

# Informing on a Rugged Landscape: Homophily versus Expertise

*Grandon Gill*

*University of South Florida, Tampa, Florida, USA*

[grandon@usf.edu](mailto:grandon@usf.edu)

## Abstract

When looking for advice, is it better to seek guidance from an expert or from others more like yourself? The paper introduces a simulation model in which a client seeks advice on how to improve fitness. Its focus is on comparing the outcomes of taking guidance from self-similar peers (homophily) and experts who base their recommendations on statistical significance.

The simulation places a collection of agents on a fitness landscape and models the informing process as the agents search for higher fitness. Four distinct agent types are developed: 1) *randomized hill climbing agents* take no advice and search for higher fitness by testing adjacent states and serve as the control case, 2) *imitative agents* look for guidance from nearby agents (mimicking homophily), 3) *expert-guided agents* are advised based upon a statistically-derived view of the landscape, and 4) *goal-setting agents* establish goals based upon observing other clients and then steadfastly pursue those goals regardless of intervening fitness levels. Of particular interest is how well each type of agent performs as the complexity of the underlying landscape varies.

The simulations described produce strikingly clear outcomes that parallel behaviors observed in real-world settings. In low-complexity environments, expert-guided agents match or outperform all other agent types. As complexity grows, however, expertise becomes fragile to the point where it can become worse than no guidance at all. Imitative agents and goal-setting agents—both of which engage in homophilic behaviors by design—track together until substantial levels of complexity are reached, at which point the goal-setting agents outperform all other agent types.

These results are important in two ways. First, they suggest an underlying rationale for the widely observed homophilic proclivities of human beings—provided we make the assumption that complex environments are routinely encountered. Second, they offer an explanation as to why practitioners frequently seem indifferent to the advice of expert research in fields—such as business and education—where the landscapes being investigated are intrinsically complex.

**Keywords:** homophily, expertise, complexity, NK landscape, multiple regression, informing, simulation.

---

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 [Publisher@InformingScience.org](mailto:Publisher@InformingScience.org) to request redistribution permission.

## Introduction

Clients in an uncertain world can benefit from being informed about the choices available to them. In some settings, such as grade school education, that process is largely dictated by an informing agent, such as a teacher. In others, key choices—such as what major to choose in college—are more likely to rest in the

hands of the client being informed. In these latter settings, an interesting question becomes: “Where should the client go for information regarding the choices that must be made?”

There would seem to be three obvious places that a client—henceforth referred to as an *agent* for the sake of consistency with standard modeling terminology—might go in search of insights. (1) The agent could seek the advice of an expert. Alternatively, the agent could seek the guidance of other agents. Within the “other agents” category, however, there is a further spectrum of choices from (2) agents very similar in characteristics to the original client (homophilic agents) to (3) agents very different from the original (heterophilic agents).

An extensive empirical research stream on the diffusion of innovations strongly supports the generalization that most informing takes place through networks of homophilic agents (Rogers, 2003, p. 307). From an information theory perspective, however, this preference seems odd. Presumably, both experts and heterophilic agents offer much greater potential for knowledge gain since their respective knowledge states differ most from that of the client seeking to be informed. Nevertheless the pattern of preference for homophily is observed over and over in the real world.

One explanation that has been proposed for homophilic preference relates to the underlying complexity of the landscape upon which agents operate (Gill, 2010). In this conceptual scheme (Gill, 2011), an agent’s state might be modeled as a collection of values (e.g., 0s and 1s), with each state being assigned its own ordinal fitness value. Where a landscape is orderly, the fitness of a particular state can be determined with a relatively compact formula. For such landscapes, general expertise with respect to the underlying formula should be of great value. At the other extreme, where a landscape is maximally complex—sometimes referred to as chaotic—knowledge of the fitness of a particular state tells you nothing about any other state on the landscape. On such landscapes, trial and error will likely be as good as any other strategy in finding high fitness states. In between the ordered and the chaotic, the rugged landscape exhibits characteristics that are less extreme. While fitness values on such a landscape do not appear to be random, neither can they be modeled with a simple formula. On such landscapes, it is proposed that homophilic agents may prove particularly successful at searching for high fitness.

The goal of the current paper is to report the results of a simulation intended to test the proposition that the benefits of homophily grow with the underlying ruggedness of the landscape. To simulate underlying landscape, the NK landscape model developed by evolutionary biologist Stuart Kauffman (1993) is used. This model was selected for two reasons: 1) it is widely used in many fields, including business, and 2) it provides a parameter (K) that allows complexity to be tuned. For purposes of comparison, an alternative cluster model of ruggedness is also tested. To simulate different agent types, algorithms defining the respective behaviors were implemented in program code.

The paper begins with a brief review of the literature relating to homophily and diffusion, intended to motivate and clarify the research questions being asked. The components of the model used in the simulation are then described; appendices providing a more detailed description of the landscapes employed and the software developed for the simulation appear at the end of the paper. The results of the simulation and accompanying sensitivity analyses are then presented, followed by a discussion of their significance. The limitations of the research and future directions for research are then identified, after which the key findings are summarized in the conclusions.

## Research on Homophily and Diffusion

Two research areas motivated the simulation model developed in this paper. The first of these involves the significance of homophily in the development of social and informing networks. The second specifically focuses on the diffusion of academic expert knowledge.

## ***Homophily and Social Contagion***

The pronounced tendency of individuals to imitate and cluster with other individuals very similar to themselves, referred to as homophily, has long been recognized. The tendency exhibits itself in two ways: individuals tend to establish linkages with similar individuals (properly referred to as homophily) and people tend to acquire common characteristics of the groups with which they associate (social contagion or social influence), thereby serving to produce increased homogeneity within the group. Because the two processes tend to produce the same end results—clusters of similar entities—distinguishing the two tendencies in empirical data sets can be challenging (Aral, Muchnik & Sundararajan, 2009, p. 21544; La Fond & Neville, 2010, p. 601).

The underlying source and purpose served by homophily has long perplexed researchers. The importance attached to the phenomenon was highlighted in a spring 2010 symposium on the “hardest problems of the social sciences” held at Harvard University. Both the significance of homophily and social contagion—along with the current lack of understanding of their root causes—proved to be a recurring theme. For example, an article describing the symposium (“Solving social sciences’ hard problems,” 2010) reported the following comments from participating researchers:

- [Professor of medical sociology and of medicine at Harvard Medical School] Christakis said he believes that in the years ahead, researchers will be able to say, with more certainty, what evolutionary advantage some attributes bring: why, for example, does emotional contagion exist? Why would it provide a selective advantage if, when you meet someone in a foul mood, it poisons your own mood, too?
- Fowler, a political scientist at the University of California-San Diego, noted that clustering is observable for a number of attributes—be it innocuous phenomena like music tastes, or attributes with more serious implications, such as obesity or alcoholism—but current social-science methods aren’t sufficient to separate homophily (people’s tendency to choose friends like themselves) from influence (friends adopting behaviors from other friends in cause-and-effect fashion).
- [Harvard economist] Claudia Goldin called for further research on the persistent problem of why women are paid less than men are, and how to level the playing field. Her own research has shown that most or all of this bias is unintentional: women self-select into fields that pay less.
- The question of where tastes come from. “If your tastes come from the people around you,” asked Christakis, “where do their tastes come from? Maybe all of a sudden one person wants something for a chance reason, and it just ripples through the network.”

The literature supporting the pervasiveness of homophily in social networks is described in McPherson, Smith-Lovin, and Cook’s (2001) seminal review of over a hundred articles related to the phenomenon. The dimensions across which homophily has been observed include:

1. Race and ethnicity
2. Sex and gender
3. Age
4. Religion
5. Education, occupation and social class
6. Network position (e.g., near the center or at the periphery; McPherson, et al., 2001, pp. 428-429)
7. Behavior
8. Attitudes, abilities, beliefs and aspirations

That we naturally seem to prefer to associate with individuals across so many dimensions of similarity has disquieting implications. This is particularly true for the first six, on the list which tend

to be out of the individual’s control (1-3) or very hard to change (4-6). For example, where high status individuals gravitate towards others of the same race and gender, how can individuals with markedly different characteristics hope to break into that elite group?

**Table 1: Recent research findings relating to homophily and social influence**

Central theme	Research Type	Findings/Implications	Reference
Homophily as determining linkages in a communications network	Simulation	Although homophily has little impact on the shortest path for information flow through a network, it can inhibit diffusion rates when other path algorithms (such as random walk) are used.	Golub & Jackson (2008, 2011)
Homophily inhibits cooperative relationships between dissimilar peers	Simulation	Homophily is an obstacle that needs to be overcome; having groups of similar individuals within an organization may be more effective in promoting cooperation than “token” individual representatives	Bacharach, Bamberger, & Vashdi (2005)
Separate homophilic clusters can form even without fixed attributes that distinguish entities	Simulation	Assuming that linkages can be established where attribute overlap exists and abandoned where it does not, relatively stable cultural clusters can form even where some destabilizing drift is present.	Centola, Gonzalez-Avella, Eguiluz, & San Miguel (2007)
Homophily helps predict the speed of diffusion in a social network	Model & Empirical	Information flows more rapidly across homophilic linkages	Choudhury, et al. (2010)
Even modest attraction of like to like can produce homophilic clusters over time.	Model & Empirical	“The dynamic interplay of choice homophily and induced homophily, compounded over many ‘generations’ of biased selection of similar individuals to structurally proximate positions, can amplify even a modest preference for similar others, via a cumulative advantage–like process, to produce striking patterns of observed homophily.”	Kossinets & Watts (2009)
Different characteristics exert different strengths of homophilic attraction	Empirical	In an online dating context, all features produced a tendency towards seeking sameness, but it was stronger from some characteristics than others	Fiore & Donath, (2005)
Different characteristics moderate the strength of homophilic effect strength	Empirical	Among pre-school and kindergarten students, homophily and social contagion effects appeared to be more pronounced for girls than for boys.	Hanish, et al. (2005)
Homophily may more likely derive from structures that bring similar people together than from a strong preference for similar others as interaction partners	Empirical	At a “mixer”, individuals did not appear to actively seek out others with identical characteristics	Ingram & Morris (2007, p. 579)
Homophily effect is stronger in establishing instrumental ties than expressive ties.	Empirical	“Instrumental ties can include ties for expertise/advice seeking, knowledge exchange, etc., and expressive ties can include those for social activities, friendship, emotional support, etc.”	Yuan, & Gay (2006)

Recent research, summarized in Table 1, has shown that computer-simulated communications patterns can arise that exhibit diffusion properties similar to those observed in homophilic networks, that stable homophilic networks can arise when only modest attractions are present, and that the relationship between entity characteristics and homophily are by no means straightforward. Homophily has also been observed at the organizational level with respect to linkages between firms (Ahuja, Polidoro, and Mitchell, 2009). None of these findings, however, address the potential benefits provided by homophily (that could explain why it is so prevalent).

Our understanding of the pervasiveness of social contagion, particularly at an unconscious level, is more recent. A well-known example of this involves weight gain. Using the meticulous health records maintained as part of the Framingham heart study, Christakis and Fowler (2009, pp. 105-112) found that when one individual in a locally connected network gains weight, there is a strong likelihood that others in the network will follow suit. Evidence of social contagion is found for many emotional states, such as happiness, and for beneficial behaviors, such as getting a flu shot. It has also been observed, however, to occur for emotional states that are negative, such as anxiety, and for behaviors that are decidedly not in the individual's best interest, such as suicide (Christakis & Fowler, 2009). It is the fact that both positive *and negative* influences appear to be imitated that makes the phenomenon so perplexing.

### ***Diffusion, Expertise, and the Research-Practice Gap***

A research-practice gap occurs in an applied field where a body of research that has been developed through academic researcher remains unapplied (and often unnoticed) by practitioners in the same field. The problem appears to be particularly acute in professional fields whose related academic research is dominated by the social sciences, such as business (e.g., Pfeffer, 2007) and education. It has also been a complaint voiced in professional fields whose reference disciplines are more closely related to the physical sciences, such as engineering, and in the life sciences. In medicine, for example, an entire journal—*Implementation Science*—has been established to address the “the scientific study of methods to promote the systematic uptake of clinical research findings and other evidence-based practices into routine practice, and hence to improve the quality and effectiveness of health care” (Implementation Science, 2011).

More generally, understanding the transmission of ideas and innovations from an expert community to a practice community is a central goal of diffusion research. Rogers (2003, p. 307) summarizes the findings of numerous studies as follows: “*Interpersonal diffusion networks are mostly homophilous*” [italics in the original]. This conclusion closely paralleled what Rogers observed in his own pioneering work as a field researcher, discovering that most farmers would not take the advice of an extension agent (i.e., the expert) to apply a particular pesticide. Instead, they would adopt it only after observing its effect on the crops of a neighboring farmer.

Drawing upon the findings of diffusion research, inadequate (and decreasing) homophily has been proposed to be a key factor in explaining why academic research increasingly fails to be incorporated into practice in the MIS field (e.g., Gill & Bhattacharjee, 2009) and for business research more generally (Gill, 2010). Supporting the argument is the trajectory along which business research evolved. Prior to the early 1960s, relatively little academic business research was conducted according to the scientific method. Instead, previous “research” was generally conducted collaboratively with local business organizations and faculty members tended to have limited academic credentials but extensive professional resumes. This environment changed radically over the next two decades, an evolution whose impetus is widely attributed to reports prepared by the *Ford* and *Carnegie Foundations* in the late 1950s that were highly critical of existing business research practices (Khurana, 2007). In response, business schools began to move away from faculty members with extensive experience in practice and towards younger researchers who had

received doctorates in more theory-driven fields such as economics and psychology. In doing so, the degree of similarity between the business academic and the full-time practitioner diminished.

## Ruggedness and the Research-Practice Gap

Although inadequate homophily between researchers and practitioners seems likely to contribute to a research practice gap, it may not be the sole contributor to that gap. One recently proposed explanation augments the “lack of homophily” explanation with an argument based on complexity. It proposes that when clients need to become informed about an environment that is complex, there is a mismatch between certain empirical practices commonly employed as part of the scientific method—most notably hypothesis testing and heavy reliance on statistical significance—and the nature of underlying phenomenon being observed (Gill, 2010). In other words, as complexity grows, experts not only fail the test of homophily, they are likely to be wrong much of the time.

To understand the basis of the argument, the concept of a *fitness landscape* is introduced. Conceptually, such a landscape is a mapping between the characteristics of an entity (including its current state) and a desired outcome, referred to as *fitness*. In evolutionary biology, for example, such landscapes can be used to map an organism’s traits (e.g., genes) to survival (e.g., Kauffman, 1993). In business, they might map the characteristics of an organization or its strategy to some desirable outcome such as potential long term profitability (e.g., Porter & Siggelkow, 2008).

What makes the fitness landscape particularly useful as a conceptual scheme is the related concept of ruggedness, a particular form of complexity. Consider a mapping between  $N$  attributes  $a_1 \dots a_N$  and fitness  $F$ . At one extreme, the relationship may be decomposable, meaning that:

$$F = f_1(a_1) + f_2(a_2) + \dots + f_N(a_N) \text{ or}$$

$$\text{Log}_X(F) = f_1(a_1) + f_2(a_2) + \dots + f_N(a_N) \text{ where } F = X^{f_1(a_1)} * X^{f_2(a_2)} * \dots * X^{f_N(a_N)}$$

Where perfect decomposability is present, an agent can maximize fitness by finding a value of  $a_i$  that maximizes each of the  $f_i(a_i)$  expressions individually. If we limit ourselves to representing each element  $a_i$  of the agent’s state  $a_1 \dots a_N$  with values of 0 or 1, such a landscape will be constrained to a single peak value for fitness. This decomposable landscape is wonderfully consistent with a research methodology that supports hypothesis testing, since the research problem of improving the value of  $F$  can be broken down meaningfully into  $N$  separate hypotheses that describe each  $f_i(a_i)$  function’s incremental contribution to fitness.

At the other extreme, fitness might instead be entirely non-decomposable, which is to say that the expression:

$$F = f(a_1, a_2, \dots, a_N)$$

cannot be meaningfully broken down further. We refer to such a relationship as maximally *rugged* or *chaotic*—a term used by Kauffman (1993)—for two reasons. First, the landscape is likely to have many local fitness peaks, making it more like a cluster of sharp peaks than a single gradually sloping hill. Second, such relationships often lead to situations where small changes in one variable can lead to large changes in fitness.

An example of a rugged relationship is the mapping between the ingredients of a recipe (attributes) and how good it tastes (fitness). If decomposable, you would be able to make a statement about a cake along the lines of: “Flour is 30% responsible for how good it tastes, sugar is 20% responsible, eggs are 10% responsible...” and so forth. What makes such a statement ludicrous is our awareness that it is the combination of ingredients, rather than the individual contributions of the ingredients themselves, that normally lead to a “fit” baking outcome.

The relevance of ruggedness to the academic-practice gap becomes apparent when we consider the fit between the research approach and the landscape being studied. Difficulties arise, in particular, where:

1. The environment being studied is rugged
2. The research methods being employed assume decomposability.

Of greatest concern are empirical research approaches that map individual attributes to outcomes, such as significance-based hypothesis testing and multivariate statistical analysis. Simply stated, if the underlying landscape being studied by a researcher happens to be rugged and such techniques are employed, the results will not only be hard to diffuse because of homophily issues, they will also be hard to diffuse because they are misleading. Substantial ruggedness necessarily means that few, if any, simple hypothesis will hold true across an entire landscape.

The development of the research-practice gap model just described hinged upon building arguments that business landscapes will tend to evolve towards ruggedness (Gill, 2010, p. 110) and that a great deal of business research explicitly or implicitly assumes decomposability (Gill, 2010, p. 325). As a consequence of this mismatch, *business practitioners might be justified in their suspicions of academic research*. There is, in fact, considerable empirical support for precisely such suspicions. In *The Black Swan*, Taleb (2007)—another participant in the previously mentioned Harvard symposium—describes the numerous failures that have resulted from following the advice of economists and financial experts, gleefully referred to as “empty suits”. He also cites Shanteau’s (1992, p. 259) landmark study of task categories where experts routinely succeed and fail, a study that specifically lists lack of decomposability as a source of failure.

If experts tend to fail when the underlying landscape is rugged, agents confronting such landscapes might consider seeking advice from other sources—such as nearby agents. At least one study—a simulation guided by patent data (Sorenson, Rivkin, & Fleming, 2006, p. 994)—offers support for the potential value of advice from nearby agents, finding “robust support for the proposition that socially proximate actors have the greatest advantage over distant actors for knowledge of moderate complexity.” Since the goal of that study was to discover a *new* “high fitness” state using existing patents as a template to guide search, the advantage of proximity did not extend to maximally rugged landscapes; for such landscapes, knowledge of the fitness of a particular state provides no information about the fitness of adjacent states. If, on the other hand, we allow agents to occupy the *same* state as observed high fitness actors, we would expect the advantage derived from observing nearby agents might continue to climb with ruggedness.

If observing the fitness of nearby neighbors can be shown to improve an agent’s fitness, we would have rational basis for both homophily and social contagion, since both encourage agents to cluster with self-similar neighbors whose fitness can then be observed. Summarizing this in the form of a research question:

*As the ruggedness of an environment increases, does mimicking the behavior of similar agents become superior to the generalized advice of experts for the purpose of identifying appropriate paths to higher fitness?*

## Research Design

The approach taken in answering the paper’s research question is simulation. To establish the setting, we imagine that many agents are placed on a fitness landscape with binary attributes  $a_1, \dots, a_N$  of the form:

$$F = f(a_1, a_2, \dots, a_N)$$

Each agent's goal is to increase fitness. The obstacle: each agent initially knows nothing about the form of the fitness function, which may range from completely decomposable to maximally rugged. In order to increase fitness, each agent can test adjacent states on the landscape (i.e., states where only one value  $a_i$  differs from the initial state). In addition, some types of agents may have access to expert advice, while others may observe the fitness of neighboring agents; in both these cases, the advice can be used in prioritizing what state to test next. The process ends when every agent on the landscape has reached a fitness peak (i.e., a state where, for every value  $i$ , moving from  $a_i$  to an adjacent value  $a_i^*$  leads to lower fitness). When the simulation is complete, a variety of measures—e.g., the number of steps required for all agents to achieve peaks and the average and cumulative fitness at the end of the run—are gathered to assess the success of the particular strategy being tested.

The design of the simulation required three separate elements:

1. A mechanism for constructing a fitness landscape with tunable ruggedness
2. A plausible mechanism for generating “expert” advice
3. Agents able to respond to homophilic and expert advice

Each of these elements is now discussed. The emphasis here is on the rationale for the specific design selected. In the “Limitations and Directions for Future Research” section, near the end of the paper, inherent weaknesses in the model and alternative approaches that could be explored in the future are presented.

### ***Fitness Landscapes***

To simulate ruggedness of the underlying environment, fitness landscapes were created that mapped a series of binary attributes to a real number between 0.0 and 1.0, representing the ordinal level of fitness for each possible combination of attributes. Two types of fitness landscape were used for the simulation, each having the same two parameters:

- $N$ : The number of attributes
- $K$ : the number of interactions between elements

The first landscape family was a standard NK landscape model, developed by evolutionary biologist Stuart Kauffman (1993). In this model,  $N$  separate functions contributing to fitness are defined, one for each attribute that impacts fitness. In addition to the specified attribute, each function depends upon  $K$  additional attributes drawn at random from the other attributes. The net effect of this process, more fully described in Appendix A, is that as  $K$  increases so does the ruggedness of the landscape. This property makes the ruggedness of the landscape tunable. The two extremes are as follows:

- $N,0$ : No interactions between variables, meaning the landscape is completely decomposable and the effects of  $a_1 \dots a_N$  are entirely independent of each other.
- $N,N-1$ : The maximally rugged case, where all attributes interact such that the effect of changing  $a_i$  cannot be determined without knowing the values of all the remaining attributes. This landscape is simulated by assigning a random value to each of the  $2^N$  possible combinations of 0,1 attributes available.

NK landscapes have been widely used for the purpose of simulating the underlying complexity of an environment. In business research, for example, such landscapes are often used to simulate or describe competitive dynamics (e.g., Levinthal, 1997; Levinthal & Warglein, 1999; Porter & Siggelkow, 2008). They have even been used in simulating the imitation and non-imitation of complex strategies (Rivkin, 2000, 2001), where it was concluded that a complex strategy combined with limited visibility could serve as a barrier to imitation.



Although NK landscapes proved to be an obvious choice for simulating a fitness landscape, their properties do not necessarily model all, or even most, real world environments. To determine if the structure of the landscape artifact is a critical driver of simulation results, a second type of landscape—referred to here as an *NK cluster landscape* or simply *cluster landscape* was tested. The cluster landscape is identical to the NK landscape at the two extremes:  $N,0$  and  $N,N-1$ . For intermediate values of  $K$ , however, it exhibits a much simpler and more transparent structure than the NK landscape, as discussed in Appendix B. As a consequence, it is easier to understand its behaviors, such as the number of peaks for a particular NK combination and the decomposability of certain variables. Most important, it is sufficiently different from an NK landscape in its intermediate  $K$  value structures that behaviors common across the two landscapes are unlikely to be pure artifacts of the particular landscape structure being tested. That, in turn, increases the likelihood that such common behaviors will generally be observed where ruggedness is present.

## **Generating Expert Advice**

Creating a plausible “expert” to guide agents involved a substantially less straightforward design decision than the generation of complex landscapes. The approach chosen began with the following design constraints:

1. *The same expert advice needed to be broadcast to all agents.* This was based upon a particular interest in understanding the research-practice gap, where the principal channel is typically the published paper available to all interested parties.
2. *Advice needed to be derived from observations of agents.* In the model employed, fitness is hidden until an agent occupies or tests a particular attribute combination. The “expert” differs from the individual agent in the ability to observe all agents and perform empirical analysis on the group as a whole.
3. *Advice was to be in the form of recommendations of values for individual attributes.* This would be a natural fit with a hypothesis testing empirical research methodology.

With these constraints in mind, it was fairly natural to choose multiple linear regression of the fitness values of all agents on the landscape to develop estimates for the fitness contribution of each individual attribute. Significance testing (at the 0.05 level) was then employed to decide when the “expert” would actively recommend a particular value for a particular attribute. Such significance testing is widely used in empirical research, a fact that has been criticized quite stridently in some circles (e.g., Ziliak & McCloskey, 2007). As a practical matter, however, in many social science fields it is rare to see empirical research published without significance tests.

What the “expert agent” is not privy to is the underlying structure of the landscape as a whole. The rationale for this design decision was prior research indicating that such structure would be very difficult to determine from empirical observations (Gill & Sincich, 2008), a finding considered at greater length in the discussion section of this paper.

## **Agent Behaviors**

Each simulation run was to be populated by a homogenous collection of agents placed on the landscape at random locations (i.e., attribute states). Four different types of agents were developed: *randomized hill climbing agents*, *imitative agents*, *expert-guided agents* and *goal-setting agents*. The local behaviors and design rationale for each type of agent are now described.

### **Randomized Hill-Climbing Agents**

Randomized hill climbing agents, henceforth referred to as *random agents* or *control agents*, were established as a control group. Each agent takes no advice and searches for higher fitness using a very simple algorithm on each move:

## Homophily vs. Expertise

- A. If all adjacent states have been tested, do nothing. This means the agent is on a peak.
- B. Randomly choose a state from the list of adjacent states that have not been tested. Determine its fitness.
- C. If the fitness of the test state is higher, move to that state.
- D. Otherwise, eliminate the test state from the list of available adjacent states.

This particular type of agent is guaranteed to reach a local peak eventually, but there is no guarantee that the peak will be high or that it will reach it quickly. It therefore provides a useful basis for comparison with other agent types.

## Imitative Agents

The imitative agent represents the first “homophilic” agent. The agent acts exactly like the random agent *except* that it takes advice from the closest agent of higher fitness. To limit the nearby agents considered, a *visibility* parameter is introduced. It can vary from 1 differing attribute (only agents in adjacent states are considered) to N differing attributes (all agents can potentially be tested for advice). The algorithm employed for each move is as follows:

- A. If all adjacent states have been tested, do nothing. This means the agent is on a peak.
- B. Identify the closest agent of higher fitness that is: i) within the visibility range, and ii) has at least one differing attribute that can be used to construct an untested adjacent state.
- C. If an agent to imitate is available, randomly chose an adjacent state from the set of available states produced by changing differing attributes one at a time. Determine its fitness.
- D. If no agent to imitate is available within the specified visibility range, randomly choose a state from the list of adjacent states that have not been tested. Determine its fitness.
- E. If the fitness of the test state is higher, move to that state.
- F. Otherwise, eliminate the test state from the list of available adjacent states.

The particular type of agent is also guaranteed to reach a local peak eventually. The hope is that relying on adjacent agents will allow it to reach the peak without as many unproductive test steps.

## Expert-Guided Agents

The expert guided agent behaves identically to the imitating agent except it gets its preferred test states from the landscape regression results rather than from adjacent agents. The algorithm employed for each move is as follows:

- A. If all adjacent states have been tested, do nothing. This means the agent is on a peak.
- B. Determine all significant attributes from the regression of attributes against agent fitness that would lead to an adjacent state that has not yet been tested.
- C. If one or more adjacent states are available based on the regression, randomly chose one of those states. Determine its fitness.
- D. If no regression attributes lead to untested adjacent states, randomly choose a state from the list of adjacent states that have not been tested. Determine its fitness.
- E. If the fitness of the test state is higher, move to that state.
- F. Otherwise, eliminate the test state from the list of available adjacent states.

The particular type of agent is also guaranteed to reach a local peak eventually. The hope is that the significance results—acquired by regressing the attributes of all agents on the landscape (independent variables) against their respective fitness (dependent variable)—will improve the efficiency of the search.

## Goal-Setting Agents

The first three types of agents share two important characteristics: i) they will never move to a state that is of lower fitness than the one they currently occupy, and ii) they will never revisit a state that they have previously occupied. The fourth class of agent was designed to be less constrained. Like the imitative agent, it relies on other agents to determine the fitness of states. Unlike the imitative agent, however, once it chooses a state to imitate it will continue to move towards that state even if required to pass through states of lower fitness or previously occupied states. In addition, although it uses the same visibility parameter as the imitative agent, it chooses the maximum fitness agent within its visibility range regardless of distance instead of choosing the closest. As a consequence, it is able to move off a local peak if a higher fitness goal becomes available. The algorithm employed for each move is as follows:

- A. Identify all agents within the visibility range and choose the highest fitness agent.
- B. If an agent is found in the visible range whose fitness is higher than the current fitness and is higher than the fitness of the current goal (if one exists), make that agent's state the new current goal.
- C. If a current goal exists, randomly chose an attribute from the current goal that is different from that of the current state and move to the associated adjacent state.
- D. If no current goal is established, act like a random agent for the current move.

The rationale behind the goal-setting agent is twofold. First, it provides insights into the consequences of requiring a strict hill climbing algorithm for the first three agent types. Second, and more importantly, goal-setting is widely viewed as one of the best tools available for increasing individual effectiveness (Locke, 2004). Furthermore, making progress towards a goal can be an important source of utility (Gill, 2008). Utility, in turn, provides the individual with an estimate of fitness (Gill, 2010). Given this relationship, it would be realistic to expect that an agent could rapidly transit through states of low fitness without experiencing diminished utility. For example, many people enjoy the mountain climbing experience despite the minute-to-minute discomfort associated with the actual climb; they know that they are moving towards the summit (the goal). These same individuals might never even consider taking the stairs up to their high rise office despite the fact that the accompanying state of exhaustion might be similar for the two activities.

## Model Parameters

Another factor influencing model design was the need to limit the number of parameters to be tested. A common problem with simulation research is that values for many parameters need to be established in order to run the simulation. Frequently, sensible choices for these parameters are not readily available, leading to the need for constant guessing or interminable sensitivity analysis. The model presented is parsimonious in its parameters, limiting itself to five of interest for each of the two landscape structures:

1. N: the number of attributes
2. K: the level of interaction (ruggedness)
3. Visibility: Ranging from 1 to N
4. Number of agents on the landscape
5. Type of agent

## Model Dependent Variables

In assessing the effectiveness of the different agent types, several different outcome measures were considered. These were collected after all agents had reached peaks for a particular simulation trial:

1. *Average fitness across the landscape*: Since, by design, a trial would not end until every agent was on a peak, this measure was an indicator of how effective the informing strategy was at ensuring that agents reached high fitness peaks, as opposed to lower peaks.
2. *Average cumulative fitness across the landscape*: This variable was intended to measure an agent's average fitness from the starting position until the simulation trial terminated. It was included so that the cost of allowing the goal-setting agent to move to lower fitness states in pursuit of a goal could be assessed against the other three strategies, which involved pure hill-climbing.
3. *Average percent of peaks occupied*: This was a measure of spread. A low value would suggest that a lot of high fitness peaks could have been missed during the search process.
4. *Average percentage of peaks above median peak fitness*: The measure—which was expected to be correlated with the average fitness measure—was also intended to reflect if a particular strategy was good at preferentially selecting high peaks.
5. *Number of steps required to complete the run*: A measure of a strategy's effectiveness in quickly locating a peak.

With the design in mind, we now turn to a brief description of the research method.

## Research Method

The research method began with the development of a software application capable of implementing the model. The details of that application are presented in Appendix C.

After testing the simulation application, a “base case” run was established, performed for both NK landscapes and cluster landscapes. This run consisted of the following parameter settings:

1. N: 10
2. K: 6
3. Visibility: 2
4. Number of agents on the landscape: 50

The N setting was consistent with other social science NK simulations (e.g., Levinthal, 1997) and was reflective of the fact that N itself tends to be less important than its value relative to K.

Moreover, earlier research using other values, such as 8, 12 and 16 (e.g., Gill & Sincich, 2008) had not observed significant sensitivity. The K setting was chosen so that it would be in the transition range between orderly and chaotic behavior. This transition tends to occur around the point where  $K=N/2$  for cluster landscapes (Kauffman, 1995, p. 57) and Kauffman argues that complex systems naturally tend to gravitate to that transition range. The visibility was set to the lowest interesting value (at  $V=1$ , goal-setting agent behavior is identical to imitative agent behavior, so  $V=2$  provided a better base case). Finally, Population=50 was chosen to ensure that the landscape (which has  $2^{10}$  or 1024 states) did not begin overly populated.

Sensitivity analysis was then performed by beginning with the base case then making the following adjustments:

- Varying K stepwise from 0 to N-1 (its full range)
- Varying visibility stepwise from 1 to N (its full range)
- Testing populations of 25 and 100. Higher population tests appeared to be unnecessary based upon the results observed.

For each set of parameters, 100 separate trials were conducted for each agent type. For each trial, the simulation set up the same landscape and initial agent positions for each of the four agent types. After each trial, results were recorded for each of the following outcome measures:

1. *Average Fitness*: Average fitness across the landscape at the end of the run. Range: 0.0 to 1.0.
2. *Cumulative Average Fitness*: Fitness for each entity accumulated from the start to the end of the run divided by the number of steps, averaged across the landscape. Range: 0.0 to 1.0.
3. *Percent Peaks Occupied*: Number of peaks occupied by at least one entity divided by the total number of peaks in the landscape. Range: 0.0 to 1.0.
4. *Percent Peaks above Median*: Percentage of entities occupying peaks with a fitness value greater than median peak fitness.
5. *Number of Steps*: Number of steps required to reach a stable state, ending the run.

## Results

The results for the base case are presented in Table 2. Comparison of NK and cluster results suggests that the landscape structure did not exert a material qualitative impact on the base case results. Percent peaks above the median correlated with average fitness and percent peaks occupied was negatively correlated, both of which were expected results (since the only way to increase average fitness is to cluster on higher peaks). For the three remaining fitness indicators, observed differences in the sample means for average fitness, steps and cumulative fitness were statistically significant (or nearly so) between agent types at thresholds below what appeared to be material. This suggested that 100 trials were more than sufficient to pick up material differences in agent performance. For these three key indicators, the differences that were statistically significant *and* material existed only between goal and random/imitator/expert (for average fitness), and between goal/imitator and random/expert (for number of steps and for cumulative fitness). In the latter case, the homophilic agents both outperformed the expert and control agents by a wide margin.

**Table 2: Results for 100 runs of the base case for each agent type  
(N=10, K=6, Visibility=2 and Number of Agents=50)**

	Value (SE)	Random	Imitator	Expert	Goal
NK Fitness Landscape	Average Fitness	.8719 (.0377)	.8761 (.0402)	.8869 (.0550)	.9100 (.0445)
	Percent Peaks Occupied	.8059 (.1110)	.5262 (.1026)	.6633 (.1301)	.3122 (.0880)
	Percent Peaks Above Median	.7156 (.0967)	.7488 (.1422)	.8029 (.0960)	.8619 (.1457)
	Number of Steps	30.8300 (6.0845)	16.0600 (4.7049)	33.3300 (19.5198)	14.0000 (4.4362)
	Cumulative Average Fitness	.5327 (.0699)	.7050 (.1269)	.5773 (.1000)	.7715 (.1091)
Cluster Landscape	Average Fitness	.9413 (.0311)	.9450 (.0304)	.9491 (.0253)	.9783 (.0239)
	Percent Peaks Occupied	.9364 (.0767)	.7619 (.1565)	.8940 (.0934)	.3152 (.1487)
	Percent Peaks Above Median	.5780 (.1106)	.6092 (.1433)	.6148 (.1421)	.9021 (.1344)
	Number of Steps	25.5900 (4.5367)	13.8200 (3.4968)	20.0800 (4.7617)	12.2300 (3.3789)
	Cumulative Average Fitness	.6479 (.0762)	.8335 (.0930)	.7232 (.1060)	.8723 (.0670)

## Sensitivity Analysis

The sensitivity analysis results for *average fitness*, *number of steps* and *cumulative average fitness* are presented graphically in the figures that follow. (Because *percent peaks occupied* and *percent peaks above median* tracked average fitness as previously described, they have been omitted.) Figures 1a and 1b show the results of varying K across its full range. Figures 2a and 2b show the results of varying visibility across its full range. Figures 3a and 3b show the results of halving and doubling the number of agents on each landscape. In each set of figures, the (a) case provides results for the NK landscape and the (b) case provides results for the cluster landscape. Looking at the qualitative behavior of the result patterns, the most evident outcomes are as follows:

- *Ruggedness (see Figures 1a and 1b)*: A number of interesting phenomena are observed as ruggedness increases across both landscapes:
  - For low values of K ( $\leq 3$ ), expert, imitator and goal-setting agents all outperform the control agent with respect to the number of steps by a similar and significant margin. As K increases, expert agent performance deteriorates until it reaches a level that is roughly equivalent to no advice at all (i.e., the control group) by the time K=6 (NK) and K=8 (cluster). Imitator agent performance deteriorates somewhat less, but does diverge from the best performer (goal-setting).
  - Average fitness attained is essentially the same for all groups but the goal-setting agents at all levels of ruggedness for the NK landscape. For the cluster landscape, the goal-setting agent outperforms the other agents after K=0. The margin first increases, then decreases. A similar pattern exists for the NK landscape (it appears less pronounced in the graph because the Y-axis encompasses a larger range).
  - The cumulative fitness of the expert agent deteriorates relative to the other agents as ruggedness increases for both landscapes.
- *Visibility (see Figures 2a and 2b)*: Control and expert performance are unaffected by visibility as a matter of design. Imitator and goal-setting agents track together on number of steps and cumulative fitness, both doing substantially better than the control and expert agents. Interestingly, the lowest visibility level (only adjacent states can be viewed), all four agents are similar with respect to the average fitness ultimately attained. As visibility increases to 3 or 4, however, goal-setting performance diverges dramatically from the rest of the pack. Beyond visibility of 3 or 4, however, few additional gains appear to be made, meaning that being able to look at very distant agents does not seem to confer a material advantage in the end-state reached for either landscape type. For number of steps and cumulative fitness, incremental benefits from additional visibility drop even faster.
- *Number of Agents (see Figure 3)*: As the number of agents on the landscape increases, both types of homophilic agents make modest gains in the number of steps and cumulative fitness. Expert agents are relatively flat with respect to average and cumulative fitness, increasing slightly with respect to the number of steps. This can be explained in terms of the increased likelihood of “unlucky” agents that happen to take a long time to reach and stabilize on peaks owing to poor random guessing. To the extent a trend exists, it is probably not material.

In summary, the two homophilic agents—the imitator and the goal-setting agent—generally outperform the other two types of agents as soon as significant complexity is present in the environment. This is generally true for speed of achieving a peak (steps), for the average fitness of peaks attained, and for the fitness accumulated over the course of the simulation. On the surface, these results seem to provide a compelling case for homophily in informing—along with a rather surprising indictment of “expertise”. The extent to which these results offer useful insights, as opposed to being artificial outcomes of the simulation artifact, is now discussed.




Figure 1a: Sensitivity analysis for ruggedness (K varying from 0 to 9) for NK landscape structure

Homophily vs. Expertise

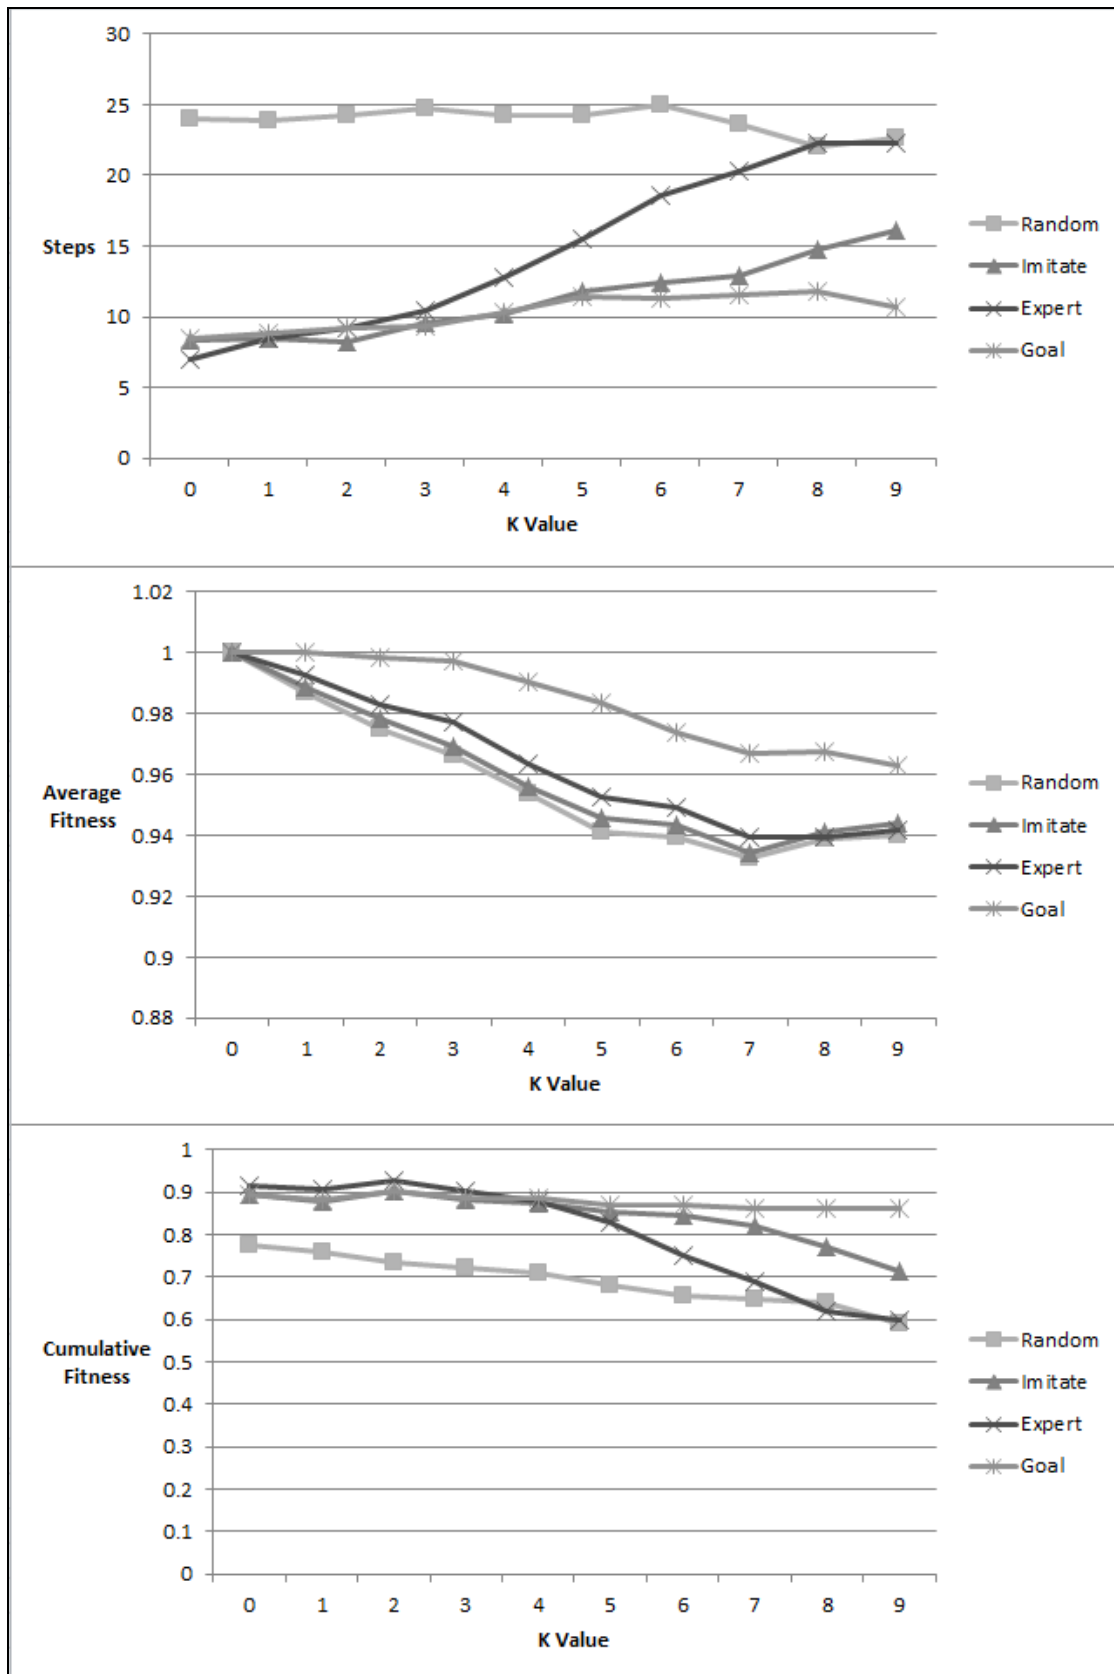


Figure 1b: Sensitivity analysis for ruggedness (K varying from 0 to 9) for cluster landscape structure






Figure 2a: Sensitivity analysis for visibility (Visibility varying from 1 to 10) for NK landscape

# Homophily vs. Expertise

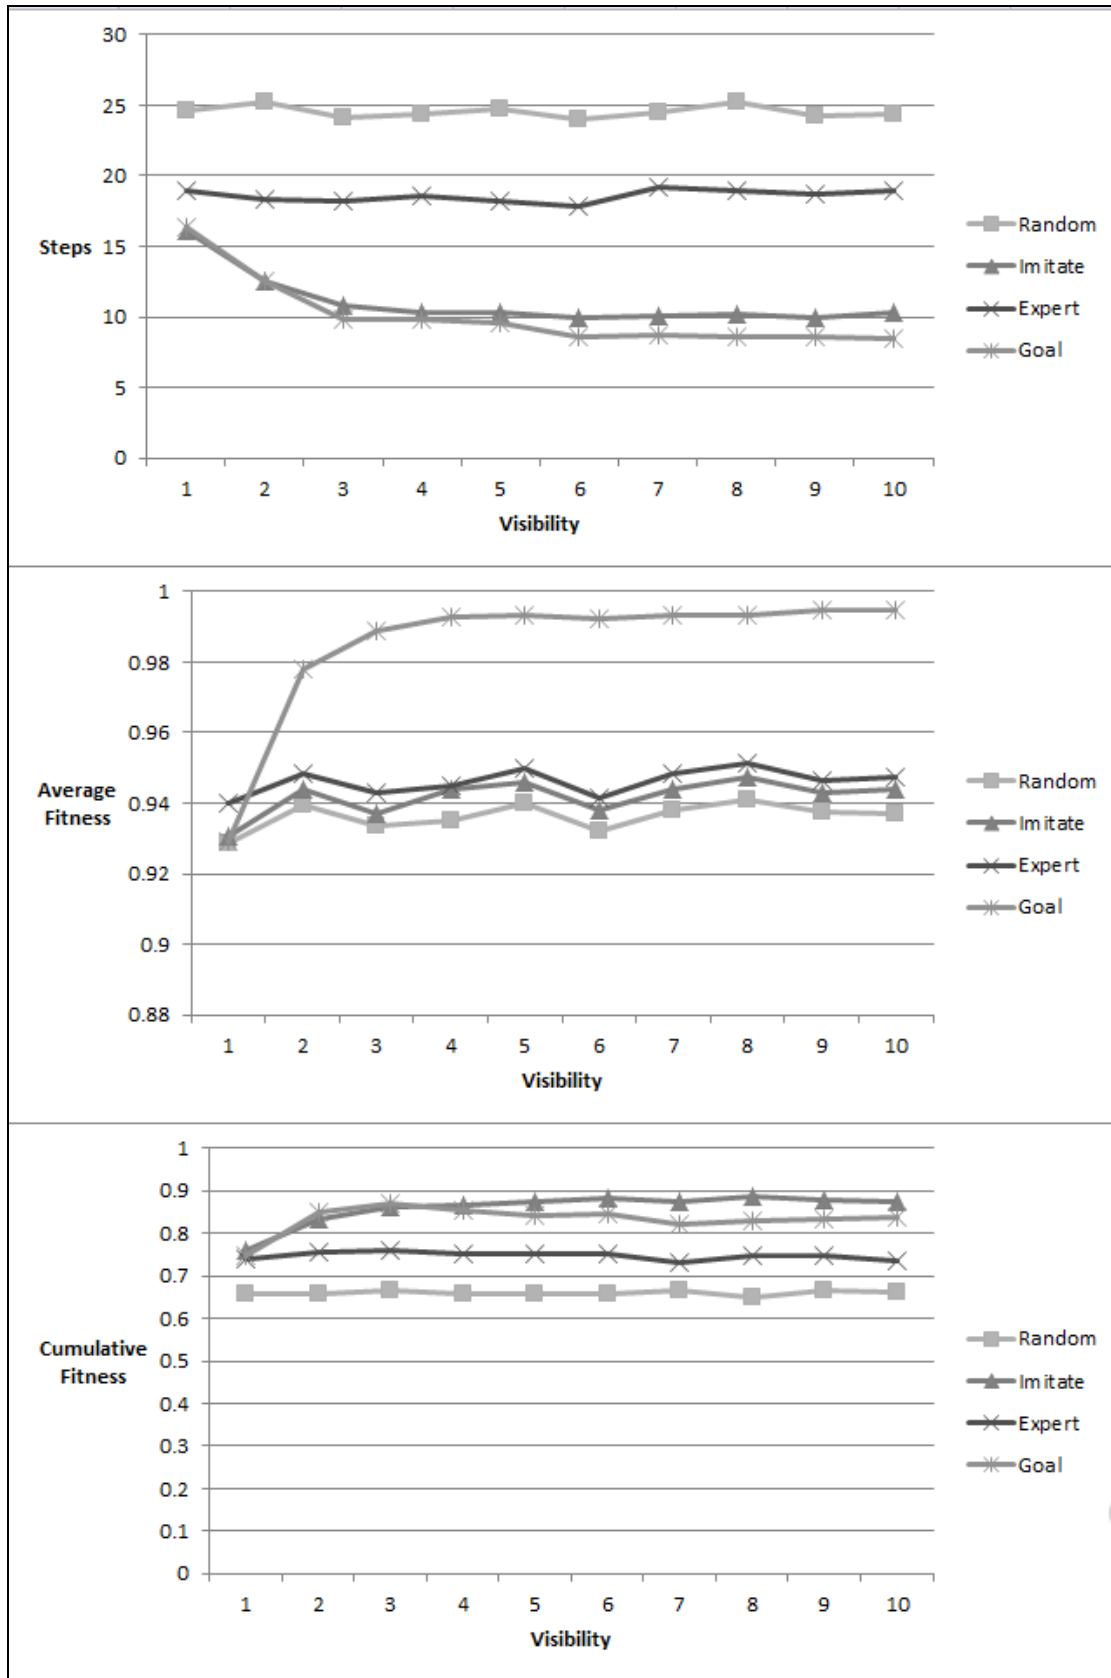


Figure 2b: Sensitivity analysis for visibility (Visibility varying from 1 to 10) for cluster landscape




Figure 3a: Sensitivity analysis for number of agents (Values of 25, 50 and 100) for NK landscape

Homophily vs. Expertise

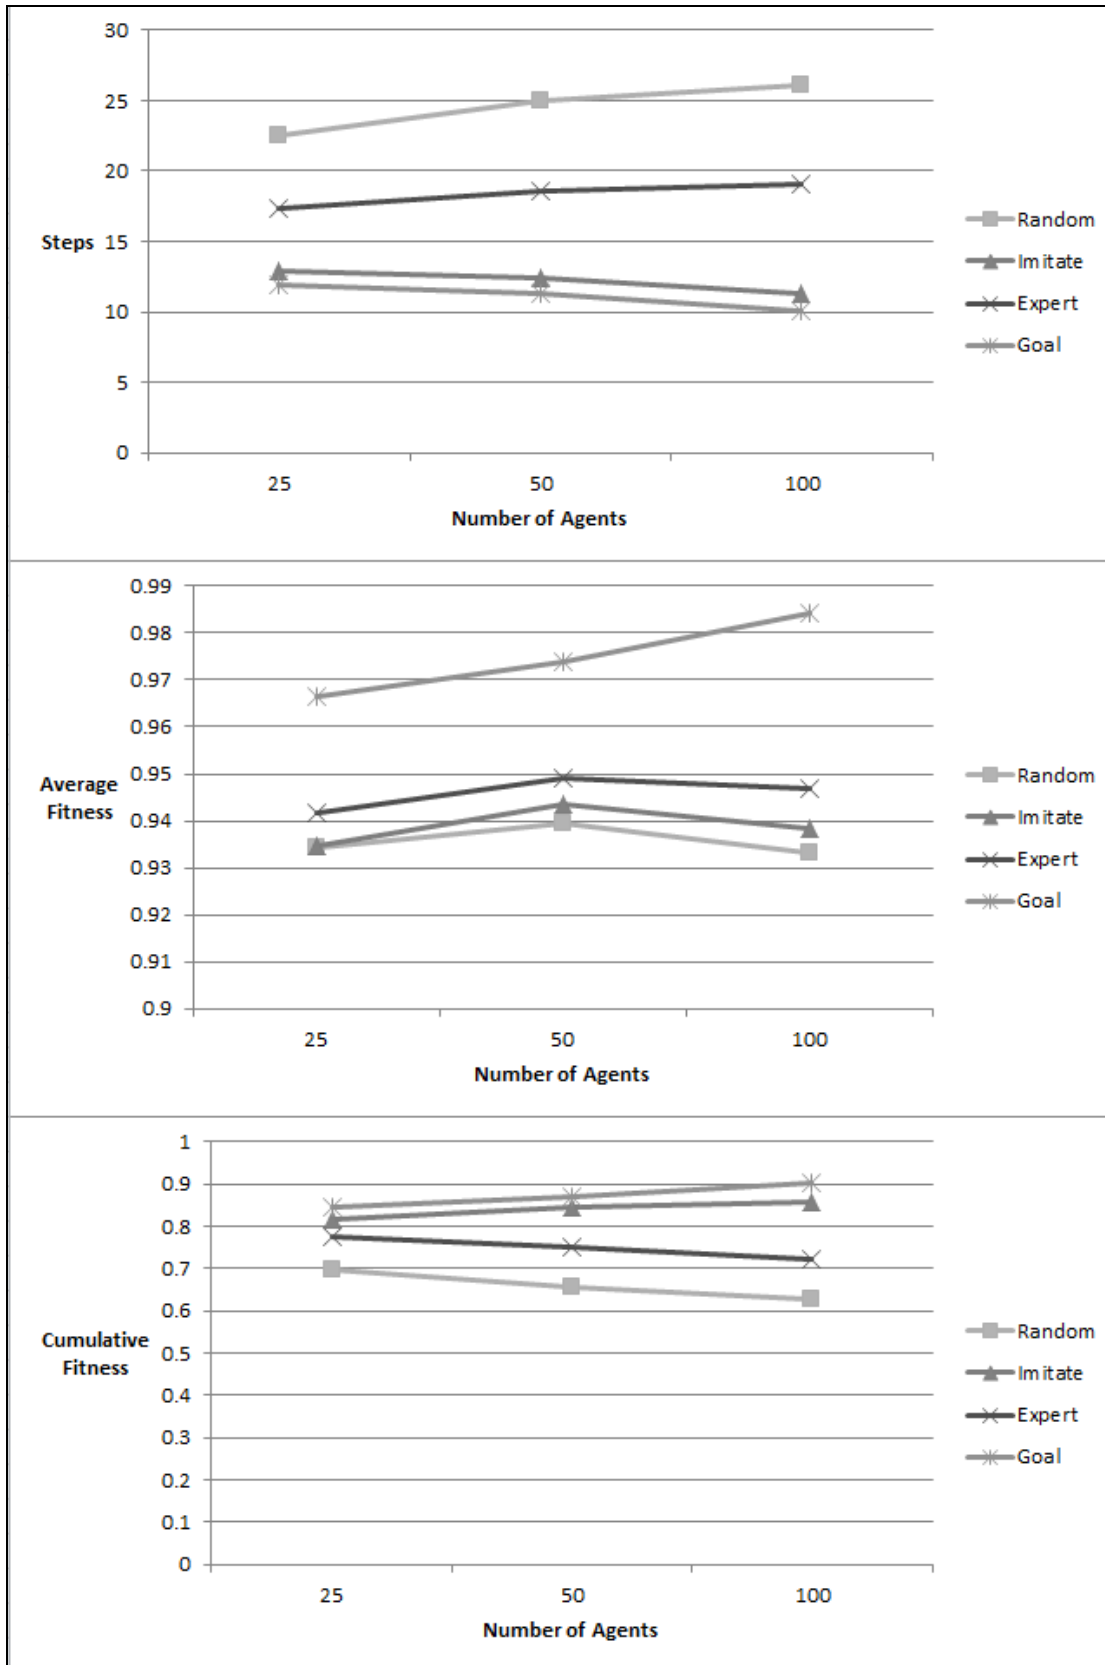


Figure 3b: Sensitivity analysis for number of agents (Values of 25, 50 and 100) for cluster landscape

## Discussion

The results just presented have a number of implications that warrant further consideration. In addition, there would seem to be two major issues that need to be examined since they could easily be viewed as a threat to the external validity of the findings:

- How realistic is the “expert” agent portrayal used in the simulation?
- To what extent is the exceptional performance of the goal-setting agent likely to generalize to a real world setting?

These three topics are now discussed.

### ***Implications of the Findings***

As previously stated in the results section, there is little ambiguity with respect to interpreting the findings of the simulation model: the two types of agents that were informed primarily through observing nearby agents—in other words, the homophilic agents—outperformed the expert and control agents in every rugged environment and on three measures reported: average fitness, speed to achieve peak fitness (steps), and cumulative fitness over the course of the run.

There are some additional implications of these findings that deserve some attention. One of these involves an evolutionary argument. Specifically, evolutionary economists argue that utility (the function that guides our decision making) must inevitably come to represent an individual’s estimate-of-fitness through the process of natural selection (Gandolfi, Gandolfi, & Barsh, 2002, p. 97). Whatever other benefits might accrue from homophily, the simulations show that it can be an effective information gathering strategy in environments that exhibit a particular type of complexity (ruggedness). If such environments are routinely encountered by a species, it would therefore represent a beneficial survival trait for individuals to move naturally closer to each other—either consciously or unconsciously—thereby enhancing the information value of observing neighbors. In the real world, there are many more than the 10 attributes used in the simulation. Thus, the safest strategy would be to group with individuals having the same uncontrollable attributes (e.g., age, ethnicity), then to mimic as many other attributes as possible so as to get the maximum likelihood that a neighboring agent’s high fitness search outcomes—in the present *or at some future date*—will be replicable. In such a world, social contagion and a conscious preference for homophily would be manifestations of the same underlying evolutionary imperative.

The nature of decomposability also suggests that a particular cluster of fitness attributes can be “nearly decomposable”—the term used by Herbert Simon (1981)—from other attributes. Agents who recognize a particular decomposition of fitness should be motivated to establish social groupings built around homophily across those clustered attributes. At the same time, they should be motivated to participate in other social groupings based upon different attribute clusters. This would provide an information gathering explanation for the “small world” phenomenon frequently observed in sociology (e.g., Watts, 2003). With respect to these small worlds, it has further been noted that the number of agents in a cluster often follow a power law distribution. Such distributions tend to occur when the attraction exerted by a particular cluster grows as agents join the cluster, i.e., “the rich get richer” (Barabasi, 2003; Gill, 2010). The fact that the rate of information gathering potential improves as agents join a cluster—since it provides more opportunities to observe agents experimenting with nearby states—could represent part of the rationale for that increasing attraction.

There is some empirical evidence that information gathering and experimentation can be accelerated when clusters of self-similar agents grow. The growth of cities has been described as an economic process in which self-sufficiency was replaced by specialization (Ridley, 2010). Different cities, however, have historically specialized in producing different goods or knowledge. That

means that a particular city or region, such as Florence during the Renaissance or today's Silicon Valley, would provide a venue in which many agents sharing similar characteristics could observe each other experimenting in close proximity. The productivity of these clusters in producing innovations should be self-evident. As Ridley (2010, p. 221) commented:

posterity will stand amazed at the thought giants like Gordon Moore and Robert Noyce, Steve Jobs and Segey Brin, Herb Boyer and Leroy Hood all lived at the same time and in the same place.

It would be foolish to discount the other economic factors that provide benefits from locating in a self-similar cluster, such as the ready availability of a workforce with suitable skills. Nevertheless, it is not necessary that idea generation be co-located with production. The fact that ideas, and not just products, emerge from these clusters supports a view that clustering may provide significant search benefits, accelerating the pace of discovery.

### ***Validity of Expertise Model***

The analysis presented in Figures 1 through 3 does not paint a very flattering portrait of expertise. In only one of the cases ( $K=0$ ) tested did the “expert” agent outperform the remaining agents at a statistically significant level; even then, it did so by a margin sufficiently narrow as to be almost immaterial. In many of the cases, expert advice proved little more beneficial than no advice at all.

One obvious criticism that can be leveled against the research method employed in this paper is that the technique used to simulate “expertise” was chosen largely as a matter of convenience (i.e., it was available as an external Excel function that could be called from the simulation, as noted in Appendix C) rather than being the best possible analytical technique. It is quite valid to argue that other statistical techniques, such as robust regression (e.g., Starbuck, 2006) and approaches more suitable to limited dependent and independent variables (e.g., Maddala, 1983), might have offered slightly more “expert” recommendations. In addition, a number of statistical techniques exist, such as structural equation modeling (e.g., Raykov & Marcoulides, 2006), that are intended to tease out more sophisticated relationships between variables. What is common to all these techniques, however, are (a) the assumption that variable impacts are decomposable unless specific interaction terms are added by the analyst, something that is usually done sparingly, and (b) the ultimate goal of using the results for hypothesis testing. Thus, in practice, the findings of these refined approaches tend to mirror standard regression results reasonably closely.

Furthermore, we should recognize that the regression approach employed *did not fail* under the conditions for which it was designed. In fact, the “expert” agent did quite well when it came to the decomposable case and, of course, that is the underlying landscape for which multiple regression algorithm is intended. It is also fair to note that in a “real world” setting, the benefits of regression analysis are likely to be even greater. Regression and other hypothesis testing techniques are intended to assess and provide an estimate of the error associated with observations. In the simulation as conducted, however, there was no error. Every time a particular combination of attributes was encountered, precisely the same fitness value was the result. Homophilic techniques, such as imitation and goal-setting could easily take false steps as a result of observational error in assessing the fitness of neighbors. The “expert” technique employed would, instead, attempt to correct for such error and alert the expert when significance could not be detected.

Given that the simulation's “expert” toolbox is only designed for low  $K$  landscapes, perhaps we should not be too concerned that it fails dramatically as landscapes move into the complex ( $K \approx N/2$ ) or chaotic ( $K \approx N-1$ ) ranges. That is probably a sensible attitude provided that experts always recognize ruggedness when it is encountered. There is a serious problem, however:

*Ruggedness does not necessarily advertise its presence.*

To justify this assertion, imagine an ideal world where our tools fully recognize the landscape structures from which observations are drawn. In such a world, where landscape ruggedness is present the tool would identify those variables that contribute to fitness in a decomposable way and would indicate zero significance for those variables that contribute only or mainly through complex interactions. Expert advice regarding the “proper” setting for individual attributes would then be limited to decomposable variables. No harm done and fitness should be reached more quickly since trial and error would not be required for those variables.

The problem is that prior research has suggested that the world we face is not so ideal. To the contrary, when significance testing or multiple regression is performed on observations drawn from complex or chaotic fitness landscapes, illusory significance values are frequently encountered (Gill & Sincich, 2008). The situation becomes particularly acute as entities migrate towards peaks—making it a particular concern for fitness landscapes, where such migration is to be expected—and occurs for the complex ( $K \approx N/2$ ) and chaotic ( $K \approx N-1$ ) ranges. In other words, a landscape can appear to be decomposable when, in fact, it is not.

A significant limitation of that earlier (Gill & Sincich, 2008) study—performed using spreadsheets to implement the simulation—was that it only tested one hill climbing strategy (each agent always chose the *best* available neighbor in moving towards fitness) and did not simulate general NK landscapes. The present study therefore offered an opportunity to extend and replicate the earlier findings using a different landscape structure and different agent behaviors.

To perform this replication, a special interface was designed that allowed the step-by-step progress on an individual run to be displayed. An example of the information gathered is illustrated in Figure 4, showing a fairly typical single simulation for an NK landscape with  $N=10$ ,  $K=6$  and 100 control agents. What is immediately evident is that as agents migrate to peaks, a substantial number of attributes immediately exhibit significance (3 after just 1 step). By the time the agents were fully distributed across the peaks (this particular trial produced 49 peaks, of which 42 were ultimately occupied), 4 attributes were deemed significant at the  $p < 0.05$  level and a respectable cross-sectional r-squared of 0.49 had been achieved (N.B. it is not unusual to see r-square values as high as 0.90 on trials with the same parameters). Well before peaks have been attained, however, significances appear and disappear, indicating that the significance values are highly unstable. To further illustrate this, a landscape with the same characteristics as the landscape in Figure 4 was replicated using the same seed for 10 trials, meaning that only the initial placement of agents varied. The results are shown in Table 3. The significance count—after just one step—ranged from 0 to 4. These results qualitatively replicate the earlier findings (i.e., Gill & Sincich, 2008), which argued that the observed significances were illusory with respect to the underlying landscape since they depend so heavily on the positioning of agents. The danger here is that significance testing results as (apparently) strong as those encountered along the way in Figure 4 could easily mislead a researcher into assuming that the underlying landscape was decomposable.

Is it plausible to speculate that a “real world” researcher might erroneously assume decomposability as just suggested? One way of considering the question is to examine the guidance we are given on research methodology. For example, Anol Bhattacharjee’s (2011) elegant open access textbook on social science research methods covers research design and methods, both qualitative and quantitative. Throughout the book, decomposable problems are largely assumed—a fact underscored by its emphasis on hypothesis testing, the nature of the block models/theories presented as illustrations and the quantitative methods described. In two places, however, the book sensibly cautions the reader about interactions (e.g., “it is not meaningful to interpret main effects if interaction effects are significant,” Bhattacharjee, 2011, p. 89). Unfortunately, we are not told *how* to identify if interactions are present or *what* the consequences of these effects will be. In particular, we are not warned that these effects can, at moderate levels of ruggedness, lead to patterns of sig-

nificance that resemble what a decomposable landscape might produce (Figure 4) and that such patterns will, most likely, fail to replicate if subsequent random samples are taken (Table 3).

Step	Average Fitness	Percent on Peaks	Percent Peaks Occupied	Percent Peaks Above Median	Significant Variables	R-Square
1	0.6103899	0.08	0.163265306	0.625	3	0.31292256
2	0.68625146	0.18	0.346938776	0.611111111	2	0.27430085
3	0.72153384	0.25	0.448979592	0.6	3	0.27457188
4	0.74408511	0.29	0.489795918	0.517241379	1	0.1709474
5	0.77208508	0.38	0.571428571	0.552631579	1	0.20242939
6	0.786479	0.42	0.612244898	0.571428571	1	0.21386061
7	0.80540651	0.52	0.714285714	0.519230769	1	0.18901762
8	0.81617343	0.57	0.755102041	0.543859649	2	0.22584693
9	0.82247229	0.61	0.755102041	0.524590164	2	0.22189453
10	0.83259305	0.71	0.836734694	0.521126761	3	0.27878808
11	0.84224429	0.77	0.836734694	0.532467532	4	0.34427224
12	0.84498314	0.8	0.836734694	0.5375	3	0.34224289
13	0.84694426	0.81	0.836734694	0.543209877	3	0.34942598
14	0.85394108	0.84	0.836734694	0.547619048	4	0.41126575
15	0.85969485	0.88	0.857142857	0.568181818	4	0.43728572
16	0.86127878	0.91	0.857142857	0.571428571	4	0.44943271
17	0.86334099	0.93	0.857142857	0.569892473	5	0.45197996
18	0.86525152	0.94	0.857142857	0.574468085	5	0.4596387
19	0.8656833	0.94	0.857142857	0.574468085	5	0.46153171
20	0.8657571	0.95	0.857142857	0.578947368	5	0.45112335
21	0.86891944	0.98	0.857142857	0.581632653	5	0.49505811
22	0.86891944	0.98	0.857142857	0.581632653	5	0.49505811
23	0.86891944	0.98	0.857142857	0.581632653	5	0.49505811
24	0.86983703	0.99	0.857142857	0.585858586	4	0.4888717
25	0.87009592	1	0.857142857	0.59	4	0.49144741

Figure 4: Step-by-step presentation of results for a simulation with N=10, K=6 and 100 control agents. The “Significant Variables” column shows the number of regression coefficients for which  $p < 0.05$ .

To bring this particular topic to a close, it is worth noting that a number of arguments have been advanced as to why we might expect certain landscapes to be rugged, most notably those typically studied by business (Gill, 2010) and education (Gill & Jones, 2010) researchers. Although it is certainly possible that the artificiality of the “expert” agent in the simulation contributed to its poor performance, it is likewise possible that its failure reflects an intrinsic weakness in empirical methods grounded in statistical significance hypothesis testing; methods routinely employed in the social sciences. The earlier-mentioned widespread failure of experts in predicting the behavior of complex environments supports the view that the failure of expertise observed in the present study—albeit somewhat artificially constructed—is not entirely farfetched.



**Table 3: Results of regression after first step of migration for 10 different runs on 44 peak NK landscape where N=10, K=6, Number of Agents = 100**

Step	Average Fitness	Percent on Peaks	Significant Variables	R-Square
1	0.628804	0.11	4	0.234517139
1	0.621097	0.1	1	0.195048756
1	0.625857	0.09	2	0.252318815
1	0.608299	0.08	2	0.173234526
1	0.623799	0.12	2	0.200305508
1	0.628255	0.08	0	0.09801366
1	0.62412	0.12	1	0.120089824
1	0.639791	0.08	2	0.266346714
1	0.644239	0.08	2	0.22644826
1	0.640507	0.1	3	0.25110181

### ***The Primacy of Goal-setting***

Another surprising aspect of the simulation was how well goal-setting agents performed relative to other types of agents. In some respects, this difference was actually understated as a consequence of how the simulation was constructed. Of particular note, each agent type could examine an adjacent state's fitness without actually occupying that state. That meant that each unsuccessful search—i.e., where the state that was being tested had lower fitness than the state occupied—took place at a cost of 1 step and without temporary loss of fitness. A plausible alternative algorithm would have required the agent to move to the lower fitness state in order to ascertain its fitness then, in the following pass, step back to the original state—at a cumulative cost of 2 steps and reduced cumulative fitness. Table 4 presents summary results using the modified 2-step algorithm just described with base case parameters. Compared with Table 2, the revised approach leads to increased steps and significantly reduces cumulative fitness for all agent types. The effect, however, is substantially less pronounced for goal-setting agents, since those agents perform such local testing only where no higher fitness goal can be established within the specified visibility range. Thus, requiring each agent to visit a state in order to test its fitness actually widens the advantage of goal-setting agents over agents of the other types.

**Table 4: Results for 100 runs of the base case on an NK landscape (N=10, K=6, Visibility=2 and Number of Agents=50) when requiring agents visit a state to test it**

Value (SE)	Random	Imitator	Expert	Goal
Average Fitness	.8748 (.0330)	.8801 (.0316)	.9017 (.0258)	.9091 (.0367)
Percent Peaks Occupied	.8299 (.1009)	.5748 (.1300)	.7000 (.1083)	.3389 (.1058)
Percent Peaks Above Median	.6848 (.0951)	.7034 (.1379)	.7878 (.1000)	.8190 (.1285)
Number of Steps	61.6200 (7.2757)	38.8200 (6.2632)	63.4500 (8.7297)	11.6400 (3.5483)
Cumulative Average Fitness	.4359 (.0429)	.5083 (.0724)	.4615 (.0517)	.8201 (.0848)

One reasonable criticism that can be made about goal-setting agent behavior is that when high visibility levels are specified it ceases to be homophilic since an agent can select distant agents rather than being limited to nearby ones. Indeed, when we set visibility to its maximum (the number of attributes in the landscape) we can be guaranteed that all goal-setting agents will ultimately end up on the same peak: the highest fitness state encountered by *any agent* over the course of its

search. There are two counter arguments that can be made in response. First, even with limited visibilities consistent with homophily, goal-setting agents outperform all other agent types.

The second argument is more subtle. It is evident that the marginal cumulative fitness and step advantages of visibility drop quickly after values of about 3 (Figures 2a and 2b). For the 10 attribute landscape, at a distance of 3 about 17% of all other states (175/1023) are visible to our agent. Returning to our earlier evolutionary arguments, if we assume that achieving additional visibility is likely to involve considerable energy expenditure on the part of an agent (who must locate more distant agents), we may expect the marginal costs of achieving additional visibility to exceed the marginal benefits at relatively short visibility distances. Once again, we see how homophily emerges as a beneficial trait for entities routinely encountering rugged landscapes.

Another implication of the goal-setting agent findings relates to the nature of expertise provided. Obviously, one of the problems facing the expert-advised agents as modeled in the simulation is that they are getting advice in the form of hypothesis tests on the effect of individual attributes. Since rugged landscapes are not structured that way, the expert constrained to describe the world in this manner is at an obvious disadvantage (see the previous section for the argument as why “real world” social scientists may be similarly constrained, thereby exhibiting a form of “trained incapacity,” Gill, 2010). There are, however, many other forms in which expertise could be accumulated and expressed. Many would not be subject to limitations presented here.

As an example, a particularly powerful manner in which experts could acquire and convey their findings is in the form of “best practice” descriptions that present a *story* describing *all* the attributes of a particular high fitness case. Where expertise is provided in this form, it could become the basis of goal-setting as a substitute for direct observation by the client agent. Here, once again, empirical evidence offers some support. Cognitive scientists have also observed that our minds are “exquisitely tuned to understand and remember stories” (Willingham, 2009, p. 51). Perhaps the story’s benefits in a complex environment guided that evolution. In business environments, at any rate, the “story” has proven to be a particularly resonant form of communication (e.g., Akerlof & Shiller, 2009; Heath & Heath, 2007; Gill, 2010). In fact, the most practitioner-focused of the business research journals, the *Harvard Business Review*, uses stories as its preferred form (Rynes, Giluk, & Brown, 2007, p. 999).

## Limitations and Directions for Future Research

It has been asserted that all models are wrong, but that they may nevertheless be useful (Serman, 2002). With respect to computer simulations of the type described in the present paper, this observation is particularly true. In designing the model described, the decision was made to keep the model as simple as possible; that is obviously a limitation. At the same time, the flexibility provided by the simulation tool also offers us the means of working around many of these limitations; these variations represent a direction for future research. Thus, limitations and directions are discussed in tandem, organized according to:

- Landscape structures (Table 5),
- Agent behaviors (Table 6), and
- Population dynamics (Table 7).

### ***Landscape Structures***

The simulation described in the present paper employed a static NK-landscape or a static NK cluster landscape with fitness distributed according to a uniform fitness distribution. The limitations imposed by these design decisions and how they might be addressed in future simulations are presented in Table 5.

**Table 5: Limitations and Future Directions for Landscape Structures**

Limitation	Future Direction	Comments
NK landscape does not necessarily reflect structures likely to be encountered in real world settings.	Develop alternative landscape constructions and determine if same behavior patterns emerge.	In describing complexity, Herbert Simon (1981) pointed out that artificial systems tend to evolve such that they consist of nearly decomposable subsystems. NK landscapes exhibit decomposability between clusters, but at a very superficial level. Multi-layer landscapes—along the lines of neural networks—would likely be required to simulate such systems.
Uniform fitness distribution is highly unlikely in real world.	Simulate landscapes where fitness is distributed according to power laws.	Power law distributions, such as the 80-20 rule, are commonly seen in fitness landscapes (Gill, 2010). Some criteria, such as “number of steps to achieve peaks” would be unaffected by this distribution, since they depend only on the ordinal value of fitness. Average fitness, on the other hand, would be hugely impacted by whether or not a very high fitness peak has been located by agents. As landscape ruggedness grows, this would likely tend to place a greater premium on strategies—including the simple “control” strategy—that visit more states before stabilizing.
Fitness is static for the landscape, making it unrepresentative of dynamic environments	Provide mechanisms for simulating continuous and step-changes to fitness	Real world environments change. The changes experienced may be gradual but they may be experienced as jolts (Gill, 2010). Any realistic landscape model needs to include the potential for both types of change. The challenge here is coming up with a plausible change model that does not force agents to behave in a particular way.
State fitness is path independent and all states are achievable, eliminating the challenge of path finding.	Build in the ability to establish “illegal” states or states that only certain adjacent states can access. Make some states irreversible once entered.	In describing the complexity of a task, the ruggedness and unfamiliarity of state fitness captured in the model present an incomplete model of fitness. At least one further dimension, involving the challenge of determining a path between a particular state and a particular goal, needs to be included (Gill & Murphy, 2011). This particular modification would likely exert greatest impact on goal-setting agents.

## Agent Behaviors

Limitations and possible modifications of agent behaviors are presented in Table 6.

**Table 6: Limitations and Future Directions for Individual Agent Behaviors**

Limitation	Future Direction	Comments
Agents are certain of fitness upon entering a state.	Add an uncertainty component to fitness.	In the real world, agents must estimate state fitness using internal functions such as utility (Gill, 2010).
Agents always seek optimum fitness and do not consider the cost of search.	Create rules by which agents stop searching once acceptable fitness is achieved. Include a cost for transitions between states.	A good example of a rule that might be established to end search is satisficing (Simon, 1955). A more realistic approach to search would also take into account costs as well as possible benefits. This may require a “time horizon” parameter.
Agents can modify all attributes that contribute to fitness.	Provide a mechanism for establishing attributes that individual agents cannot modify.	Some attributes that may impact an agent’s fitness, such as age and ethnicity, are not normally under the agent’s control. When these attributes participate in interactive clusters, agents may find themselves searching profoundly different landscapes.
“Expert” agents are mechanical	Create alternative types of expertise.	Empirically mapping attributes to fitness makes for unimaginative experts. Other forms of expertise, such as knowledge of specific cases or underlying fitness structure could make for more competitive experts.

## Population Behaviors

In the existing model, it is assumed that each run is populated by agents of the same type, that the fitness of a state is unaffected by the agents that occupy it, and that the main impact of agents on other agents is through the provision of information. Some limitations of, and possible modifications to, these assumptions are presented in Table 7.

**Table 7: Limitations and Future Directions for Population Assumptions**

Limitation	Future Direction	Comments
Populations are homogeneous.	Allow agents of different types to populate the environment.	It is well established in the diffusion literature that individual agents behave differently. For example, an agent might be an innovator, an early adopter, a late adopter, or a laggard (Rogers, 2003). Such heterogeneous populations might well lead to more effective informing, a fact that simulations could verify. A similar diversity is likely to exist with respect to visibility, with some agents connecting to a vastly larger collection of agents than others (Gladwell, 2000).
Fitness of state is unaffected by the number of agents occupying it.	Create some states where fitness is reduced and increased by occupancy.	A state whose fitness is derived from supplying a particularly resource is likely to see fitness falling as more agents occupy it. On the other hand, states or clusters exhibiting network effects (e.g., adoption of a particular communications medium) are likely to see large increases as agents occupying the state increase.
Agents report the fitness of the state they occupy with complete accuracy.	Add uncertainty to agent reports of fitness and, potentially, the possibility that agents will misreport their fitness.	Even if an agent has perfect knowledge of the fitness of a particular state, there is no guarantee that it will be in the agent's capability or best interest to convey that fitness accurately to neighbors. Moreover, particularly where the fitness of a particular state is influenced by occupancy, real world agents may have an incentive to overreport or underreport their fitness.
Agents do not share historical path information with other agents.	Particularly in fitness landscapes where all states are not accessible, the ability to share history with other agents might prove beneficial.	Real world informing processes not only convey information about where to go but also how to get there. The popularity of certain types of books, such as biographies and "how to" books underscores this fact.

## Next Steps

As is invariably the case with simulations, the challenge presented by the dozen possible enhancements to the present model just proposed (Tables 4-6) is that they each involve incorporating more parameters and more specific assumptions into simulation. As a practical matter, considering the impact of all these modifications—and the many other changes than could be envisioned—would require a time consuming research program. To assess whether the benefits of such an undertaking would justify the effort it would require, we need to turn to the conclusions and consider what this exercise has already taught us.

## Conclusions

When an agent seeks to be informed about the environment, the decision of where to look for information is a critical one. In the study presented in this paper, that agent has three choices:

- *The environment itself*: The agent can directly test possible alternatives and determine their relative fitness.
- *An expert*: The agent can consult an expert whose knowledge—expressed as a series of simple hypotheses about the impact of each attribute on fitness—is derived from surveying the fitness of all the agents on the landscape.
- *Nearby agents*: The agent can look at nearby agents and determine if their fitness warrants moving towards their attribute values.

What the present study shows, quite convincingly, is that as the underlying complexity (ruggedness) of the environment (landscape) grows, nearby agents—when available—become a better informing source than either the expert or the environment itself. The simulation also demonstrates that the benefits of looking at far-away agents versus nearby agents drops off quite quickly as similarity decreases. Simply stated, the results make the case that homophily can be an effective informing strategy.

Many of the limitations of the simulation approach employed are described in the previous section and need not be repeated here. What is particularly interesting about these results, however, is the degree to which behaviors consistent with the simulation are observed in real world settings. Experts in many fields lament about the failure of their ideas to diffuse to practice. These findings show why we might have evolved to be suspicious of such expertise in a complex world. Homophily is widespread. This is consistent with an agent seeking to benefit from the informing advantage of being surrounded by other agents facing similar or identical environments. Social contagion breeds more perfect homogeneity. Paradoxically, this would hold true for traits generally considered as non-beneficial as well those considered beneficial. If an individual is close to a number of obese people, then he or she may need to become obese in order to achieve the maximum informing value from observing their experiences. It is possible that in a very rugged world, where a single misstep can lead to an agent's demise, that the information gathering benefits flowing from homophily may, on balance, exceed the cost of adopting traits that—on the surface—seem to reduce fitness. In such a world, the internal utility function that we use to compare the estimated fitness of our choices could easily evolve to favor conscious and unconscious imitation.

There is, of course, a price to be paid for informing through homophily. In such a world, fashions, fads and information cascades are inevitable. By encouraging us to cluster with others around a small number of peaks, it limits our willingness to strike out and explore other peaks. That can have serious consequences in terms of the overall fitness of the population. A single shock in a highly clustered world, such as the blight that led to the Irish potato famine, can be disastrous. Furthermore, fitness is often distributed according to a power law, meaning that a few choice spots are *much* higher than their surroundings would suggest. If we do not discover these peaks, the overall fitness of the population suffers. For this reason, we might expect an evolutionary stable strategy (ESS) to emerge whereby a small but relatively stable proportion of the population develops a trait in which the drive to explore and innovate—thereby expanding the range of known peaks available—exceeds the impulse of homophily. Diffusion studies (e.g., Rogers, 2003) suggest that precisely this type of diversity exists.

It would be both imprudent and incorrect to assert that the present study concludes that homophily leads to more effective informing than expertise. To begin with, the findings hold only for rugged landscapes; if the environment is such that attributes contribute to fitness independent of each other, it makes sense to go for expertise every time! Similarly, the findings only apply to a particular type of expertise; in fields such as medicine, experts are already well aware of the fragility of findings derived purely from sampling large populations of individuals and discount it heavily. There is also no reason to believe that expertise derived from knowledge of best practices or from a detailed understanding of the underlying mechanisms driving the structure of the land-

scape will be similarly fragile. In applied social science fields such as business and education, however, the landscapes being studied can be rugged (Gill, 2010; Gill & Jones, 2010) and studies that rely heavily on significance testing of observational data sprout like noxious weeds in the literature. As researchers in these fields, we should therefore not be surprised when practice fails to embrace our findings with the enthusiasm that we had eagerly anticipated.

## References

- Ahuja, G., Polidoro, F., Jr., & Mitchell, W. (2009). Structural homophily or social asymmetry? The formation of alliances by poorly embedded firms, *Strategic Management Journal*, 30, 941–958.
- 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.
- Altenberg, L. (1995). Genome growth and the evolution of the genotype–phenotype map. In W. Banzhaf & F. H. Eckman (Eds.), *Evolution and biocomputation* (pp. 205–259). Berlin: Springer-Verlag.
- Altenberg, L. (1997). NK fitness landscapes. In T. Back, D. Fogel, & Z. Michalewicz (Eds.), *The handbook of evolutionary computation* (pp. B2.7:5—B2.7:10). Oxford, UK: Oxford University Press.
- Aral, S., Muchnik, L., & Sundararajan, S. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *PNAS*, 106(51), 21544–21549. Retrieved from <http://www.pnas.org/cgi/doi/10.1073/pnas.0908800106>
- Bacharach, S., Bamberger, P., & Vashdi, D. (2005). Diversity and homophily at work: Supportive relations among white and African-American peers. *Academy of Management Journal*, 48(4), 619–644.
- Barabasi, A. L. (2002). *Linked*. New York, NY: Plume.
- Bhattacharjee, A. (2011). *Social science research: Principles, methods and practices*. Global Text Project. Retrieved from [http://dl.dropbox.com/u/31779972/Social\\_Science\\_Research.pdf](http://dl.dropbox.com/u/31779972/Social_Science_Research.pdf)
- Centola, D., Gonzalez-Avella, J., Eguiluz, V., & San Miguel, M. (2007). Homophily, cultural drift and the co-evolution of cultural groups. *Journal of Conflict Resolution*, 51(6), 905-929. doi: 10.1177/0022002707307632
- Choudhury, M., Sundaram, H., John, A., Seligmann, D., & Kelliher, A. (2010). "Birds of a feather": Does user homophily impact information diffusion in social media? *arXiv:1006.1702v1*. Retrieved from: <http://arxiv.org/abs/1006.1702>
- Christakis, N. A., & Fowler, J. H. (2009). *Connected: The surprising power of social networks and how they shape our lives*. New York, NY: Little Brown.
- Fiore, A., & Donath, J. (2005). Homophily in online dating: When do you like someone like yourself? *CHI 2005, April 2–7, 2005, Portland, Oregon, USA*, 1371-1374. doi:10.1145/1056808.1056919
- Frenken, K. (2006). A fitness landscape approach to technological complexity, modularity, and vertical disintegration. *Structural Change and Economic Dynamics*, 17, 288–305.
- Gandolfi, A. E., Gandolfi, A. S., & Barash, D. (2002). *Economics as an evolutionary science: From utility to fitness*. New Brunswick, NJ: Transaction Publishers.
- Gill, T. G. (2008). 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. (2010). *Informing business: Research and education on a rugged landscape*. Santa Rosa, CA: Informing Science Press.
- Gill, T. G. (2011). When what is useful is not necessarily true: The underappreciated conceptual scheme. *Informing Science: The International Journal of an Emerging Transdiscipline*, 14, 1-32. Retrieved from <http://www.inform.nu/Articles/Vol14/ISJv14p001-032Gill589.pdf>

- 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., & Jones, J. (2009). A tale of three classes: Case studies in course complexity. *Journal of IT Education*, 9, 1-29. Retrieved from <http://www.jite.org/documents/Vol9/JITEv9p001-029Gill717.pdf>
- Gill, T. G., & Murphy, W. (2011). Task complexity and design science. *9th International Conference on Education and Information Systems, Technologies and Applications: EISTA 2011*, Orlando, FL, July 19th-22<sup>nd</sup>.
- 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, NY: Back Bay Books.
- Golub, B., & Jackson, M. (2008). How homophily affects diffusion and learning in networks, *arXiv:0811.4013v2*. Retrieved from: <http://arxiv.org/abs/0811.4013>
- Golub, B., & Jackson, M. (2011). Network structure and the speed of learning: Measuring homophily based on its consequences. *SSRN*. Retrieved from <http://ssrn.com/abstract=1784542>
- Hanish, L. Martin, C., Fabes, R., Leonard, S., & Herzog, M. (2005). Exposure to externalizing peers in early childhood: Homophily and peer contagion processes. *Journal of Abnormal Child Psychology*, 33(3), 267-281.
- Heath, C., & Heath, D. (2007). *Made to stick*. New York, NY: Random House.
- Implementation Science. (2011). Retrieved 11/4/2011 from <http://www.implementationscience.com/>
- Ingram, P., & Morris, M. W. (2007). Do people mix at mixers? Structure, homophily, and the “life of the party.” *Administrative Science Quarterly*, 52, 558-585.
- Kauffman, S. A. (1993). *The origins of order*. Oxford, UK: Oxford University Press.
- Kauffman, S. A. (1995). *At home in the universe: The search for laws of self-organization and complexity*. Oxford, UK: Oxford University Press.
- Khurana, R. (2007). *From higher aims to hired hands*. Princeton, NJ: Princeton University Press.
- Kossinets, G., & Watts, D. (2009). Origins of homophily in an evolving social network. *American Journal of Sociology*, 115(2), 405-450.
- La Fond, T., & Neville, J. (2010). Randomization tests for distinguishing social influence and homophily effects. *WWW 2010, April 26-30, 2010, Raleigh, North Carolina, USA*, 601-610.
- Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, 43(7), 934-950.
- Levinthal, D. A., & Warglien, M. (1999). Landscape design: Designing for local action in complex worlds, *Organization Science*, 10(3), 342-357.
- Locke, E. (2004). Goal-setting theory and its applications to the world of business. *Academy of Management Executive*, 18(4), 124-125.
- Maddala, G. S. (1983). *Limited-dependent and qualitative variables in econometrics*. Cambridge, UK: Cambridge University Press.
- McPherson, M., Smith-Lovin, L. & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27, 415-444.
- Pfeffer, J. (2007). A modest proposal: How we might change the process and product of managerial research. *Academy of Management Journal*, 50(6), 1334-1345.
- Porter, M., & Siggelkow, N. (2008). Contextuality within activity systems and sustainability of competitive advantage. *Academy of Management Perspectives*, 22(2), 34-56.

## Homophily vs. Expertise

- Raykov, T., & Marcoulides, G. A. (2006). *A first course in structural equation modeling* (2nd ed.). Mahwah, NJ: Earlbaum.
- Ridley, M. (2010). *The rational optimist: How prosperity evolves*. New York, NY: Harper Perennial.
- Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science*, 46(6), 824–844.
- Rivkin, J. W. (2001). Reproducing knowledge: Replication without imitation at moderate complexity. *Organization Science*, 12(3), 274–293.
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). New York, NY: Free 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.
- Shanteau, J. (1992). Competence in experts: The role of task characteristics. *Organizational Behavior and Human Decision Processes*, 53, 252-266.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 129-138.
- Simon, H. A. (1981). *The sciences of the artificial* (2nd ed.). Cambridge: MIT Press.
- Solving social sciences' hard problems*. (2010, April 27). *Harvard Magazine*. Retrieved 11/4/2011 from <http://harvardmagazine.com/2010/04/social-sciences-hard-problems>
- Sorenson, O., Rivkin, J. W., & Fleming, L. (2006). Complexity, networks and knowledge flow, *Research Policy*, 35, 994–1017.
- Starbuck, W. H. (2006). *The production of knowledge*. Oxford, UK: Oxford University Press.
- Sterman, J. D. (2002). All models are wrong: Reflections on becoming a systems scientist. *System Dynamics Review*, 18(4), 501–531.
- Taleb, N. N. (2007). *The black swan*. New York, NY: Random House.
- Watts, D. (2003). *Six degrees: The science of a connected age*. New York, NY: Norton.
- Willingham, D. T. (2009). *Why don't students like school?* San Francisco, CA: Jossey Bass.
- Yuan, Y., & Gay, G. (2006). Homophily of network ties and bonding and bridging social capital in computer-mediated distributed teams. *Journal of Computer-Mediated Communication*, 11(4), article 9. Retrieved from: <http://jcmc.indiana.edu/vol11/issue4/yuan.html>
- Ziliak, S. T., & McCloskey, D. N. (2008). *The cult of statistical significance*. Ann Arbor, MI: University of Michigan Press.



## Appendix A

### Construction of NK Landscapes

The NK landscape, developed by evolutionary biologist Stuart Kauffman and described most fully in his book *Origins of Order* (Kauffman, 1993) is the artifact commonly used to simulate and investigate the properties of rugged fitness landscapes. This appendix summarizes its conceptual basis and construction.

#### ***Fitness Landscapes***

As described in the body of the paper, the general fitness landscape maps a set of attributes,  $a_1, \dots, a_N$  to a value that roughly corresponds to the desirability of the combination. In its most general form, this can be expressed as:

$$\text{Fitness} = f(a_1, a_2, \dots, a_N)$$

In a biological setting, the “fitness” value might refer to an entity’s likelihood of surviving to the next generation. In other settings, it might reflect some other outcome, such as the value of a firm. What is common to all fitness landscapes, however, is the notion that—over time—entities will attempt to migrate from combinations of low fitness towards combinations of higher fitness. In genetics, attribute changes occur through processes such as mutation and swapping that occurs through sexual reproduction. In settings where the entity is an intelligent agent, changes may occur through decision making.

In the most general fitness landscape, each attribute  $a_i$  affecting fitness may take on multiple discrete values or may be continuous. Similarly, there are no restrictions placed on the shape of the fitness function. Landscapes this general, however, have few general properties that can be used as a basis for further investigation.

#### ***NK Landscape***

The NK landscape is an attempt to enforce some structure on the general fitness landscape drawing upon Kauffman’s insights into the interactions that typically seem to take place between genes. Since these specific interactions are beyond the scope of the present paper, we turn directly to the structure of the model.

The NK landscape model was intended to be parsimonious with respect to parameters and tunable with respect to ruggedness, which is essentially synonymous with complexity. The landscape is formed by first assuming that there are  $N$  attributes, each of which is binary (0,1) in nature. For each attribute, a dominant fitness function is defined that depends on the attribute value itself plus the value of  $K$  additional attributes, i.e.,

$$f_1(a_1, K \text{ of the additional attributes})$$

$$f_2(a_2, K \text{ of the additional attributes})$$

...

$$f_N(a_N, K \text{ of the additional attributes})$$

The fitness value for a particular state is then determined by adding these values together, i.e.,:

$$\text{Fitness} = f_1 + f_2 + \dots + f_N$$

Since each function takes  $K+1$  binary attribute values as arguments, a function can take on  $2^{K+1}$  possible values. To simulate the individual functions, a random number can be assigned to each

combination. The relative impact of the individual functions on fitness can then be determined in some arbitrary manner, such as randomly assigning weights.

An important property of NK landscapes is the presence or absence of multiple peaks. A peak is defined as a state where modifying the value of *any single attribute* produces a decline in fitness. Peaks are important because most schemes for increasing fitness involve changing one attribute at a time. Thus, entities guided by fitness in their search strategy tend to find peaks “sticky”. Where multiple peaks exist on a landscape, it is therefore possible for entities to become stuck on local peaks whose fitness is far lower than that of the landscape’s maximum peak.

The values of K can vary from 0 (no interactions) to N-1 (every attribute interacts with every other attribute). These two end points can be simplified. Assuming  $w_1, w_2, \dots, w_N$  are the weights assigned to each fitness function, the K=0 case can be written as a simple linear combination of weighted attribute values:

$$w_1a_1 + w_2a_2 + \dots + w_Na_N$$

Since an arbitrary constant could be added without impacting the ordinal value of fitness, this type of landscape has the same type of structure used in multiple regression models. It should also be evident that this landscape has a single peak, since a best value (0 or 1) can be determined for each  $a_i$  depending upon whether  $w_i$  is positive or negative.

At the extreme where K=N-1, it should be evident that each fitness function  $f_1, \dots, f_N$  has the same set of arguments  $a_1, \dots, a_N$ , and therefore there is no reason not to combine them into a single function with random values assigned to each of the  $2^N$  attribute combinations. It is also possible to estimate the number of peaks in such a landscape. Since each combination has a value assigned randomly, each element has a  $1/(N+1)$  probability of being a fitness peak. This is based upon the fact that the total number of adjacent states is N, and 1 additional state (the value itself) must also be considered. There are, however,  $2^N$  possible combinations in the landscape. Thus, the number of peaks can be estimated by the formula:

$$2^N/(N+1)$$

In the states between 0 and N-1, no simple formula for estimating the number of peaks and other properties has been reported. As a consequence, much of Kauffman’s (1993, 1995) work has involved simulating these landscapes to examine their properties. What becomes qualitatively clear from these simulations is that:

- As K approaches N, the number of peaks increases
- As K approaches N, the peaks become more uniformly distributed across the landscape

Thus, we can expect that the most successful search strategies may vary considerably depending upon the K value.

### ***NK Landscape Variations***

Many possible variations on the NK landscape are possible and have been used in research. Since they have not been employed in the present paper, they are sketched out only briefly here, with an eye towards possible extensions of the work in the future.

Kauffman (1993, 1995) introduces the idea of co-evolving NK landscapes. In this model, the fitness landscape for one species has some attributes that depend upon the population of one or more different species. As a consequence, if the populations of those other species changes (as a consequence of migrating to a different fitness state), the fitness landscape of the original species may become altered—in some cases radically. Such coupled systems can reach a stable state or could remain in a constant state of change depending on the nature of the landscapes and the cou-

plings between them. This type of landscape could therefore be employed to study a dynamic, rather than a static, model of fitness.

An NK variation that is somewhat more general than the original model has also been proposed (Altenberg, 1995, 1997; Frenken, 2006). In Kauffman's model, the number of individual functions that contribute top fitness (i.e.,  $f_1, \dots, f_N$ ) is constrained to be the same as the number of attributes. There is, however, no obvious conceptual reason for this 1:1 relationship between functions and attributes. Instead, it is proposed that the functions should refer to however many specific characteristics are known to affect fitness, while the functions should map attributes to these characteristics. To use cooking as an example, if our tongue responds to 5 sensations—Wikipedia refers to them as sweet, bitter, umami (savory), sour, and salty—we might choose to map ingredients (attributes) to these values, rather than having a separate function for each ingredient. This also opens up the possibility of multi-layer landscapes.

Another NK variation involves forcing decomposability between the various functions  $f_1, \dots, f_N$  by limiting the variables each function includes. For example, suppose we partition the landscape at some middle value  $P$ , where both  $P > K$  and  $N - P > K$ . We can then specify the following:

For functions  $f_1, \dots, f_P$ , only values  $a_1 \dots a_P$  may be used as arguments

For functions  $f_{1+P}, \dots, f_N$ , only values  $a_{1+P} \dots a_N$  may be used as arguments

By making this separation, we can then treat the process of maximizing fitness for  $a_1, \dots, a_P$  as an independent problem from that of maximizing fitness for  $a_{1+P}, \dots, a_N$ . This would simplify our search process (see Appendix B) and would allow us to incorporate any knowledge we have regarding the decomposability of a specific landscape into our simulation.

Naturally, extensions to the NK landscape such as the ones described come with the price tag of added parameters to be explored. Before asserting the generalizability of any simulation findings, however, it would make sense to explore behaviors on as many different structures as possible.

## Appendix B

### Construction and Behaviors of NK Cluster Landscapes

The NK cluster landscape, used as an alternative to the NK landscape to provide a contrasting mechanism for simulating increasing ruggedness, was also inspired by Kauffman (1995), although it is not clear that it was ever used for research purposes. It creates the simplest possible landscape based upon decomposable clusters that consist of either interacting sets of variables or individual variables. In this appendix, we briefly explore its construction and some of its key behaviors.

#### Constructing an NK Cluster Landscape

NK landscapes, generally speaking, define a mapping between binary (0 or 1) attributes and fitness, although non-binary attributes can also be supported. The number of attributes is specified by the attribute  $N$ . In addition, there are  $K$  interrelationships, with possible values ranging from 0 to  $N-1$ . Whereas the NK landscape (see Appendix A) implements these relationships through a series of  $N$  functions, each having  $K+1$  variable arguments, the cluster landscape simply breaks the top level fitness function into one or more separate functions, each of which has one or more arguments. The functions are decomposable, since no attribute is allowed to appear in more than one function. One way to establish how the fitness function decomposes is through a process that Kauffman (1995) likens to that of connecting buttons with threads, illustrated in Figure A2.1.




Figure A2.1: Button Analogy to Landscape, with 10 attributes and 5 links ( $N=10$ ,  $K=5$ )

As more and more threads are attached between buttons, clusters of multiple buttons begin to emerge. As an illustration, in Figure A2.1 the clusters are 7, 8, 1-2, 3-4-5-6, and 9-10. These clusters are important because they define the structure of the fitness function. For example, if we

refer the value of each attribute (button) as  $v_1 \dots v_{10}$ , then the figure would imply a fitness function of the form:

$$\text{Fitness} = f_A(v_1, v_2) + f_B(v_3, v_4, v_5, v_6) + f_C(v_7) + f_D(v_8) + f_E(v_9, v_{10})$$

Furthermore, since each  $v$  can only take on two values,  $f_C$  and  $f_D$  can be replaced with simple coefficients. To simulate relationships, the same approach Kauffman (1993) used for NK landscapes can be applied: random numbers are generated to simulate the fitness of each combination in a cluster. Thus, the clusters associated with  $f_A$  and  $f_E$  would each have 4 associated values randomly assigned (for the 0-0, 0-1, 1-0 and 1-1 combinations of the associated attributes and  $f_B$  would 16 combinations (for the 0-0-0-0, 0-0-0-1, ..., 1-1-1-0, 1-1-1-1 combinations). It should be readily evident that as cluster size grows, so does the number of combinations—with  $2^X$  combinations for clusters of size  $X$ .

Omitting any “thread” that connects a button to a cluster it is already connected to, it should be apparent that when  $K$  (number of threads) reaches  $N-1$ , every button belongs to the same large cluster. At that point, landscape fitness can be simulated by assigning a random number to all  $2^N$  possible combinations in the landscape, making it the same as the  $N, N-1$  case for the regular NK landscape. This is the “maximally rugged” or chaotic landscape. On the other hand, when  $K=0$ , then we have a fully decomposable landscape, which also happens to be the same as the NK landscape case.

By running simulations on the computer, Kauffman (1995, p. 57) shows that as the number of threads approaches a value just under  $N/2$ , we start to see maximum cluster size increasing dramatically as additional threads are added. This makes sense; where the network is loosely connected, it will be dominated by two button connections. Once reaching a certain size, however, a single large cluster will start to form and the likelihood that new connections will join buttons to that particular cluster will become increasingly high. Kauffman attributes particular significance to this region, which he describes as the boundary between chaos and order.

### ***Peaks in a Cluster Landscape***


Because of the simple construction of a cluster landscape, it is relatively easy to estimate the number of fitness peaks. Suppose, for example, that a cluster has  $M$  elements. That implies that any individual element has  $M$  adjacent neighbors, plus itself—meaning it must be the largest of  $M+1$  elements. Since the NK cluster model described assigns cluster values randomly, the number of peaks can be estimated by the same formula used for chaotic NK landscapes (see Appendix A), i.e.,

$$2^M/(M+1)$$

This can be demonstrated for the 2-way link, where the formula produces an estimate of  $4/3$ . Assume that we rank-order fitness (e.g., 1 is highest and 4 is lowest). Then, two distinct peaks exist only where the two highest values are at opposite corners (i.e., 1,1 and 0,0 or 0,1 and 1,0). As illustrated in Figure A2.2, this occurs for only 8 of the 24 possible orderings of fitness meet this criterion, meaning the average number of peaks is  $(16*1 + 8*2)/24$ , or  $4/3$ .

Depending upon how a cluster landscape is constructed, the estimated number of peaks can vary considerable. Suppose that  $K < N/2$  (the point at which some large clusters necessarily begin form). Then, in the case where the landscape consists of nothing but two way links, then the number of peaks will be given by the formula:

$$(4/3)^K$$



**Figure A2.2: Possible 2-Variable Orderings of Fitness.  
Ellipses Identify where 2-Peak Combinations Occur**

On the other hand, if all  $K$  elements happen to be connected as one large cluster, the number of peaks would be:

$$2^{K+1}/(K+2)$$

To get a sense of how these two relate to each other, we can make some gross simplifying approximations, namely:

$$(4/3)^2 \approx 2, \text{ meaning } (4/3)^K \approx 2^{K/2}$$

$$2^{K+1}/(K+2) \approx 2^K$$

Thus, the minimum number of peaks is roughly equal to the square root of the maximum number of peaks for a given  $K$ . For  $K \ll N/2$ , our best estimate is probably our lower bound. As  $K$  approaches  $N/2$ , our best estimate will rapidly move from the minimum to the maximum value. For  $N/2 \ll K \leq N-1$ , our best estimate becomes the upper bound.

### **Search in a Cluster Landscape**

The number of peaks in the landscape is important because peaks—by their very nature—tend to be sticky in the face of incremental adaptation. The impact of  $NK$  landscape ruggedness on search, however, can be even greater. Recalling that each cluster is simulated randomly in an  $NK$  landscape, the only way to determine the peak value in a cluster is to examine each element. Consider, then, a 100,0 landscape (completely decomposable). In this case, we would need to incrementally test each attribute's impact of fitness independently, requiring 100 tests.

Now consider a 100,20 landscape. If that landscape consisted of 20 two-attribute pairs (each having 4 possible combinations, as per the grids in Figure A2.2), the total number of tests would be:

$$20 * 4 \text{ (for the 40 attributes participating in pairs)} + 60 \text{ (for the remaining variables)} \\ = 140$$

On the other hand, suppose that all 20 links happened to participate in a single cluster of 21 attributes. Then the number of combinations would be:

$$2^{21} \text{ (combinations within the cluster)} + 79 \text{ (for the remaining variables)} \\ \approx 2 \text{ million}$$

This difference in search is substantially higher than the difference in the estimated number of peaks, which is:

$$(4/3)^{20} \approx 315 \text{ and } 2^{21}/22 \approx 95,000$$

To conclude, it is also important to note that the type of efficient search and peak identification computed here would be available *only where knowledge of how the fitness landscape decomposes is available*. Lacking such knowledge of how variables are clustered, systematic search is likely to be nearly impossible for fitness landscapes impacted by many attributes. It therefore follows that such knowledge of landscape structure is quite likely the most valuable type of information to have when attempting to predict the likely cost of search.

### Comparison of Cluster Peaks to NK Landscape Peaks

For comparison purposes, the number of actual peaks (averaged over 100 runs) for NK and cluster landscapes as K varies from 0 to 9 is shown in Figure A2.3.




Figure A2.3: Comparison of landscape peaks for NK and cluster landscapes with N=10

## Appendix C

### Overview of Software Used in Simulation

In order to achieve maximum flexibility in customizing landscapes and agent behavior, the software used to develop the simulation whose results are presented in this paper was programmed from scratch using the C# programming language and the *Visual Studio C# 2010 Express Edition* development environment. The assumptions and interface employed are now briefly described. Researchers who would like to acquire a no-cost copy of the software and source code licensed under the Creative Commons should contact the author directly.

#### Assumptions

Even as described by Kauffman (1993, 1995) and others, there is a certain amount of flexibility associated with developing the model. Key assumptions made in construction are now presented.

#### Landscape

NK and cluster landscapes were created using the procedures described in Appendix A and Appendix B. The following assumptions were required:

- To ensure maximum comparability between the NK and cluster landscapes, fitness was normalized for each landscape such that the lowest fitness state was set to 0.0 and the highest fitness state was set to 1.0.
- A uniform distribution was used to generate fitness values. This should have no impact on hill climbing—which depends on ordinal values—but does mean that average fitness and cumulative average fitness values could differ for other distributions. Having a different distribution of fitness would also likely have some effect on the regression results used to direct expert recommendations.
- The landscape establishes the fitness of all states when it is created (using a random seed). After that point, it is assumed to be static. Peaks are also identified at that time and stored in a collection.
- Separate seed values were available for creating landscapes, distributing agents prior to movement, and controlling agent movement. During comparison simulation runs, this meant that for each trial, the four types of agents all began on the same landscape.
- For convenience, the integer position of each state on the landscape also identifies its attributes, using the base 2 bits of the position (e.g., position 0 has all attributes in state 0). 64 bit integers are used to hold position, placing a theoretical limit on N. As a practical matter, once N values in the high 20s are reached, the memory requirements for the landscape would likely exceed available RAM on most computers.

#### Agents

The algorithms controlling agent behaviors are described in the research design section of the paper. In addition to the 4 basic types (control, imitating, expert-guided, and goal-setting), a separate set of 4 agents required to step into a state in order to ascertain its fitness, then return to the initial state if the new fitness was lower (a.k.a. 2-step agents) were developed. Since their qualitative behavior was quite similar to that of the original agents, only one summary of their behavior (Table 4 in the body of the paper) is included.

Other relevant design issues are described as follows:



- Agents always move by randomly selecting values from a list of available states (i.e., adjacent states that have not yet been visited).
  - Expert-guided agents always exhaust all expert recommendations before trying available states that have not been recommended.
  - Imitator agents always try to match the behavior of the nearest higher fitness agent before trying other available states.
  - Goal-setting agents always move to a state that moves them one attribute closer to the selected goal. The attribute changed is selected randomly from the list of attributes that differ between the goal and the agent. In the event no higher fitness goal is visible, it reverts to control agent behavior.
- The only circumstance under which any agent will stop testing is when all adjacent states have been tested.

## Population

Populations of agents were assumed to be homogeneous and the number of agents was controlled by parameter. Other key assumptions:

- Agents were randomly placed on the landscape, with attributes determined by the bits in their numeric position (see description of landscape assumptions).
- Regression analysis used to develop expert recommendations was performed using an Interop call to the LINEST function in Microsoft Excel. Coefficients with a significance value  $< 0.05$  were identified, as were attributes that were the same across the landscape—which starts happening quickly in decomposable landscapes ( $K=0$ ).
- For the imitator and goal-setting populations, an insertion sort—ordered by diminishing fitness—was used to identify suitable neighbors.
  - For goal-setting agents, the first value higher fitness encountered that was within the visibility range was established as the goal.
  - For imitator agents, the search of the sorted list continued until the closest higher fitness neighbor was determined (i.e., an agent with a distance of 1 differing attribute would be chosen in preference to a higher fitness agent with 2 differing attributes).
- The simulation halts when all agents in the population are located on peaks—even if some individual agents have not yet verified that the state they occupy is a peak. This could potentially lead to a number of cycling steps near the end of the 2-step run (with some agents testing off peak while others return to their peak). That is why the base case used agents who remained in place while testing, moving only when higher fitness was encountered.

## Interface

The interface of the simulation makes heavy use of the .NET DataGridView control supplied with Visual C#. A variety of different types of runs can be made with the software:

- *Landscape features*: Allows values such as number of peaks to be determined for simulated landscapes of a particular type (NK or cluster) and NK combination. These can be averaged over a specified number of runs.

## Homophily vs. Expertise

- *Single pass steps*: Performs a single run with the specified type of landscape and agent, performing a regression analysis on each step (see Figure A3.1)
- *Single agent type runs*: Performs a specified number of trials on a particular landscape and agent, providing the results of each run and the average values.
- *Key variables*: Performs a specified number of trials on all four agents, using the same specified landscape (randomly generated for each run) across agents.
- *Full simulation*: Performs a “key variables” run starting with the base case (specified values for N, K, number of runs and number agents) then varying:
  - K from 0 to N-1
  - Visibility from 1 to N
  - Number of agents—testing the original number divided by two and times two.

The results of these simulations could then be pasted into MS-Excel spreadsheets to produce the graphs and tables used in the paper.

The screenshot shows the 'Complex Landscape Simulator' window. The 'Single Pass Steps' tab is active. The parameters are set to Agent: Control, Landscape: NK, N: 10, K: 5, Agents: 100, Passes: 100, and Visibility: 2. The main window displays a table of simulation results for 13 steps.

Run	Actual Peaks	Estimated ...							
	44	0							
Step	AverageFit...	Percent on ...	Percent Pe...	Percent Pe...	Max Peak ...	Significant ...	Significant a...	R-Square	
1	0.6405070...	0.1	0.2272727...	0.6	False	3	3	0.2511018...	
2	0.6741074...	0.12	0.2727272...	0.5833333...	False	2	2	0.2572538...	
3	0.7122101...	0.17	0.3181818...	0.6470588...	True	3	3	0.2288488...	
4	0.7353373...	0.21	0.3636363...	0.7142857...	True	2	2	0.2403318...	
5	0.7596462...	0.27	0.4318181...	0.7407407...	True	1	1	0.2079719...	
6	0.7838546...	0.34	0.5	0.7352941...	True	3	3	0.2380242...	
7	0.7959489...	0.38	0.5454545...	0.7368421...	True	4	4	0.3255485...	
8	0.8082992...	0.44	0.5909090...	0.7045454...	True	5	5	0.3430188...	
9	0.8215327...	0.48	0.6136363...	0.6875	True	4	4	0.3240142...	
10	0.8311100...	0.56	0.6363636...	0.7321428...	True	6	6	0.3719550...	
11	0.8424992...	0.61	0.6818181...	0.734375	True	5	5	0.3634255...	
12	0.8523113...	0.72	0.7045454...	0.7222222...	True	5	5	0.4069211...	
13	0.8587698...	0.79	0.7045454...	0.6962025...	True	5	5	0.3859317...	

**Figure A3.1: Interface for the simulation software used in the paper. Results displayed are for a single pass step simulation on an NK landscape for N=10, K=6 and 100 agents.**

## 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 Information Technology Education*.