unable to make sense of our research, many of us take the position that the purity of our efforts should not be sullied by conveying it in a form suitable for a practitioner audience. Consider the following quote:

... I would not trade off academic writing style and demand that we write in a manner that is simple, concise and clear... Instead I would expect that we educate our practitioners to appreciate brilliant intellectual efforts! Several IS phenomena are hard to understand and may demand difficult and esoteric language because they cannot be couched in the "common language". Still, the message of a text written in an esoteric language can

be relevant. For example, the fashionable use of Heidegger in understanding design or use of IT is neither possible nor useful unless the reader can work through Heidegger's thick concepts and ideas. My nightmare would be to emasculate Heidegger and dress him into HBR format. (Lyytinen, 1999, pp. 26-27)

I have to admit that *my* nightmare would be to find myself trapped in a room filled with MIS research, unable to leave until I identified an indisputably brilliant intellectual effort. But, subjective assessments of research quality notwithstanding, if we are so self-absorbed that we believe we have no responsibility to frame our research in a form suitable for practice, the usefulness of our research—however true our theory may be—is moot.

The Drawbacks of Conceptual Schemes

The drawbacks of conceptual schemes are considerably simpler to specify because they are precisely the same as their strengths, namely:

1. Conceptual schemes are specifically intended to produce decisions, and are therefore likely to be of consequence
2. Conceptual schemes are typically “sticky”, meaning that once in place they can be very difficult to dislodge.

Unfortunately, these drawbacks can be HUGE. For example, it would be nearly impossible to name a war that was not justified based on one or more conceptual schemes that had little—if anything—to do with truth. Every invalid scientific paradigm that has ever existed was, by definition, a conceptual scheme. All the prejudices, biases, and irrational thought patterns that influence our actions in an undesirable way have, at their heart, flawed conceptual schemes.

The stickiness of conceptual schemes—even verifiably false ones—proves equally problematic. Bain (2004, pp. 22-23) describes a study in which physics students refused to give up their Aristotelian models of motion despite numerous experimental demonstrations that both proved that their pre-conceived models were wrong and offering the (correct) Newtonian model. Heath and Heath (2007) specifically explore the degree to which “urban myths” diffuse and are believed to be true despite their obvious implausibility.

Part of what makes conceptual schemes so sticky may be the fact that, unlike theory, they are not expected to be entirely true in all situations. The “exception proves the rule” adage has no place in theory testing; in conceptual schemes it represents a reasonable compromise. For this reason, demonstrating that a conceptual scheme is not true will *not* cause it to be abandoned under most circumstances. In fact, according to Kuhn (1970), falsification alone *never* leads to the abandonment of a paradigm, even in the hard sciences.

Striking an Balance between Theory and Conceptual Schemes

Given the extraordinary potential for damage of false conceptual schemes (e.g., witness the holocaust), it is not surprising that most academic researchers studying rugged domains choose to focus on theory. Indeed, a marked preference for avoiding Type I error (asserting something that is not true) at the risk of making Type II error (failing to assert something that is true) characterizes many of the social sciences (Ziliak & McCloskey, 2008). In the MIS field, this attitude was expressed in an article on improving our research relevance that, in discussing the topic of measuring user satisfaction with a system, made the following assertion:

To answer these questions correctly requires proper measures. However, measurement is a complex endeavor and poor measures are worse than no measures at all as they provide incorrect data for prediction, monitoring and evaluation. (Benbasat & Zmud, 1999, p. 10)

It is very easy to read a statement like this without giving it a second thought. But it is only valid to the extent that poor models are worse than no models at all. For example, imagine yourself walking down a dark two lane highway when suddenly you are blinded by incoming headlights in the lane where you are walking. You need to jump out of its way, either to the right or to the left. You are so startled, however, that you momentarily forget whether you are in the U.S. or the U.K.—meaning you are not sure whether jumping to the right (or left) will put you on the safe shoulder (good choice) or into another lane of oncoming traffic that you have not checked (poor choice). As it turns out, however, the worst possible choice is to remain in your current position because you can't decide which of the two alternatives is better, left or right.

As I mentioned earlier, economists have been providing public policy makers and investors with models that have proven to be breathtakingly inaccurate in their predictions (e.g., Taleb, 2007). Are we worse off than if we made these decisions with no model whatsoever? For that matter, is it even possible for an individual to hold “no model whatsoever” while making decisions?

Given that ignoring conceptual schemes may be tantamount to abdicating nearly all academic influence on practice in many disciplines, we need to strike a balance. What now follows are five recommendations that derive from the analysis I have presented.

Recommendation #1: Emphasize the top of the inverted U-curve

In an ideal world, as academic researchers we would always target the *True & Useful* region of Figure 3, where theory and conceptual schemes intersect. In nearly all cases, this intersection will involve topics where practitioners are uncertain, i.e., the top of the Figure 4 inverted U-curve of informing potential. Thus, before choosing research topics to investigate, it would be sensible to identify questions where practitioners do not already hold strong opinions. One way to accomplish this would be through holding focus groups with practitioners before engaging in a research project. Another might be to conduct a survey, where practitioners were asked to agree/disagree with clearly worded hypotheses that might be candidates for investigation.

Naturally, there will be situations where results are found that contradict strongly held practitioner views. A central point of this essay is that such opinions *cannot* be overcome by published theory alone. To be effective in such informing situations, we will need far greater leverage in the informing process (see Recommendations #2 through #5).

Recommendation #2: Respect the power of prevailing models

In my experience, there are few topics that evoke the typical academic's contempt more rapidly than “junk science”. I can think of any number of reasons why such an attitude makes sense. Junk science can be created and disseminated without rigorous peer review; consider the many popular unsubstantiated trade books that fill the airport bookseller's stalls. Junk science is often created as a money-making enterprise (with that aspect cleverly hidden); consider the “research” proliferated (and suppressed) by the tobacco industry in its heyday. Junk science frequently panders to the prejudices of the client; consider the recommendations of the diet industry. But where we should really be concerned is the way that junk science, in all its forms, can interfere with our ability to inform. Once a client holds a conceptual scheme—whether it be a product of junk science or rigorous investigation—before a new model can be adopted, the existing model must be confronted.

What I would therefore propose is that if we wish our research to be useful, we need to make an effort to understand the existing universe of conceptual schemes that is available to our practitioner clients, including those that may be the worst sort of popular tripe. In many respects, the least appropriate thing we can do is to dismiss such ideas as being beneath our notice. By doing so, we can easily undermine our own ability to be useful. For example, the late Everett Rogers (2003)

began his famous book *Diffusion of Innovation* with a story about the attempts to introduce better sanitation practices in a Peruvian village. Central to this effort was having residents boil their (unsafe) water before using it for drinking. Unfortunately, in this community, the prevailing conceptual scheme that had evolved over many generations equated the notion of “cold” with “healthy” and “hot” with “unhealthy”. As a consequence, boiling water was equated to unhealthy by most of the community. After an extensive campaign, only about 5% of the population followed the recommendation. In discussions with two key adopters—the “exceptions” to the general rule—it was found that one came from another village that did not hold with the hot/cold model and another was already sick, so boiling water before drinking it was consistent with existing social norms. We run the same risk if we attempt to disseminate our research to practice without understanding their existing models.

It would also make sense to use existing models—whatever their origins—as potential research targets. By doing so, we can become aware of their limitations and can explain these to the practitioner clients we are intending to inform with other findings. We must also accept the very real possibility (more likely a probability, in my opinion) that some of these non-academic models will prove valid under many circumstances and deserve more credit than we might otherwise give them. Just because the model involves riders and elephants or cheese that mysteriously moves does not mean that it cannot be very useful to the right person under the right circumstances.

Recommendation #3: Foster bi-directional channels to practice

As was noted earlier, an advantage that consultants have over (non-consulting) academics is that they typically engage in extensive face time with their clients, while academics tend to set publication as the goal to which they most frequently attend. Bi-directional informing channels allow the informer to assess pre-existing client models as part of the informing process. That frequently allows Figure 4 issues (what the client already knows and what the client is initially unwilling to believe) to be side-stepped.

The types of bi-directional channels that academics might develop have been discussed at length elsewhere (e.g., Gill, 2010). These include encouraging:

- Faculty consulting engagements, treating them in a manner equivalent to research grants.
- Action research, where the researcher is an active participant in the project being investigated.
- Case writing, where the researcher documents complex situations in the field, either for the purpose of developing a teaching artifact or for a research project.
- Professional doctoral programs, temporarily bringing high level practitioners into the academic environment in order that they become exposed to research findings and, later, become available to disseminate research to practice.
- Extended part-time doctoral programs, where practitioners—perhaps right out of an undergraduate or master’s program—work and, at the same time, take a course or so each year, providing a channel to practice that could last a decade or more.
- Researcher sabbaticals with practice.
- Participation in professional service and networking organizations.

Whether we continue to focus our research on truth (at the risk of being useful only occasionally) or on being useful (at the risk of finding ourselves asserting things that later prove to be false), most informing to practice will not take place absent the establishment of channels such as these.

Recommendation #4: Value the creation of sticky messages

If a substantial portion of our mission evolves to dislodging ill-thought-out conceptual schemes introduced by charlatans or just by mistake, we will not succeed if the message we convey is not equally sticky or more so.

Heath and Heath's (2007) earlier mentioned SUCCEs framework (simple, unexpected, concrete, credible, emotional, and stories) is a good starting point for the creation of such messages. The final item, stories, may prove to be particularly important in this context. That a message framed as a story has unique resonance has been noted in many different contexts (e.g., Akerlof & Shiller, 2009; Gill, 2010; Guber, 2007; Rynes, Giluk, & Brown, 2007; Willingham, 2009). Part of this stickiness may stem from two characteristics possessed by stories in general:

1. The elements of causality, conflict, communications, and character embedded in many stories provide built-in linkages that facilitate cognitive storage and later retrieval (Willingham, 2009).
2. Because a story does not directly attack the listener's models, and because we are used to hearing both factual and fictional stories, it is less likely to be treated as a challenge to existing knowledge (e.g., the left hand side of the Figure 4 informing potential inverted-U). Instead, a client may pick and choose what is considered germane to his or her specific situation.

Thus, conducting research in a manner that leads to the creation of sticky stories should ultimately increase the likelihood that what has been discovered will become useful to other clients.

It is important to emphasize that a sticky conceptual scheme—whether it be in story or other form—will not, in all likelihood, be sufficient to ensure informing in a rugged landscape. Complex ideas have never diffused well on their own. If combined with the other recommendations, particularly the establishment and maintenance of rich bi-directional channels to clients, then there is a high likelihood that the preconditions for informing will be met.

Recommendation #5: Partner with the wisdom of practice

Following recommendations 1, 2, & 4 would present an obvious danger: that poorly validated, but nevertheless sticky, conceptual schemes will escape into practice, influencing them to take poor or self-serving actions that negatively impact their organizations or the world at large. We, as academics, have been trained to hate Type 1 errors. What happens if our advice is taken and it happens to be wrong?

If we are to reconcile our fear of making mistakes with our desire to have impact, we can only do so by reimagining the informing systems in which we participate. Conducting research in our individual silos and communicating our knowledge mainly to students and academic peers, it is relatively easy to cast ourselves in the role of informer. In doing so, we naturally conceive of practitioners in the complementary role of intended recipient for what we have learned. Under such a conceptual scheme, the peril of misinforming is clearly one that needs to be avoided.

There are two distinct problems with the informer-recipient model of this particular informing system. First, it seems to misstate how the process actually occurs in the real world. At least in the case of management ideas, to the extent that informing takes place across the academia-practice chasm, the bulk of knowledge appears to be moving *from practice to academia* (Barley, Meyer, & Gash, 1988). Second, symptomatic of our previously mentioned tendency towards self-absorption, such a model tends to foster the paternal view that practitioners are roughly equivalent to students, needing to be protected from any errors we might make.

If we are to be effective in informing practice, we need to recognize that in a rugged landscape, most “truth” will tend to be local. That means coming to terms with the fact that the degree of truth underlying our conceptual schemes will vary according to the context experienced by a particular practitioner. What we also need to recognize, however, is that the practitioners with whom we *exchange* ideas are not passive vessels waiting for us to provide a conceptual scheme that they can use. Rather, they will interpret and shape any scheme that we offer to fit the local reality that they have experienced. They will accept those elements consistent with their experience, reject those elements that are not and, if our message is sticky, carefully consider those elements that they had not previously thought about. The researcher/practitioner relationship, therefore, should not be viewed in terms of sender/receiver; even informer/client may misstate the directionality of informing. Rather, to work effectively, we must view the system as a true partnership in which both sides become equally enriched through the informing process. In MIS, we have long known that the most effective systems generally emerge when developers and users collaborate in the creation of the system. There is little reason to expect that a conceptual scheme developed through research conducted in the absence of expert involvement will fare any better an information system foisted on disengaged users.

Conclusions

The central point of this essay has been that what is true is not always useful, while what is useful is not always true. As academics, we should not reflexively opt for truth without, at least, considering the implications of usefulness as a goal. That I have chosen to label models that focus on truth as *theory*, while those that focus on usefulness as *conceptual schemes* is largely a matter of convenience. If you happen to disagree with the labels that I have used, do not let that obscure the fundamental point that I have tried to make.

The second key point I have emphasized is the role that the ruggedness of the domain we study plays in striking an appropriate balance. On a rugged landscape, the relationships that hold true around one combination of factors may be largely irrelevant to the entity focusing on an alternative peak; as Michael Porter (1985) asserts, the company that seeks cost leadership in an industry needs to play by a very different set of rules than the company following a differentiation strategy. On a rugged landscape, relationships that are true one minute may cease to be valid the next, as changes in the environment alter the pathways to success. On a rugged landscape, the size of the theory needed to describe the entire topography is larger than what we can reasonably expect to communicate. In contrast, the size of a conceptual scheme is determined by the capabilities of the channels being employed and by the cognitive capacity and pre-existing models of the client, not by the underlying landscape function. For these reasons, conceptual schemes seem particularly well suited to rugged landscapes.

Conceptual schemes differ from theory in a number of ways. Conceptual schemes need to be resonant in form (a.k.a., sticky), since a scheme that communicates poorly is unlikely to be used. Theory takes whatever forms the truth demands. Conceptual schemes need to be malleable in their interpretation, customizable to the existing models of each client. The best theory will be unambiguous. Conceptual schemes adapt to a changing world. Theory is supposed to reflect unchanging principles. Conceptual schemes are best created in partnership with the clients who will apply them. The development of theory makes no such demands.

The choice of truth vs. usefulness will always be a matter of degree. For academic researchers in particular, establishing the truth of what we say will always be a major concern. I would not have it any other way. Models developed without concern for truth are far too likely to misinform; sadly, they are often applied to manipulate clients, to disinform them. As academics, we must never relax our guard against such snake-oil. In doing so, however, we must also acknowledge that slavish devotion to truth is almost as bad. By allowing ourselves to become completely uncon-

cerned with usefulness in any practical way (even as we pay lip service to it), we do not solve the problem of bad conceptual schemes reaching practice. We simply abdicate responsibility for such schemes and, in so doing, allow them to circulate unhampered by our inspection and intervention. To the outside observer, such as a state funding agency, it is very hard to justify a continuing stream of resources to a research endeavor that dismisses the goal of usefulness based upon its principled commitment to truth. Thus, for those of us who prefer not to finance our search for truth entirely out of our personal pockets, the time has come for the pendulum to swing back in the direction of usefulness. It is time that we grant conceptual schemes the respect that they deserve, even as we continue to relish our role of debunking those schemes whose application is proven harmful.

References

- AAAS (2010). *Q & A on evolution and intelligent design*. American Association for the Advancement of Science. Retrieved 11/9/2010 from http://www.aaas.org/news/press_room/evolution/qanda.shtml
- Akerlof, G. A., & Shiller, R. J. (2009). *Animal spirits: How human psychology drives the economy and why it matters for global capitalism*. Princeton, NJ: Princeton University Press.
- Ariely, D. (2009). The end of rational economics. *Harvard Business Review*, July-August, 1-7 (HBS reprint page numbers).
- Aune, B. (1987). Conceptual relativism. *Philosophical Perspectives*, 1(Metaphysics), 269-288.
- Bain, K. (2004). *What the best college teachers do*. Cambridge, MA: Harvard University Press.
- Barley, S. R., Meyer, G. W., & Gash, D. C. (1988). Cultures of culture: Academics, practitioners and the pragmatics of normative control. *Administrative Science Quarterly*, 33(1), 24-60.
- Baskerville, R. L., & Myers, M. D. (2009). Fashion waves in information systems research and practice. *MIS Quarterly*, 33(4), 647-662.
- Benbasat, I., & Zmud, R. (1999). Empirical research in information systems: The practice of relevance. *MIS Quarterly*, 23(1), 3-16.
- Dane, E. (2010). Reconsidering the trade-off between expertise and flexibility: A cognitive entrenchment perspective. *Academy of Management Review*, 35(4), 579-603.
- Davidson, D. (1973/1974). On the very idea of a conceptual scheme. *Proceedings and Addresses of the American Philosophical Society*, 47, 5-20.
- Davis, M. (1971). That's interesting!: Towards a phenomenology of sociology and a sociology of phenomenology. *Philosophy of the Social Sciences*, 1, 309-344.
- Gill, T. G. (2010). *Informing business: Research and education on a rugged landscape*. Santa Rosa, CA: Informing Science Press.
- Gill, T. G., & Bhattacharjee, A. (2009). Whom are we informing? Issues and recommendations for MIS research from an informing sciences perspective. *MIS Quarterly*, 3(2), 217-235.
- Gill, T. G., & Sincich, A. (2008). Illusions of significance in a rugged landscape. *Informing Science: the International Journal of an Emerging Transdiscipline*, 11, 197-226. Retrieved from <http://inform.nu/Articles/Vol11/ISJv11p197-226GillIllusions.pdf>
- Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly*, 30(3), 611-642.
- Guber, P. (2007). The four truths of the storyteller. *Harvard Business Review*, December, 52-59.
- Haidt, J. (2006). *The happiness hypothesis: Finding modern truth in ancient wisdom*. New York, NY: Basic Books.
- Haidt, J. (2010, October 16). What the Tea Partiers really want. *Wall Street Journal*.
- Hayden, T. (2009). What Darwin didn't know. *Smithsonian*, 39(11), 41-48.

- Heath, C., & Heath, D. (2007). *Made to stick*. New York, NY: Random House.
- Heath, C., & Heath, D. (2010). *Switch: How to change things when change is hard*. New York, NY: Broadway Books.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75-106.
- Hubbard, R., & Armstrong, J. S. (1992). Are null results becoming an endangered species in marketing? *Marketing Letters*, 3(2), 127-136.
- Johnston, S. J. (2003). The future of IT Consulting: Q&A with Richard L. Nolan. *HBS Working Knowledge*, 31 March. Retrieved from <http://hbswk.hbs.edu/item/3402.html>
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.
- Kauffman, S.A. (1993). *The origins of order*. Oxford, UK: Oxford University Press.
- Ko, D., Kirsch, L. J., & King, W. R. (2005). Antecedents of knowledge transfer from consultants to clients in enterprise system implementations. *MIS Quarterly*, 29(1), 59-85.
- Kuhn, T. S. (1970). *The structure of scientific revolutions* (2nd ed., enlarged). Chicago, IL: University of Chicago Press.
- Lyytinen, K. (1999). Empirical research in information systems: On the relevance of practice in thinking of IS research. *MIS Quarterly*, 23(1), 25-28.
- Maslow, A. H. (1943). A theory of human motivation. *Psychological Review*, 50, 370-396.
- McKelvey, B. (2002). Model-centered organization science epistemology. In J. A. C. Baum (Ed.), *Companion to organizations* (pp. 752–780). Thousand Oaks, CA: Sage.
- Mitra, S., Sambamurthy, V., & Westerman, G. (2010). *How do CIOs measure and communicate IT performance?* SIM Advanced Practices Council.
- Myers, M. D., & Baskerville, R. L. (2009). Commentary on Gill and Bhattacharjee: Is there an informing crisis? *MIS Quarterly*, 33(4), 663-665.
- Nolan, R., & Bennigson, L. (2003). Information technology consulting. *HBS Working Paper #03-069*.
- Orlikowski, W. J. & Baroudi, J. J. (1991). Studying information technology in organizations: Research approaches and assumptions. *Information Systems Research*, 2(1), 1-28.
- Porter, M. E. (1985). *Competitive advantage*. New York, NY: Free Press.
- Roethlisberger, F. (1977). *The elusive phenomena*. Boston, MA: HBS Press.
- Rynes, S. L., Giluk, T. L., & Brown, K. G. (2007). The very separate worlds of academic and practitioner publications in human resources management: Implications for evidence-based management. *Academy of Management Journal*, 50(5), 987-1008.
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). New York, NY: Free Press.
- Shanteau, J. (1992). Competence in experts: The role of task characteristics. *Organizational Behavior and Human Decision Processes*, 53, 252-266.
- Taleb, N. N. (2007). *The Black Swan*. New York, NY: Random House.
- Tetlock, P. E. (1999). Theory driven reasoning about plausible pasts and probable futures in world politics: Are we prisoners of our preconceptions? *American Journal of Political Science*, 43(2), 335-366.
- Thaler, R. H. (2000). From homo economicus to homo sapiens. *Economic Perspectives*, 14(1), 133-141.
- Weinberg, P. (2001). The relevance of MIS research to the business community. *Communications of the Association for Information Systems*, 6(10), 99-101.
- Willingham, D. T. (2009). *Why don't students like school?* San Francisco, CA: Jossey Bass.

Ziliak, S. T., & McCloskey, D. N. (2008). *The cult of statistical significance*. Ann Arbor, MI: University of Michigan Press.

Acknowledgements

I would like to thank Zbigniew Gackowski for his careful reading and commentary on the article prior to its final publication. I would also like to thank the 5 anonymous reviewers for the InSITE 2011 conference for their suggestions and encouraging remarks.

Biography



Grandon Gill is a Professor in the Information Systems and Decision Sciences department at the University of South Florida. He holds a doctorate in Management Information Systems from Harvard Business School, where he also received his M.B.A. His principal research areas are the impacts of complexity on decision-making and IS education, and he has published many articles describing how technologies and innovative pedagogies can be combined to increase the effectiveness of teaching across a broad range of IS topics. Currently, he is Editor-in-Chief of *Informing Science: The International Journal of an Emerging Transdiscipline* and an Editor of the *Journal of IT Education*.