Research Themes in Complex Informing

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

Eli Cohen
**Informing Sciences Institute,
Santa Rosa, CA, USA**

ggill@coba.usf.edu

EliCohen@InformingScience.org

Abstract

This paper introduces a six paper series that examines the manner in which complexity impacts the informing process. Two of these papers specifically consider how the objective complexity of the domain being studied changes the nature of the solution—with domains consisting of many interacting elements and changing criteria for success tending to produce highly rugged fitness landscapes that violate the normal assumptions of decomposability that we make in our research and incorporate into our theories. A third paper considers how the concept of utility impacts informing and is impacted by complexity. Another pair of papers examines, first, how the structure of a client's mental models—referred to as structural complexity—changes with repetitive task performance and, then, how various cognitive filters work in concert with this structure during the informing process. The final paper specifically considers client-to-client informing processes, using the diffusion literature to argue that they are indispensible to complex informing and further mapping different models to different levels of informing complexity. Collectively, it is hoped that these papers will spur further research into the processes that enable complex informing.

Keywords: research methods, rigor, relevance, rugged fitness landscapes, adaptation, utility, learning, informing sciences, complexity, bias, decomposability.

Introduction

If every message that is passed from sender to client were received and understood correctly and unambiguously, there would be no need for the informing sciences. However, informing is complex and complexity interjects itself into many informing situations. Where it does, it can increase the challenges of effective informing in many ways. How do we present our message so the client accepts it? Does the client understand the message as we intended it to be understood? Importantly, how do we ensure that our messages, once understood by the client, will have the impact that we desire? How do we create messages that inform multiple clients? Can we get clients to pass our messages among themselves? All these questions become more difficult when complexity is present.

Let us consider a model for the informing process that builds on the Shannon and Weaver Model of Communications (1949) and Wilson's (1981) model of information-seeking behavior, shown as Figure 1.

It is useful to relate to Figure 1 as we introduce the following papers. These papers all relate to different facets of task complexity within a given context.
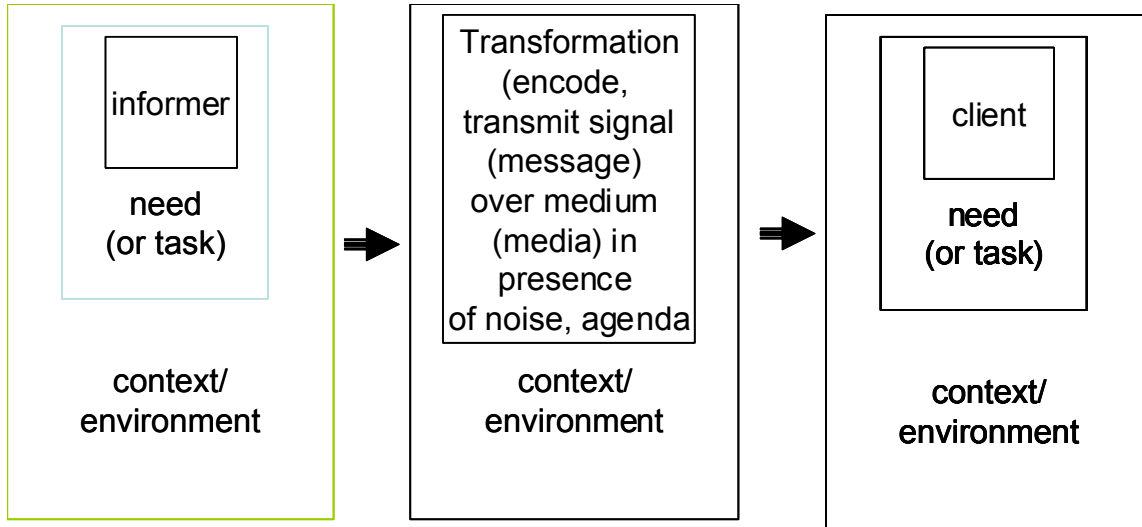
Figure 1. The Informing Science Model builds on the Shannon and Weaver (1949) and Wilson (1981) models. The emphasis in this model is on the context of the client, the informer, and the transformation of message between the two.

Despite its potential importance to informing processes, the topic of complexity has not heretofore been systematically investigated within the informing sciences literature. The series of six papers being introduced here represents our attempt to introduce some key complexity-related concepts and to promote further research in the area of complex informing. Because, as it turns out, complexity—in its various forms—can exert a sizeable impact on informing. For example, our research suggests the following:

- The level of complexity associated with an informing system can dramatically alter the forms of research that are most appropriate for investigating the system's behavior and the types of results we can expect to find (Gill, 2008b).

- That same complexity can cause traditional quantitative techniques to yield misleading results about the systems we are investigating, making us overconfident in the relationships that we discover and suggesting the presence of relationships that are pure illusions (Gill & Sincich, 2008).

- Changing levels of complexity in a client's mental models can exert a strong influence on how clients value information and in the degree to which inconsistencies in client decision-making are likely to be experienced (Gill, 2008a).

- Complexity changes that occur as we acquire expertise can dramatically impact how clients represent and share information, leading to substantial challenges when experts with one type of expertise attempt to share their knowledge with experts having alternative types of expertise (Gill, 2008e)

- The complexity of information being shared changes the nature of the barriers to informing that we encounter when a single client is being informed (Gill, 2008d).

- Where complex information is to be shared between two dissimilar communities, client-to-client informing processes are likely to be far more critical to the success of the process than sender-to-client processes (Gill, 2008c).

Many of these findings presented are not original. Rather, they are distilled from a broad range of research fields that include economics, evolutionary biology, cognitive science, decision theory,

psychology, sociology, and management. The hoped-for contribution of this series of papers, then, is to bring research into complex informing processes to the informing sciences mainstream.

What is complexity? Unfortunately, the concept of complexity is far from simple. As a result, before we consider how the papers tie together, we need to clarify what we mean by complexity.

Introduction to Task Complexity

The initial inspiration for the series of articles being presented is a 2006 article in *Informing Science* titled "Task Complexity and Informing Science: A Synthesis" (Gill & Hicks, 2006). Based upon their analysis of several hundred articles that either defined or applied the task complexity construct, the authors concluded that there was no possible consensus definition. Instead, there were five general classes of complexity definitions, each of which could be useful under certain circumstances: experienced, information processing, problem space, structure, and objective. These classes are summarized in Table 1.

Of the five classes, the first two define task complexity in terms of its consequences. They prove not to be particularly useful for our present purposes. The remaining three, however, can each contribute to informing in their own way. *Problem space complexity* constructs, for example, characterize the nature of the task performer's pre-existing knowledge and mental models. For example, if the task performer were a computer program, problem space complexity metrics might include the number of lines of code, the number of paths through that code, the amount of memory required to hold the code, the nature of the variables used to hold information during processing, the initial values of these variables, and so forth. *Structural complexity*, on the other hand, takes a more qualitative look at the problem space and attempts to characterize the types of knowledge to be used in task performance. For example, does the task performer apply task-specific rules or rely upon more general common sense or analogy? Is there a clear task goal or is task performance guided by a more general collection of vague goals? Is performance automatic, or does the task performer need to make conscious decisions during each step in the task?

Table 1: Task Complexity Classes (summarized from Gill & Hicks, 2006)

Name	Form of Definition	Description/Example
Experienced	Task Complexity → Psychological state	If an individual perceives a task to be difficult, then the task is complex
Information Processing	Task Complexity → Information Processing Activity	If a task a task produces high information processing, then the task is complex
Problem Space	Problem Space Attributes → Task Complexity	*Example*: A task's complexity is defined by the minimum size of the computer program required to perform the task.
Structure	Lack of Structure → Task Complexity	*Example:* The more routine a task, the less complex it is
Objective	Task Characteristics → Task Complexity	*Example:* A task's complexity is determined by the number of task elements, the degree of interrelationship between the elements and the degree to which task objectives are changing (Wood, 1986).

Objective complexity, which proposes task complexity to be a strict function of task characteristics, seems—intuitively at least—as if it might be the most useful class of the bunch. Its particular strength is that it does not depend upon knowing the mental models of the task performer. Also, it

is intended to predict task complexity, unlike the first two classes, which merely report its presence.

However, it has never been very clear what objective complexity was good for. The problem becomes particularly acute when task performers are not told how to perform a task. Countless experiments, for example, find that as a task becomes more objectively complex, performers start choosing ways to simplify how they perform the task. That breaks down any desired relationship between objective complexity and experienced or information processing measures. Moreover, the task may not specify what tools are to be used in performing it. Imagine how the nature of what is experienced by the task performer can change depending upon the presence or absence of a computer. It also turns out that it is very hard to measure the level of objective complexity, although attempts to devise formulas have been made (e.g., Wood, 1986).

For these reasons, objective complexity never gained much traction as a construct. Yet, it turns out that the objective complexity construct proposed by Wood (1986) is very good for predicting the ruggedness of a fitness landscape. The importance of that is the subject of the first paper in the series.

Reflections on
Researching the Rugged Fitness Landscape

The first paper in the series (Gill, 2008b) draws upon the concept of a rugged fitness landscape, pioneered by evolutionary biologist Stuart Kauffman (1993). It is based upon the concept of a fitness function, which is a mapping between a set of attributes and the resulting desirability or survivability of the associated entity (i.e., its fitness). The concept of a fitness function is widely distributed across many disciplines. In computer science, for example, it appears as the static evaluator function often encountered in game playing programs. In economics and finance, it is called the utility function. In the military, annual performance ratings of personnel are referred to as fitness reports. Kauffman's research focused on understanding the overall behavior of such functions across their entire domain, referred to as their *fitness landscape*. What he found was that as the number of attributes and their interrelatedness grows, the landscape becomes increasingly more rugged. What ruggedness means is that instead of having a single fitness peak, achievable by tuning the fitness of each attribute independently, you have many local peaks—each of which depend upon specific combinations of attributes.

Conceptually, it is useful to think of a rugged landscape as being similar to rough terrain, such as that illustrated

Figure 2: The rugged landscape of Bryce Canyon, Utah, USA (photo by Grandon Gill)

in Figure 2. As is the case in that illustration, attempting to reach peak fitness by continuously moving upwards is unlikely to get you to the highest peak—but it will get you to the nearest local peak in fitness. Unlike the illustration, however, a fitness landscape is not two dimensional. Rather, it has as many dimensions as the fitness function has arguments.

The particular challenge presented by rugged landscapes is that the impact of variables upon fitness is not decomposable. In Figure 2, for example, the elevation function has two arguments, the N-S dimension (latitude) and the E-W dimension (longitude). As should be readily evident, however, if we assume fitness to be the same as altitude, travelling north will sometimes take us up (to higher fitness) and sometimes take us down (to lower fitness), the same being true for travelling east. Thus, the rules regarding what direction to travel so as to improve fitness depend entirely upon where you are in the landscape. Or, stated another way, the more rugged the landscape, the more likely its behavior will be localized and not particularly generalizable.

An example of such localized behavior presented in the paper involves the design of three courses, each of which was characterized as exemplary using a variety of criteria—the details of which are presented in the paper. What is interesting about these courses is that there is not a single design characteristic that they share in common (See Table 2). In a decomposable landscape, this would suggest that none of the characteristics listed make a particularly strongly contribution to course fitness. While that is possible, it seems quite unlikely given the range of characteristics presented. In a rugged fitness landscape, on the other hand, the conclusion that we would draw would be that the characteristics all (or nearly all) contribute to fitness, but only in combination with other characteristics. For example, mandatory attendance might not make sense in combination with flexible deadlines, since classroom presentations and assignments could then easily get out-of-synch.

Table 2: Cross-Course Comparison (originally from Gill & Jones, 2008)

	Class A	Class B	Class C
Classroom Lectures	No	Yes	Minimal
Multimedia Lectures	Yes	No	No
Moderated Classroom Discussions	Optional	No	Yes
Paired Student Problem-solving	No	Yes	No
Student Presentations	No	No	Yes
Deadline Flexibility	Yes	No	No
Mandatory Attendance	No	Yes	Yes
Examinations	No	Yes	No
Outside Class Projects	Yes	No	Yes
Level of Performance Feedback	High	High	Low
Grade Subjectivity	Low	Low	High
Student Level	Undergraduate	Undergraduate	Graduate
Source	Evolved	Designed	Designed
Instructor	Instructor A	Instructor B	Instructor A
Instructor Experience with Course Subject Matter	High	Low	High
Evaluations (a possible indicator of fitness)	Outstanding	Outstanding	Outstanding

There are four main conclusions supported by the paper:

1. That ruggedness in a fitness landscape fundamentally changes how we need to understand the underlying process behind the landscape. Decomposable landscapes lend themselves to being explained by *attractive theory*, which is to say theory that is compact, generalizable, and can be reproduced. Rugged landscapes, on the other hand, tend to yield *ugly theory*, which is to say large and full of qualifications, not generalizable, and hard to reproduce outside of the situation in which the original observation was made.

2. That high objective complexity in a task domain produces a rugged fitness landscape that impacts how we perform the task.

3. That the typical informing system meets all the prerequisites for objective complexity—many attributes that can affect the fitness of the system (e.g., motivation sources, client learning styles, client and sender problem space characteristics, delivery system characteristics) and many characteristics that need to be matched with other characteristics (e.g., client information display preferences with system knowledge presentation format) creating numerous interrelationships.

4. That our understanding of informing systems fitness is, therefore, more likely to be advanced by deep observations of the behavior and history of specific system instances than by quantitative analysis of survey data consisting of a subset of attributes and the estimated fitness of many systems.

These conclusions are quite significant in their implications for informing systems research. They come, however, with two potential Achilles heels: 1) if an informing system exists on such a rugged landscape and quantitative methods are unlikely to yield useful results, how is it that the MIS field—which focuses on an overlapping subject area—so often finds statistically significant results in its own quantitative investigations (e.g., studies of the technology acceptance model), and 2) if rugged fitness landscapes are so ubiquitous, how is it that the most widely used fitness function in the social sciences, utility, is not generally perceived to exist on a rugged, multi-peaked landscape.

Understanding the how a rugged landscape influences statistical relationships is the subject of the second paper. How utility and the rugged fitness landscape relate is the subject of the third.

Illusions of Significance in a Rugged Landscape

The second paper (Gill & Sincich, 2008) addresses the question of what happens when quantitative research methods, most particularly multiple linear regressions, are applied to a rugged fitness landscape under the incorrect assumption that the landscape is decomposable. In a practical sense, it is an attempt to understand what happens when we are wrong about the fundamental nature of the landscape we are studying, and therefore use inappropriate tools to study it.
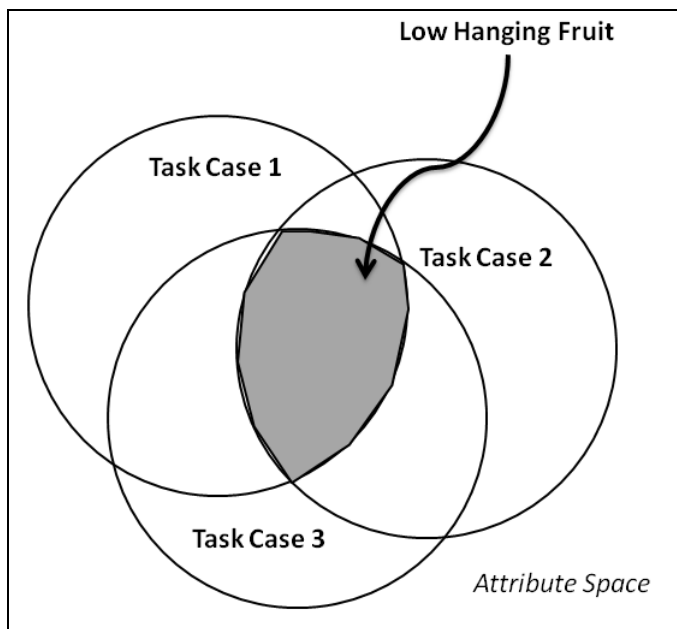

Figure 3: Low Hanging Fruit in Attribute Space (from Gill, 2008e), where the shaded area indicates variables that act consistently across all the included task cases.

Our prediction in performing the analysis was that certain variables would contribute decomposably to fitness and would, therefore, be picked up readily by quantitative tools. These would be the "low hanging fruit" of research, as shown in Figure 3. We expected, on the other hand, that variables that only affected fitness through interactions with other variables would be overlooked. Since MIS research frequently incor-

porates fairly obvious findings into its conclusions (e.g., top management support increases the likelihood of system adoption, intention to use typically precedes system use, to name a just a few of the gems), the explanation that some variables act globally (decomposable) while some act in concert with other variables (based on local interactions) seemed a plausible answer for the statistically significant results so often obtained by the MIS field.

Although the analyses performed were not inconsistent with the proposed explanation, they also introduced an entirely new twist into the problem. The models that were employed assumed that entities on the rugged landscape would attempt to maximize their fitness and would, therefore, tend to migrate to local peaks. This is comparable to mechanisms proposed by Kauffman (1993) in his fitness landscape studies. It is also completely consistent with processes such as utility maximization that are axiomatic in fields such as economics and finance.

What we observed was unexpected. Specifically, as soon as such migrations towards local fitness peaks started to occur, major errors in statistical significance estimates and coefficient estimates begin to surface. In fact, even in entirely random landscapes with 8 independent variables, migration to the local fitness peaks led to significant relationships—often at levels of $p<0.001$—being detected for an average of two variables. Where the underlying process being simulated was a mixture of decomposable and non-decomposable relationships, as it would have been in our Figure 3 model, the impact of migration was even worse. In essence, nearly all statistical tests and estimates began to fail in the direction of showing false significances (although, in a few cases, actual significances were also hidden).

The problem that was identified was not the fault of the statistical algorithms being employed. Rather, it resulted from conducting certain types of analysis (multiple linear regression was used in the examples, but related techniques such as factor analysis, structural equation modeling, and even less sophisticated significance tests would also be affected) on observations that were inconsistent with the model we were fitting them to. In fact, if we recognized the underlying fitness landscape as being rugged, it probably would never have occurred to us to conduct such tests. The example used in the paper is that of a cookbook. Assuming we had a taste tester available, it would be possible to devise a regression for which the dependent variable was taste ranking (fitness) for each recipe while the independent variables were dummies signifying the presence or absence of each ingredient in the recipe. While it would be easy to construct an regression in this manner, we would have little reason to do so because we intuitively understand that each recipe in the cookbook is, itself, likely to be on or close to a fitness peak—the result of years of experimentation by the authors. Even if a particular ingredient seemed to be a contributor to fitness overall (e.g., the coefficient for garlic was positive and had a significant p-value), it would therefore be extremely unlikely that adding it to a recipe where it was absent (e.g., an angel food cake recipe) would enhance the taste fitness of the new recipe.

The concern raised in the conclusion of the paper is that we have no such intuition regarding the non-decomposability of the contributors to fitness of informing systems. Moreover, academic rewards are particularly high for the development of attractive theory, which tends to depend upon the underlying fitness landscape being decomposable, not rugged. As a result, there is considerable incentive for reviewers to ignore the possibility that the landscape they are studying is quite rugged. And, unfortunately, the statistical behavior of their results will not necessarily be treated as evidence against the assumption of decomposability.

A Psychologically Plausible Goal-Based Utility Function

The third paper in the series (Gill, 2008a) focuses on the economic concept of utility. It is designed to act as a bridge between the first two papers—which focus on ruggedness—and the

fourth and fifth papers—which focus on how our mental models change with learning and informing.

Utility is the term economists use to describe the level of satisfaction that we achieve, or expect to achieve (expected utility), as a consequence of decisions. A utility value is therefore a subjective measure of the fitness of our choices, its arguments being the quantities for a specific basket of goods and services that we choose. As such, it has the characteristics of a typical fitness function.

The behavior of the utility function is also fundamental to the mathematical underpinnings of some of the most successful disciplines in the social sciences, including both economics and finance. As such, its existence presents a challenge to the rugged fitness landscape model of the first paper since, for the most part, any potential ruggedness in the utility fitness landscape tends to be ignored by these disciplines. The question then becomes: how can these disciplines be so successful if the underlying function that they base their theories on behaves in a manner inconsistent with their axioms.

In considering this question, the third paper makes a number of points. First, it is not so much that economics and the decision sciences are unaware of utility's capacity to exist on a rugged fitness landscape. Rather, the mathematical intractability of that assumption interferes with theory development and therefore tends to be assumed way. Second, there is a great deal of evidence, developed in psychological studies, that utility exhibits considerable ruggedness. Third, the question of motivation, particularly intrinsic motivation, is nearly absent from the utility literature, in spite of the fact that it is a major source of satisfaction. Fourth, utility changes as we learn—a process almost entirely ignored by economists—and the resulting shape of the utility fitness landscape is likely to change during that process as well.

All four of these points can be illustrated with a typical framing experiment, taken directly from the paper. Consider the following pair of choices (adapted directly from Tversky & Kahneman, 1988, p. 173-174):

First choice problem:

Assume yourself richer by $300 than you are today. You have to choose between:

> A. A sure gain of $100
> B. 50% chance to gain $200 and 50% to gain nothing

Second choice problem:

Assume yourself richer by $500 than you are today. You have to choose between:

> A. A sure loss of $100
> B. 50% chance to lose nothing, 50% to lose $200

The actual results of this experiment were that most (72%) chose A in the first problem, while most (64%) chose B in the second. This presents a bit of a problem for any utility model that assumes consistent choice, since—in fact—options A and B in both choice problems lead to precisely the same outcome: a $400 gain in choice A, a 50-50 distribution between $300 and $500 in the case of choice B. This inconsistency suggests a more rugged utility landscape exists than is generally assumed. Economists have been subjected to reports of similar inconsistencies for over 25 years, and Daniel Kahneman, a psychologist who pioneered research in this area with the late phychologist Amos Tversky, received the 2002 Nobel Prize in Economics for his efforts. Thus, economists cannot claim ignorance of the phenomena; they must be assuming them away.

How this example illustrates the relationship of utility to learning is more subtle, yet more central to the theme of the paper. Consider the following, however. Once you learn about framing ex-

periments, you quickly realize that they invariably follow the same pattern. Two scenarios are presented with two choices for each. Although dissimilar on the surface, the two scenarios nonetheless lead to identical outcome pairs when you dig deeper. It is therefore irrational to choose anything but the pair of choices that lead to the same outcome for both scenarios. In other words, A and A or B and B are rational in our example; combinations of A and B across choices are not. Having learned this about framing problems, would you ever again choose irrationally when presented a framing problem?

The paper presents a model that views utility in a fundamentally different way from the economic perspective. First, using the large and successful goal setting literature developed in psychology and management as evidence, it proposes that utility is better viewed as acting through goals than being a direct function of acquired goods and services. This serves to tie utility to intrinsic motivation, which is fundamental to goal setting models. Second, it proposes that as we learn the nature of the goals change.

The model, which is focused on the utility associated with performing a task (thereby relating it to the task component of an informing system), argues that as we perform a task and learn about it, we move from satisfying very generic goals—hard wired into all of us by evolution—to very task specific goals and, finally, to the single goal of making progress towards completing the task. This progress is shown in Table 3, taken from the paper.

Table 3: Utility sources for task performance (from Gill, 2008a)

Source of Utility	Description
Task progress	Intrinsic utility derives from the process of completing the task. Because the task has become routine, the conscious satisfaction of individual task-related goals is no longer required. In this stage, neoclassical models of the utility of information are likely to be valid.
Task goals	Intrinsic utility derives from satisfying and progressing towards a series of specific goals that are presented as part of the task. These goals may serve to direct activities or act as targeted levels of achievement. In this stage, task-related utility is best predicted by the goal setting models so extensively documented in management and psychology. Inconsistencies in decision-making and biases will tend to lose impact in this region as goals become increasingly well established. As the task is repeated, specific goals and tradeoffs between goals become increasingly automated, ultimately leading to migration towards the task progress stage.
Generic drives and desires	Intrinsic utility derives from satisfying and progressing towards generic drives and desires that are present, to varying degrees, in all individuals. At this level, which is applicable mainly to highly unfamiliar and learning-oriented tasks, utility tends to be highly subject to framing issues and cognitive biases of the sort that economists and decision scientists have identified in experiments— often inconsistent with "rational" models of behavior. As task-specific goals are learned, the utility migrates towards the task goals stage.

Generic goals and, initially, task goals will tend to be numerous and exhibit many interactions with each other. As a consequence, we would expect the utility landscape to be quite rugged during the early stages of learning. In the middle of the task goals phase, however, we start to identify and concentrate on single utility peaks. A single peak focus eliminates many of the challenges presented by ruggedness; therefore, at this point, we expect many of the ambiguities in utility should cease to be significant. By the time the task progress stage has been reached, utility derives directly from observed indications of task progress. At this point, the function behaves much as the economic model would expect it to perform.

Evidence for the proposed model of utility is gathered from many sources, including psychology, cognitive science, economics, evolutionary biology, neuroscience, and computer science. The utility model is therefore also proposed as a case study in how a transdiscipline, such as the informing sciences, can advance thought by bringing together disparate views of the same problem.

Utility learning, however, is only one aspect of the overall learning processes that take place within the mind of the client. To develop a more complete view of learning and its relationship to informing, we need to apply the same principles to broader model of client cognition: the problem space. That is the focus of the fourth paper, which addresses the relationship of structural complexity to informing.

Structural Complexity and Effective Informing

The fourth paper of the series (Gill, 2008e) considers the relationship between structural complexity and informing. It does this using the model of a problem space, which extends the notion of utility developed in the third paper (the *goal space*) to the knowledge representation scheme used for the task (the *state space*) and the mechanisms available for transforming and manipulation that knowledge (the *operator space*). In each case, the model posits four levels of structure, ranging from general purpose (Level 4) to highly compiled, task-specific knowledge (Level 1). For the goal space, this maps well to the utility model—with the Task Goals region of Table 3 being divided into its rugged fitness regions (Level 3) and decomposable regions (Level 2).

In structural complexity, the assumption is that the more specialized the knowledge used in performing the task, the lower the structural complexity. As a consequence, many of the findings reviewed in the paper are drawn from the large literature involving the acquisition of expertise. One nearly universal finding of this literature is that with practice, the contents of our problem spaces move from higher to lower levels—true for the state space (chunking), the operator space (automization), and, as demonstrated in the third paper, for the goal space (learned utility and specific peak focus).

Another way of looking at expertise involves the performance of task cases. A task case represents a specific instance of a task, which may only require a small fraction of the problem space required for the task in general. For example, when a patient walks into a physician's office with a problem, it represents the initiation of a specific task case. The problem space of the physician, on the other hand, is expected to be much larger—since he or she must be able to handle many different task cases as part of the job.

Using this perspective, illustrated in Figure 4 (taken from the paper), experts can be differentiated from non experts based upon the level of knowledge typically applied to task cases. The nature of the goal space, however, also exerts an important influence on the role played by task cases. As illustrated in Figure 5, knowledge specific to the problem space is likely to divide into to two forms: core knowledge that is applicable to all or most task cases and case-specific knowledge, applicable to a single or small number of task cases. Where the goal space exhibits a rugged fitness landscape—which, from the first paper (Gill, 2008b), we anticipate under conditions of high objective complexity—task cases correspond to local fitness peaks and knowledge of these peaks will tend to dominate our problem space contents. Where objective complexity is low (small number of task components, largely decomposable relationships), we expect the relative amount of core knowledge to be much greater.
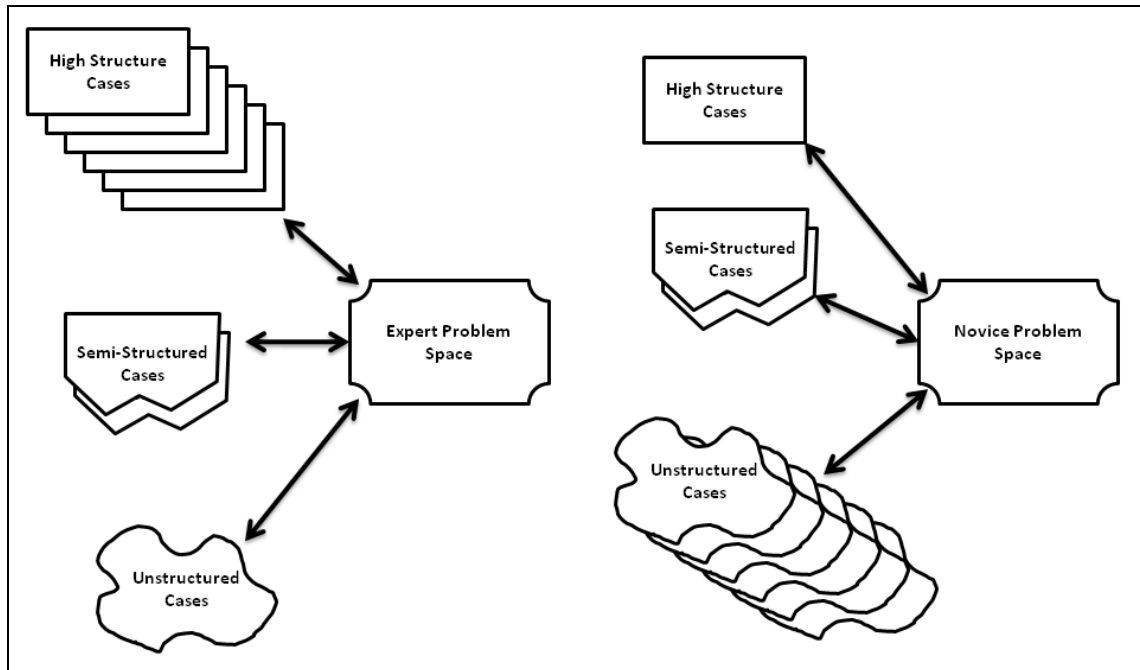
Figure 4: Expert vs. Novice Problem Spaces (from Gill, 2008e). For the expert, the majority of task cases would be expected to involve low problem space levels, with the occasional higher levels required for less familiar cases. For the novice, the situation is reversed.
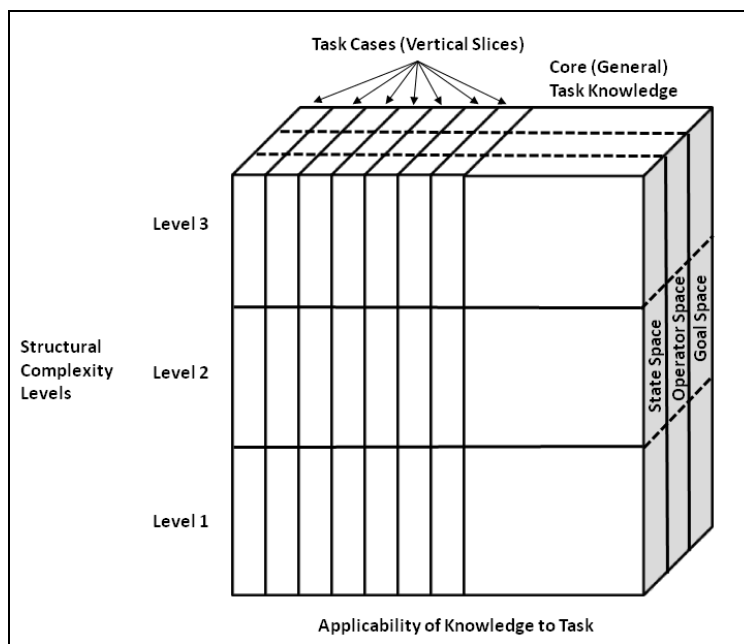

Figure 5: Problem Space Structural Complexity Model (from Gill, 2008e). The relative size of the task cases area, compared to core knowledge, is likely to depend heavily on objective task complexity. Where such complexity is low, core knowledge should represent the vast majority of the problem space. Where objective complexity is high, task case knowledge is likely to predominate.

Once again, we can use the cookbook as an analogy. Ow-Owing to its high objective complexity (many possible ingredients that strongly interact with each other in producing a final product), we'd expect that expertise in the cooking domain would tend to be dominated by knowledge of how to prepare specific recipes, rather than by core knowledge applicable to all cooking activities (e.g., organic chemistry). One natural outcome of this type of problem space would be a tendency of experts to specialize, thereby reducing the number of task cases that must be considered and increasing the relative quantity of core knowledge applicable to the subset of task cases chosen. In the cooking domain, for example, by choosing to become a

pastry chef, the expert increases the amount of knowledge that may be considered core to all task cases. Obviously, medicine and academia are two other examples of areas where such specialization has occurred.

A particularly important concept emphasized by the paper is the distinction between two types of knowledge: *episteme* (theory) and *phronesis* (practical wisdom). This leads to a further distinction that is very important from an informing perspective, that of the academic-expert and the practitioner-expert. The academic-expert is characterized by high levels of episteme, which exists at high problem space levels by virtue of the fact that the knowledge relates to tasks that the expert does not routinely practice. The practitioner-expert, on the other hand, has a knowledge distribution that is dominated by lower, more highly compiled, chunks. This hypothetical distinction is illustrated in Figure 6 (from the paper).
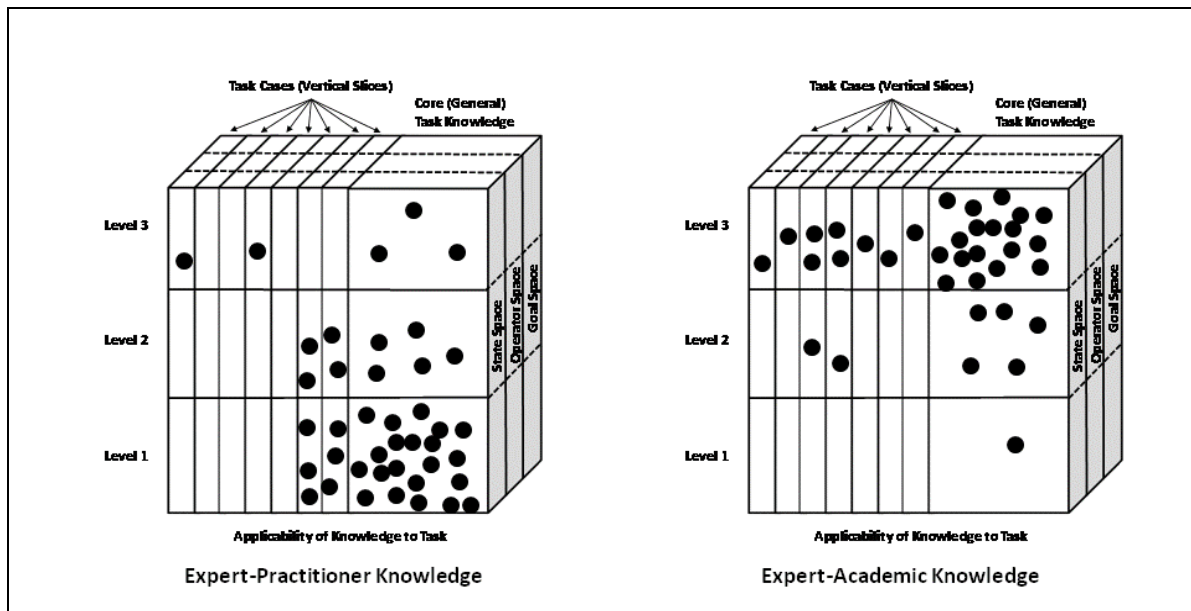

Figure 6: Illustrative Patterns of Practitioner and Academic Knowledge Distributions (from Gill, 2008e). The practitioner mainly utilizes knowledge that has been highly compiled, as a consequence of frequent practice. Such knowledge, however, is likely to be concentrated around familiar task cases. The academic, in contrast, has knowledge that is mainly in symbolic, higher level forms, but that knowledge is likely to support a broader range of task cases.

A central theme of the paper, then, is the types of informing barriers that these distinct knowledge representations present when the academic-expert attempts to communicate with the practitioner-expert. In addition, because the goal space is a function of utility, and utility is closely related to motivation (Gill, 2008a), the paper proposes a three of "laws" related to informing. Specifically:

The Law of Abandoned Expertise: Clients will resist any task-related informing activities that require relinquishing existing expertise in their problem space.

The Law of Limited Visibility: In the absence of concrete negative performance feedback or external pressures, an individual will gradually come to view the entire goal space in terms of the peak that he or she has reached or is presently climbing. The phenomenon will be particularly pronounced in a goal space where multiple peaks exist, but only one provides feedback.

The Law of Low Hanging Fruit: Within a rugged goal space, those problem space attributes that enhance or detract from goal fitness decomposably across nearly all fitness

peaks will tend to obscure equally important contributors to fitness that only act upon certain specific peaks.

The implications of these laws, and the evidence for their existence, are discussed more thoroughly in the paper. The challenges presented by these laws, particularly the last two, grows as the goal space becomes more rugged.

The structural complexity model proposed in the paper essentially acts as a bridge between problem space complexity—which is a function of the specific contents of a problem space—and objective complexity—which largely determines the shape of our goal space. In developing the model, however, two interesting questions emerge:

1. In the study of utility (Gill, 2008a), it was specifically proposed that anomalies in preference that exist at high levels of structural complexity will tend to exert a much lower impact at lower levels of structural complexity. Is the same true for the problem space in general?

2. Given that the obstacles that present themselves when expert-academics and expert-practitioners attempt to inform each other, how is it that any such informing ever takes place?

The first of these questions is addressed in the fifth paper, which proposes a general client-side informing model. The second is addressed in the final paper, which considers the role played by client-to-client informing in the overall informing process.

The Single Client Resonance Model: Beyond Rigor and Relevance

The fifth paper (Gill, 2008d) specifically addresses the role of informing as it relates to the client's mental models. It starts with the premise that a message's impact on a client's mental models cannot be entirely explained by the quality (*rigor*) of the contents of the message and its potential value (*relevance*) to the client. Instead, a third quality, *resonance*—using the term introduced in Gill and Bhattacherjee (2007)—is needed to explain the message's ability to acquire the client's attention and enter the client's problem space without being distorted.

The paper starts with a case study illustrating how rigor and relevance are insufficient to ensure effective informing. In reviewing the informing sciences literature on the subject, it then identifies a model that captures many of the issues relating to resonance: Jamieson and Hyland's (2006) filter model. That model proposes that information on its way to the client must first pass through as series of filters, described as:

1) information biases,
2) cognitive biases,
3) risk biases, and
4) uncertainty biases.

Information biases reflect mechanisms that modify incoming information to align it with existing client preferences, usually applied at an unconscious level. Cognitive biases are employed to simplify decision making so as to keep it within the limits of the client's bounded processing resources. Risk and uncertainty biases derive from client attitudes towards risk and uncertainty, respectively. Because these filters, particularly the first two, can serve to block information altogether, they can act as severe impediments to resonance.

The filter model is a very useful model. Using the problem space model, however, it can be extended in two important ways. First, given the importance of motivation and other visceral factors

in the informing process—as discussed at length in the paper and in the goal-based utility model paper (Gill, 2008a) as well—any discussion of filters would need to incorporate these aspects of the informing task. Second, as was also discussed in the utility model, we would not expect filters to act uniformly on all levels of the problem space. Thus, it is possible to propose that certain filters are more likely to be applicable to certain levels of structural complexity. The resulting *Single Client Resonance Model* (from the paper) is presented in Figure 7.
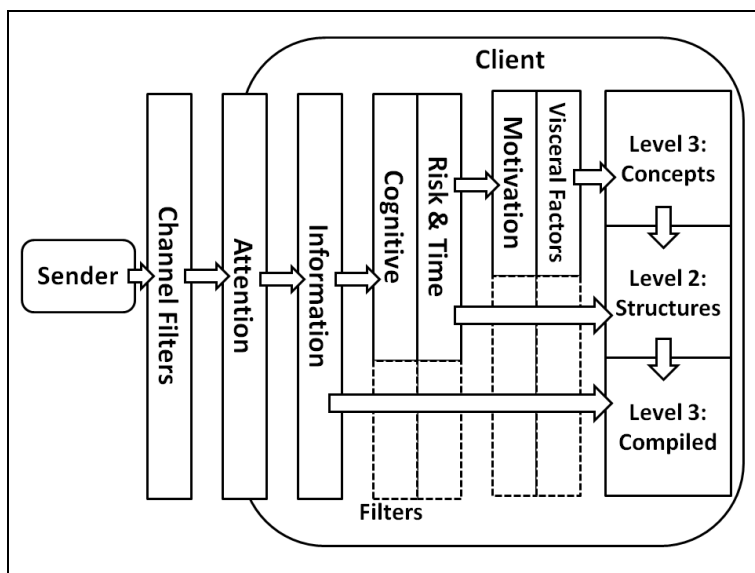

Figure 7: Single Client Resonance Model (from Gill, 2008d). In this model, filters are expected to exert differential effects depending upon the problem space level to which a message is targeted.

The majority of the paper focuses on reviewing findings that support the existence and specify the expected activities of the different types of filters. At the end of the paper, however, the proposed model is contrasted with the SUCCESs model of communication "stickiness" proposed by Heath and Heath (2007). A very strong correspondence is found, further supporting the overall structure of the model.

A significant difference between the SUCCESs model and the proposed Single Client Resonance Model relates to the complexity of the messages involved. Whereas the SUCCESs model is largely built around simple messages (the first S in SUCCESs refers to "simple"), the Single Client Resonance Model is particularly concerned with complex informing situations. As such, it will tend to be most useful in situations where the sender has a relatively high level of knowledge regarding the client's mental models.

In presenting the Single Client Resonance Model, the paper specifically notes that the earlier definition of resonance that had been proposed (Gill & Bhattacherjee, 2007) was actually in two parts: the first involving a message's ability to impact the mental models of an individual client, the second involving its ability to initiate subsequent client-to-client informing. The final paper therefore addresses the second of these aspects of resonance.

Resonance within the Client-to-Client System: Criticality, Cascades, and Tipping Points

A long and distinguished literature on the diffusion of innovation (e.g., Rogers, 2003) finds that few, if any, complex innovations ever make their way into a client community without considerable word-of-mouth communications. We refer to these communications more generally as client-to-client informing. It, therefore, follows that any attempt to understand the process of complex informing without considering client-to-client informing is likely to be of very limited practical value.

Unfortunately, client-to-client informing processes have received scant attention in the informing sciences literature. The overriding purpose of the final paper (Gill, 2008c) is, therefore, to provide

an overview of some these processes and identify some of the common and distinct characteristics of each.

Three general mechanisms of client informing are presented:

1. *Criticality*: This model is based on the concept of a critical system, most commonly used in the context of nuclear engineering. The simplest of the three models, it could be described as client-sender motivated communications, since it is applicable only when one client who possesses the information is strongly motivated to inform other clients about it.
2. *Information Cascade*: Introduced originally in economic theory, this model is normally presented in terms of clients making a choice between two options for which information about prior client adoptions is available. Although often applied to products (e.g., VCR formats, movies), it can also been applied to pure informing situations, such as the enrollment decision made between alternative classes or the choice of a research topic. It can be characterized as client-recipient motivated informing, since it is the potential recipient who actively decides which option to pursue.
3. *Tipping Point*: Building upon assumptions presented in Gladwell's (2000) widely read book <u>The Tipping Point</u>, this model is typical of general diffusion models (e.g., Rogers, 2003) that examine how innovations—including ideas—migrate through communities. It could be characterized as a social-task model, since informing is motivated by both task performance-related criteria and by social criteria.

In each case, a simple simulation was devised to examine the properties of the model. The goals of these simulations were to: 1) identify the parameters that each model requires, 2) identify characteristic behaviors of the model, 3) determine model sensitivity to parameter choices, and 4) determine the level of random volatility that each model produced.

The most significant conclusions of the paper were that all the models exhibited the typical s-curve of diffusion processes, signifying gradual early adoption, followed by rapid diffusion, followed by a tailing off as maximum penetration is reached. They also all exhibited sensitivity to certain key parameters that could dramatically impact the ultimate level of informing penetration that could be expected. From a practical standpoint, they all required values for a sufficiently large number of parameters so that applying the models to real world situations would likely be difficult. We are, therefore, probably limited to qualitative insights into typical system behaviors for the foreseeable future. Finally, all the models have specific domains of applicability with respect to: 1) whether the informing is driven by client-senders (individuals wishing to spread the information within the client
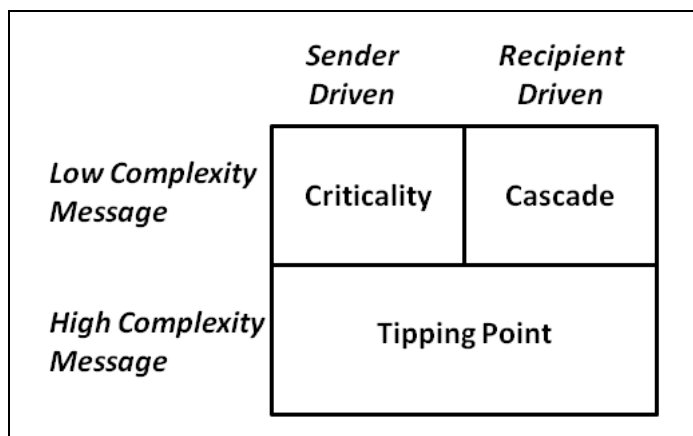

Figure 2: Mapping model domains to who drives informing and message complexity (from Gill, 2008c). In general, the Tipping Point model seems most applicable to situations of high informing complexity. This does not imply that complex decisions cannot be impacted by other models (particular the Information Cascades model). Rather, it suggests that for such decisions—such as which political candidate to vote for—a client may decide to base the decision on what other people are doing, rather than on attempting the complex reasoning necessary to come up with an individual choice.

community) or client-receivers (individuals who want to acquire the information within the client community), and 2) the level of complexity of the information being conveyed. The most appropriate domains are summarized in Figure 8 (from the paper).

The paper concludes by pointing out that the multi-client informing process is likely to be heavily impacted by the rapidly growing body of findings in the area of network theory and that combining network theory with our understanding of complex informing is likely to prove a fruitful research domain in the future.

Conclusions

The variety of reference disciplines for Informing Science is great and remarkable. It draws from research originally conducted in university departments whose research seemingly has nothing to do with one another. It includes, to name a few, psychology, epidemiology, evolutionary biology, and philosophy. While information technology certainly has its important place in informing, the focus of the research presented here is with the clients' needs and the characteristics of their tasks.

The theory introduced here and in the following papers is distinctive in that it extends Informing Science with a focus on the task and its impact on the client. The papers explore task complexity from a variety of epistemologies. It also extends beyond the basic Informing Science model with the introduction of client-to-client communications. The existing Informing Science model does not directly relate to such "informing", yet we need to recognize the important role these processes can play.

Clearly, there are many components to the Informing Science model and much more research that can be and is being done. These papers make a solid and substantial contribution to Informing Science, particularly in exploring the complexity of the client/task interaction. The next big challenge for informing science is to unify these models with the large body of existing research that is built around technology-enabled informing systems.
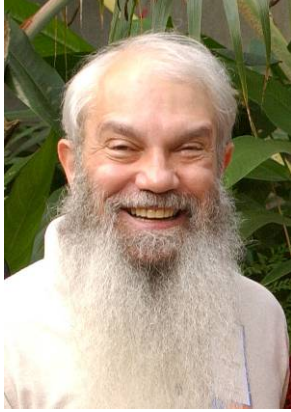
References

Gill, T. G. (2008a). A psychologically plausible goal-based utility function. *Informing Science: the International Journal of an Emerging Transdiscipline*, *11*, 227-252. Retrieved from http://inform.nu/Articles/Vol11/ISJv11p227-252Gill220.pdf

Gill, T. G. (2008b). Reflections on researching the rugged fitness landscape. *Informing Science: the International Journal of an Emerging Transdiscipline*, *11*, 165-196. Retrieved from http://inform.nu/Articles/Vol11/ISJv11p165-196Gill219.pdf

Gill, T. G. (2008c). Resonance within the client-to-client system: Criticality, cascades, and tipping points. *Informing Science: the International Journal of an Emerging Transdiscipline*, *11*, 311-348. Retrieved from http://inform.nu/Articles/Vol11/ISJv11p311-348Gill221.pdf

Gill, T. G. (2008d). The single client resonance model: Beyond rigor and relevance. *Informing Science: the International Journal of an Emerging Transdiscipline*, *11*, 281-310. Retrieved from http://inform.nu/Articles/Vol11/ISJv11p281-310Gill222.pdf

Gill, T. G. (2008e). Structural complexity and effective informing. *Informing Science: the International Journal of an Emerging Transdiscipline*, *11*, 253-279. Retrieved from http://inform.nu/Articles/Vol11/ISJv11p253-279Gill223.pdf

Gill, T. G., & Bhatacherjee, A. (2007). The informing sciences at a crossroads: The role of the client. *Informing Science: the International Journal of an Emerging Transdiscipline*, *10*, 17-39.

Gill, T. G., & Hicks, R. (2006). Task complexity and informing science: A synthesis. *Informing Science: the International Journal of an Emerging Transdiscipline*, *9*, 1-30.

Gill, T. G., & Jones, J. (2008). A tale of three courses: Case studies in course complexity. *Proceedings of the DSI Annual Meeting, Baltimore, MD.* 2091-2096.

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

Gladwell, M. (2000). *The tipping point*. New York: Back Bay Books.

Heath, C., & Heath, D. (2007). *Made to stick.* New York, NY: Random House.

Jamieson, K., & Hyland, P. (2006). Good intuition or fear and uncertainty: The effects of bias on information systems selection decisions. *Informing Science: the International Journal of an Emerging Transdiscipline*, *9*, 49-69.

Kauffman, S. A. (1993). *Origins of order: Self organization and selection in evolution.* Oxford, U.K.: Oxford University Press.

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). New York: Free Press.

Shannon, C. E., & Weaver, W. (1949). *The mathematical theory of communication.* Urbana, Illinois :The University of Illinois Press.

Tversky, A., & Kahneman, D. (1988). Rational choice and the framing of decisions. In D. E. Bell, A. Raiffa, & A. Tversky (Eds.). (1988). *Decision making: Descriptive, normative and prescriptive interactions* (pp. 167-192). Cambridge, UK: Cambridge University Press.

Wilson, T. D. (1981). On user studies and information needs. *Journal of Documentation*, 37(1), 3-15.

Wood, R. (1986). Task complexity: Definition of the construct. *Organizational Behavior and Human Decision Processes, 37*, 60-82.

Biographies

Grandon Gill is an Associate 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 an Editor of the *Journal of IT Education* and an Associate Editor for the *Decision Science Journal of Innovative Education.*

Eli Cohen founded the [Informing Science Institute](#) (ISI), an international organization of over 500 members from over 60 countries. The institute publishes eight journals and several dozen books, all of which are available online to everyone without charge (as well as in paper format). The organization holds one or more international conferences each year. ISI is an organization of colleagues mentoring fellow colleagues. It draws together people who teach, research, and use information technologies to inform clients (regardless of academic discipline) to share their knowledge with others.

Dr Cohen's background is multi-disciplinary. He holds degrees in and has published research in Management Information Systems, Psychology, Statistics, Mathematics, and Education. He has taught in Poland, Slovenia, South Africa, Australia, and the US. In addition, he has conducted seminars in a large number of countries, including Fiji, New Zealand, Australia, Hong Kong, Singapore, Malaysia, Thailand, and Cyprus.