Two Faces of Search: Alternative Generation and Alternative Evaluation

Thorbjørn Knudsen* and Daniel A. Levinthal**

At its core, a behavioral theory of choice has two fundamental attributes that distinguish it from traditional economic models of decision-making. One attribute is that choice sets are not available ex-ante to actors but must be constructed. This notion is well established in our models of learning and adaptation. The second fundamental postulate is that the evaluation of alternatives is likely to be imperfect. Despite the enshrinement of the notion of bounded rationality in the organizations literature, this second postulate in fact has been largely ignored in our formal models of learning and adaptation. We develop a structure with which to capture the imperfect evaluation of alternatives at the individual level and then explore the implications of alternative organizational structures, comprising such individual actors, on organizational decision-making.

*University of Southern Denmark
**Wharton School, University of Pennsylvania, Philadelphia PA 19104.

A central building block of the behavioral theory of the firm is the notion of bounded rationality (Simon, 1955). In contrast to the optimizing agent of neo-classical economics, Simon offered the satisficing decision-maker. The set of alternative actions are not presumed to be laid out in their entirety ex-ante, but must be discovered or searched. This facet of the behavioral theory of the firm (March and Simon, 1958; Cyert and March, 1963; Nelson and Winter, 1982) is by now well-established. However, another critical facet of bounded rationality has been largely ignored in this tradition and that is, how alternatives, once identified, are to be evaluated.

Simon suggested that, rather than optimizing a utility function, individuals search for alternatives until they identify one that satisfies some minimum performance criteria --- i.e,. in his words, individuals engage in satisficing behavior. Central to this perspective is the view that choice alternatives are considered in a sequential manner and that the process of the sequential evaluation of alternatives stops well short of some latent optimal possible option. What is less salient, though considered in the original discussion, is how actors are to evaluate the proposed solutions or alternatives. How is an actor to know whether a given alternative in fact "satisfices" or not? Simon notes that there may be uncertainty as to whether a particular alternative yields a state of nature that is in the satisfactory set or not, but the text suggests that this indeterminacy may be resolved by identifying a new alternative that does not suffer this risk. Yet, this discussion points to an important lacuna in both this early and subsequent development of behavioral theories of individuals and firms.[1]

[1] In contrast, the question of what constitutes an appropriate threshold, or aspiration level, has received considerable attention in the literature (cf., Greve, 2003; Lant, 1992). This boundary of what constitutes satisfactory or non-satisfactory performance is defined by a comparison of current performance to prior

While ideas of search are central in behavioral theories of the firm (March and Simon, 1958: Cyert and March, 1963), the mechanisms by which these alternatives are evaluated are less clearly developed (Gavetti and Levinthal, 2000). Models of adaptive search generally have the following characteristics. Some space of possible alternatives is sampled. The realization from this "draw" is then compared either to the current status quo action or in other cases to an aspiration level (e.g., Levinthal and March, 1981; Nelson and Winter, 1982; Lant and Mezias, 1990). When the space of alternatives constitutes attributes such as prices (Stigler, 1961; Nelson, 1970), assuming that quality attributes are well-defined and equal, then the model does not seem to require any elaboration. However, consider other possible spaces of alternatives, such as the space of possible spouses or the set of possible new production technologies for a factory. When presented with a new alternative from one of these sorts of "spaces", how is one to recognize a satisfactory solution when one is confronted with one? When is one to stop searching for alternatives? Despite the neglect of uncertain evaluation in our models of adaptive search, it is an important characteristic of many task environments.

Another important gap in our formal models of search is that the work tends to be remarkably non-organizational, though see March (1991), Lin and Carley (1997), Seshadri and Shapira (2003), and Rivkin and Siggelkow (2003) for important exceptions. While the label of organizations may be invoked, often times the formal structure corresponds to a model of individual-level problem solving. We draw on Christensen and Knudsen's (2004) recent extension of the Sah and Stiglitz (1986) characterization of

---

performance, as well as to the performance of others who are viewed by the focal actor as belonging to his or hers reference group.

organizational architectures to provide a framework with which we can consider the impact of alternative organizational structures on search processes.

At a basic level, evaluation of alternatives can suffer from two possible errors: Type I errors of rejecting a superior alternative and Type II errors of accepting an inferior alternative. As shown on the work on economic architecture (Sah and Stiglitz, 1986; Christensen and Knudsen, 2004), different organizational structures vary in their proclivity to make one type of error or the other. In particular, hierarchical structures, in which a proposal needs to be validated by successive ranks of the hierarchy in order to be approved, will tend to reduce the likelihood that an inferior alternative will be adopted --- i.e., hierarchy reduces errors of Type II. In contrast, what Sah and Stiglitz (1986) term polyarchies, a flat organizational structure in which approval by any one actor in a series of decision-makers is sufficient for an alternative to be approved, will tend to minimize the probability of rejecting a superior alternative --- i.e., polyarchy reduces errors of Type I. Christensen and Knudsen (2004) provide a general graph theoretic treatment of these structures that allows one to consider the full range of organizational architectures that range between these two extreme forms and thereby allow one to specify structures that trade-off these two types of errors as the relative degree of hierarchy and polyarchy shifts and, further, to examine the change in the overall reliability of the organizational structure as the number of actors within the organization changes.

Using this analytical platform, we examine how alternative organizations of evaluators would move on a space of possible alternatives. In particular, we use the structure of fitness landscapes (Wright, 1931; Kaufman, 1993) to characterize a sense in which alternatives are more or less proximate to one another.[2] As characterized by Levinthal (1997), a process of local search is modeled as examining, at random, one of the adjacent points in the space of alternatives. The value of points in adjacent locations in fitness landscapes, as developed by Kaufman (1993), are correlated, with the degree of correlation being "tuned" by the intensity of the interdependencies among the N attributes that contribute to the fitness of a given alternative. Changing the level of interdependencies also impacts the overall structure of the landscape in that the number of local peaks increases with the degree of interdependencies (Kaufman, 1993). The presence of local peaks poses particular challenges to a process of local search, as a decision maker at a local peak will be unable to identify superior alternatives that may be present on the broader landscape.

While the structure of fitness landscapes has been used to consider the issue of organizational adaptation (cf., Levinthal, 1997; Rivkin, 2000), as with much of the broader literature on search processes, the issue of how alternatives are to be evaluated has been under-developed. There have been some recent studies (cf., Rivkin and Siggelkow, 2003; Dosi et al., 2003; Ethiraj and Levinthal, 2003) that examine how the allocation of decisions across actors within an organization and the degree to which actors make decisions based on the parochial concerns of their local sub-unit or the payoff to the broader organization effects the adaptive capabilities of an organization. However, in these analyses, there is no uncertainty as to the payoff implications of the choices being made; rather, decision-making processes are impacted by the perspective (local versus global; one sub-unit versus another) taken by the actors. Closer in spirit to

[2] While it is common to refer to the value ascribed to the phenotype as fitness in this structure, the value may better be thought of as a kind of technical performance measure, as opposed to corresponding to a measure of birth and death rates as in the standard biological use of the term fitness.

our effort is Gavetti and Levinthal (2000), who contrast evaluation that is “off-line”, in which assessment is done of the basis of actors’ cognitive model of the fitness landscape, and “on-line”, in which the evaluation occurs subsequent to experience with the actual alternative.

We try to incorporate both “faces” of a search process: the sequential evaluation of alternatives and the uncertain evaluation of those alternatives that are identified. Whilst the first topic has been studied extensively in previous models of adaptive search, the second topic of uncertain evaluation has been rather neglected. We find that highly accurate evaluation of alternatives results in search processes being greatly influenced by the happenstance of the order in which alternatives are identified by the actor. As a consequence, highly accurate evaluators become trapped by their random starting position in the landscape of alternative actions. Perfect evaluation leads to the rapid identification of a local peak and the persistence in that location across time. In contrast, imperfect evaluation leads to a more robust search process that is not as influenced by the happenstance of ones starting position in this landscape. Furthermore, we find that populations of moderately imperfect evaluators pay a surprisingly modest penalty in terms of the variability in their performance, either cross-sectionally or across time.

We also consider, albeit in a rather stylized manner, how the structure of organizational evaluation of alternatives impacts these dynamics. Organizations that are hierarchical in structure, even if composed of imperfect evaluators, tend to replicate the conservatism of perfect evaluators and become trapped by local peaks. Hybrid forms, consisting of a mixture of polyarchy and hierarchy, effectively balance the dual imperatives of exploration and exploitation (Holland, 1975; March, 1991). Finally, the less able (or, conversely, the more able) are individual evaluators, the more attractive are organizational forms that tend towards hierarchy (polyarchy) as the hierarchical structure tends to compensate for the high error rates of less able individual evaluators (or, conversely, the variance induced by the polyarchy forms tends to compensate for the overly precise judgments of more able evaluators).

Model Structure

The model structure has three basic elements: the characterization of individual evaluation of alternatives, how individual evaluators are aggregated into an organizational form, and the specification of the task environment or the space of alternatives.

Individual Evaluation of Alternatives

Individual agents are characterized as being able to distinguish between a proposed action alternative and the status quo with more or less reliability. A perfect evaluator would, with certainty, distinguish between inferior and superior alternatives no matter how small the value differences are among two proposals. However, decision makers are unlikely to conform to such high standards. Actors are likely to make errors in identifying which, among a pair of alternatives are, in fact, superior. However, one would expect that the likelihood of making a false classification is a decreasing function of the actual differences in value between the alternatives. That is, one may frequently mis-classify pairs of alternatives that vary in payoff by only a small amount. In contrast, if the payoff to the two alternatives is substantially different, then the probability of making a mis-classification, while not zero, would certainly be less than in the former case.

These properties are reflected in the screening functions represented in Figure 1. The horizontal axis indicates the actual difference in payoffs between a currently held alternative and a proposed alternative (current fitness minus new fitness), ranging from large negative differences in value to large positive values. The vertical axis indicates the probability that an agent would accept the proposed alternative. Obviously, an intelligent screening function should have an upward slope such that superior alternatives are more likely to be accepted than inferior alternatives. In the extreme, with a perfect evaluator, the curve would have a point of discontinuity at zero, such that proposed alternatives with a payoff less than the current alternative (yielding a negative fitness difference) would be rejected with probability 1 and those with higher payoff (positive fitness difference) accepted with certainty. We specify a family of linear screening functions where the slope of the line, indicated by the variable α, can be interpreted as the screening capability of the agent. A steeper slope, or higher value of α, implies that the probability of accepting a proposal is more sensitive to changes in its actual value. Within the class of linear screening functions, we restrict our attention to those that are unbiased in that if there is no difference in payoff between the proposed alternative and the current action (the value of fitness difference is zero), then the actor is equally likely to accept or reject the proposed alternative.[3] As α becomes arbitrary large ($\alpha \to \infty$), the screening function approximates that of a perfect evaluator.

---

[3] One might imagine that actors have a "status-quo" bias in which case when faced with a new alternative that yields the same payoff as the current alternative they would be inclined to reject the proposed alternative. We have analyzed screening functions with this property, essentially such a "bias" simply shifts the y-intercept of these curves and generates qualitatively similar results to the ones provided here.

Insert Figure 1 here

Aggregation of Individual Evaluators Into an Organizational Form

Individual agents can be aggregated into more complex organizational forms. In particular, organizations can be characterized by the number of agents within them, but also more subtly by the nature of decision authority within them. Following Sah and Stiglitz (1986), we focus on whether or not a given actor has the authority to approve or reject a proposed alternative, or is merely authorized to pass the proposed initiative along within a broader chain of command. In particular, consider Figure 2 which represents the flow of decisions in six distinct organizational forms. In what we term a hierarchy, a proposal is initially considered by the agent at the far left in the figure. If the proposal is rejected by that agent, it is eliminated from further consideration (indicated in the figure by the dashed vertical line from that decision node). Alternatively, if the proposal is approved by that individual, then it is passed rightward to the next individual in the chain of command. A proposal is acted upon only if it is positively screened by all six agents. The agent to the far right effectively sitting at the top of this hierarchy, only views proposals that have been successfully vetted by the lower-level actors and has final say as to whether the organization adopts those proposed initiatives that reach his or her attention. In a hierarchical evaluation of a stack of job applications, for example, proposals are eliminated at each level in the hierarchy and a diminished stack moved up

---

We examine the no-bias condition simply to eliminate the need to introduce another parameter in the subsequent analysis.

to the next level. The agent at the top of the hierarchy will only see a very small stack of proposals.

Insert Figure 2 here

At the other extreme is the polyarchy structure. Here, proposed alternatives can be adopted by any of the six decision makers; an alternative is only dismissed if all decision makers in succession reject it. It is this contrast between the conservative (rejection oriented) hierarchical structure and the pro-acceptance oriented polyarchy structure that Sah and Stiglitz (1986) consider. Christensen and Knudsen (2004) develop an analytical structure that allows them to consider a wide range of hybrid forms that lie intermediate to these two extremes.[4] Figure 2 illustrates 4 intermediate forms that range from nearly hierarchical (Hybrid 1) to nearly polyarchy (Hybrids 3 and 4).

In analyzing the role of alternative organizational forms, we wish to distinguish between the effect of individual differences in screening ability and the impact of the structure of the relationship among evaluators within the organization. Therefore, we treat organizations as being homogenous in the screening ability of the individual agents that comprise the organization, though we examine the impact of varying this homogenous level.

Figure 3 indicates the effective screening properties of six alternative organizational forms, all comprised of six agents with a α value of 10. Using methods

[4] The critical analytical challenge is to specify the implied organizational screening function that results from a set of individuals of a given ability (or individual screening function as in Figure 1) that are organized in a particular structure (as suggested by Figure 2). Christensen and Knudsen (2004) derive this mathematical relationship.

outlined in Christensen & Knudsen (2004), we derived an organizational level screening function, *F*, which is a mathematical representation of the flow of decisions in an organizational form (as shown in Figure 2). In order to examine the effect of changing organizational structure, we assume that all members in an organization have identical ability. That is, we assume that the individual level screening function, $f(x)$, the probability that an individual accepts an alternative, is the same for all members of an organization. In this case, the organizational level screening function, *F*, is a polynomial in the individual level screening function, $f(x)$. For example, accepting an alternative in the six-member hierarchy requires that all of its six members accept the alternative. Therefore, the organizational level screening function of the hierarchy (shown in Figure 3) is given by $F= f(x)^6$, which is the probability that this structure accepts the alternative in question. In a similar way, it is easy to see that the organizational level screening function of the polyarchy (shown in Figure 3) is given by $F= 1-(1-f(x))^6$, i.e., the probability that at least one out of the six polyarchy members accepts an alternative. The screening functions of the four hybrids, shown in Appendix 1, were derived in a similar way. As a point of reference, in Figure 3 we also include for comparison the evaluation function of a single perfect evaluator.[5]

We see that the screening function of the polyarchy lies everywhere above the screening functions of alternative forms. This implies that the polyarchy for any given fitness difference is, as we suggested above, more prone to accept alternatives, even those with a negative value --- a Type II error. That is, the polyarchy reduces Type I error at the expense of increasing Type II error. Conversely, hierarchies are very unlikely to

[5] The issue of organizational form is not relevant in the case of perfect evaluators as each perfect evaluator in the organization would simply replicate the decision of others.

mistakenly accept an inferior alternative,(a Type II error) with the probability of acceptance being near zero for alternatives with a negative fitness difference. However, that same caution causes the hierarchical structure to reject many superior alternatives, i.e., alternatives with a positive fitness difference --- a Type I error. Hierarchy reduces Type II error and increases Type I error. Interestingly, the screening function of hybrid forms will trade-off the effects of polyarchy and hierarchy: hybrids both reduce Type II and Type I error (in particular, this effect is apparent in hybrids 2 and 3).

Insert Figure 3 here

Specification of the Task Environment

The final basic element of the model structure concerns the task environment in which agents (and organizations) operate. Prior work examining the effect of alternative organizational forms on the effectiveness of screening alternatives treats the process of alternative generation as being purely random draws from some fixed distribution of possibilities (Sah and Stiglitz, 1986; Christensen and Knudsen, 2004). However, research on organizational search processes (March and Simon, 1958; Cyert and March, 1967) has emphasized the spatial location of the set of possible alternatives, with the notions of local and distant search being central in theoretical (March and Simon, 1958; Nelson and Winter, 1982) and empirical analyses (Podolny and Stuart, 1996; Rosenkopf and Nerkar, 2001). The imagery of spatial location is given clear expression in work on search in fitness landscapes (Levinthal, 1997).

The generation of alternatives is not purely random, but is likely to reflect the availability of options in the neighborhood of the organization's current practices. Building on Levinthal (1997) and related work (Rivkin, 2000; Rivkin and Siggelkow, 2003; Dosi et al., 2003) the task environment of fitness landscapes (Kauffman, 1993) is used to characterize a space of alternatives, where alternatives vary along any one of N dimensions and the correlation among distinct alternatives can be "tuned" by manipulating how interdependent these N elements are in determining the overall payoff.

If attributes of a policy contribute to performance in a relatively independent manner, the landscape of policy alternatives is relatively smooth, changing one attribute of the policy only affects the performance contribution of that attribute in isolation. In contrast, if policy attributes have a high degree of interdependence, then changing even just one attribute may have broader repercussions and affect the performance contributions of other attributes. As a result, a "landscape" of alternatives in which there is a high degree of interdependence will exhibit a relatively low-level of correlation, with even modest shifts in attributes leading to a pronounced change in overall value. Related to this issue of degree of correlation among neighboring alternatives is the number of peaks in this performance landscape. With no interdependence among the policy attributes, there is a single peak in the landscape corresponding to the optimal setting of each of the individual policy attributes. As interdependence increases, the performance surface will exhibit local peaks of configurations of policy attributes that exhibit some degree of internal consistency (Kauffman, 1993).

More formally, we specify alternatives as consisting of N attributes, a_1, …, a_N. For simplicity, it is assumed that each attribute can take on two states. A performance landscape is a mapping of any possible vector of firm choices $A = (a_1, a_2, \ldots, a_N)$ to performance values V(A). We create performance landscapes with a variant of the NK-model (Kauffman, 1993, see Sorenson 2002 for a review of these models in the organizations literature). The value of each individual attribute a_i is affected by both the state of that attribute itself and the states of a number of other attributes $\mathbf{a_{-i}}$. Denote the value of attribute a_i by $c_i(a_i, \mathbf{a_{-i}})$. For each landscape, the particular value of an attribute, c_i, is determined by drawing randomly from a uniform distribution over the unit interval, i.e., $c_i(a_i, \mathbf{a_{-i}}) \sim u[0, 1]$. The value of a given set of alternatives A is then given by:

$$V = [c_1(a_1, \mathbf{a_{-1}}) + c_2(a_2, \mathbf{a_{-2}}) + \ldots + c_N(a_N, a_{\mathbf{-N}})].$$

The identity of $\mathbf{a_{-i}}$, i.e., the set of alternatives that affect each attribute a_i, is given by the interaction structure of the firm's decision problem (i.e., the variable K).

Unless indicated otherwise, our results reflect the average of 100 entities searching on each of a 100 distinct landscapes. Each of these landscapes has the same structure in terms of K, the degree of interdependence among attributes in contributing to performance, but represents a distinct draw on the common underlying probability generating structure. To enhance the comparison across these families of landscapes, we normalize the performance level on each surface so that average performance equals 0.5 and maximum performance equals 1.

Analysis

To provide some initial understanding of the nature of the adaptive search process we model here, we first consider the behavior of individual agents and then, in the subsequent analysis, model the behavior of alternative organizational structures.

Individual Evaluators

Figure 4 indicates the performance of two types of agents who vary according to the accuracy of their evaluation function. For the sake of a base-line comparison, we model one agent as being a perfect evaluator; in this setting, only alternatives that enhance the actual payoff will be accepted. In contrast, the other agent ($\alpha = 10$) exhibits some intelligence in evaluation (i.e., $\alpha > 0$), with the probability of accepting a more favorable alternative increasing as a linear function of the performance increases associated with that alternative; however, this agent will at times mistakenly accept alternatives that in fact offer inferior performance and in other instances reject alternatives that could enhance the organization's performance (i.e., α is finite). We see that the perfect evaluator quickly asymptotes in the performance that is achieved, while the imperfect evaluator not only outperforms the perfect evaluator, but, if additional periods are examined, continues to exhibit modest but steady performance improvement. Perfect evaluation leads to the rapid identification of a local peak and the perfect evaluation function will lead the actor to maintain that position for the remainder of the simulation, while imperfect evaluation leads to persistence in search.

Insert Figure 4 here

We would expect, however, that imperfect evaluation would suffer from two possible downsides. First, it is natural to expect that an imperfect evaluator would experience a slower rate of ascent in initial performance gains as an imperfect evaluator, by definition, will at times make downward moves.[6] Even though this conjecture was confirmed, the difference in the initial rate of progress between the imperfect and the perfect evaluator is extremely slight in the comparison shown in Figure 4. However, around period 40, we start to see a divergence in the two performance curves as the performance of the imperfect evaluator continues on an upward gradient while that of the perfect evaluator begins to asymptote.

The other "penalty" that imperfect evaluation might exhibit is with respect to a limited ability to maintain, over extended periods of time, the attractive alternatives that have been identified. Given the noise in his or her evaluation process, even if a global peak is identified, there is a chance of mistakenly being seduced off of it by an alternative that appears superior. Figure 5 illustrates the emergence of the distribution of agents across the performance landscape, where the 1024 distinct locations in the landscape are ranked ordered from 1 (the global peak) to the lowest value (1024).[7] In the initial period, locations are randomly arrayed and hence the distribution of agents across locations is quite flat. Rapidly, we see the emergence of clusters of agents aggregating on specific locations in the landscape. We see a particular massing of agents on the global peak, although given the imperfect evaluations, there is some dispersion of agents around this peak --- a "cloud" of agents as if it were. In contrast, in Figure 6 with perfect evaluations, we see greater cross-sectional dispersion among agents as "columns" of agents aligned on distinct local peaks.

---

[6] This downward movement is masked by the results in Figure 3 as the figure provides the results for the average performance over a set of runs (10,000 agents: 100 distinct landscapes with 100 agents on each). Examination of individual runs does reveal instances of such downward "walks".

[7] In Figures 5 - 8, we analyze the behavior over a single, randomly specified performance landscape (10,000 agents). This allows us to identify a fixed population of 2^N alternatives, or 1024 given N has a value of 10 in our analysis. These figures were generated by assigning the average number of agents at each of the 1024 locations for each of the 250 time steps.

Insert Figures 5 and 6 here

Figures 7 and 8 provide more direct evidence regarding the ability of the two populations of agents to identify relatively attractive locations in the performance landscape. We see that the histogram of agents' locations across different performance levels is much more dispersed in the case of imperfect evaluators. Imperfect evaluators tend to mass at locals with the highest level of performance, but occasionally visit discrete locales associated with rather low levels of performance. While imperfect evaluators tend to mass amongst the highest performing locations, there are fewer perfect evaluators at the very highest performing locations. Perfect evaluators tend to spread out evenly at intermediate performing locations.

Insert Figures 7 and 8 here

Building on this analysis, we can convey a more literal sense of the imagery of "clouds" of agents forming more or less tighter clusters of movement around different performance peaks in the landscape. The histograms in Figures 7 and 8 indicate the number of agents in the two populations that mass at different locales, but these figures do not convey a sense of the instability or turbulence in the populations of imperfect evaluators. Imperfect evaluators are, on average, finding attractive locations in the space of alternatives, but are they wandering at some, possibly high, rate of velocity on the performance landscape?

Figure 9 creates a panel of images that conveys a sense of the dispersion of agents across locations in terms of the size of the "clouds" of agents massing at a location and the flux in these constellations of agents. We measure flux, or turbulence, by the ratio of the number of different agents who visit a particular point in a given period of time (the last 10% of the run in this analysis) to the average number of agents at that location during that same time interval. In Figure 9, this measure of turbulence is indicated by color with a darker color indicating a more turbulent setting. As the population of agents "cools" down and this ratio approaches one, we use an increasingly lighter color. The size of the circle represents the number of agents at that location.

Highly imperfect evaluators (α equal to 1), lead to both a very diffuse population (the circles are numerous in number and tend to be of modest size) and a very turbulent structure with many different agents visiting a given locale (the clouds are very dark). Moderately able evaluators (α values of 3, 5, and 10) result in agents clustering on superior locations, with the agents forming a few large constellations around the superior points in the performance landscape. Clearly, the degree of turbulence diminishes as the agents' evaluation ability increases, with the color of the "clouds" shifting from dark to light. However, we see that with populations of agents who are highly accurate evaluators (α equal to 50 and a population of perfect evaluators) the clouds tend to be "frozen" in that there is little (in the case of α equal 50) or no (in the case of perfect evaluators) turbulence, but that such "freezing" tends to result in a number of distinct constellations of agents, many of which are not associated with particularly attractive points in the performance landscape.

Thus, imperfect evaluators appear not to wander too far off from the attractive peaks that they identify. That is, the "clouds" of agents around the peaks stay tightly clustered. We do not often see a situation in which slightly inferior alternatives are adopted and, then from this new lower base, even more inferior alternatives are mistakenly adopted. It is certainly possible for agents to take such a two-step "walk" from an attractive peak and on occasion they will do so. However, the fact that the screening process, while imperfect, is nonetheless intelligent, in that more favorable alternatives are more likely to be accepted than less favorable ones, implies that mistakes, walks away from superior alternatives, will tend to be self-correcting. After accepting an inferior alternative that takes him or her away from an attractive peak, it is more likely that the subsequent move will be back to this same peak rather than a move that takes the agent even further away from this location. Ironically, the agents with an evaluation mechanism more prone to error exhibit less cross-sectional variability than the population of perfect evaluators. While it is true that perfect evaluators will exhibit greater reliability in terms of period to period location and performance differences, we see from Figure 9 that the degree of instability in the behavior of the imperfect evaluator is rather modest. Furthermore, this population convergence among the imperfect evaluators occurs at the highest performing locations in the landscape.

Robustness

This result that the performance of imperfect evaluators can exceed that of a perfect evaluator is not a "knife-edge" property of the model. The critical factor regarding the robustness of the analysis relates to the structure of the task environment.

We have examined in this baseline analysis a landscape with a moderate degree of interdependence ($K = 3$). In the limit, with no interdependencies (i.e., $K = 0$), the landscape would have just one peak corresponding to the globally maximum fitness level. In such a setting, the fact that perfect evaluators reliably identify their local peak and sustain their position on this peak over time can only serve to enhance performance. However, as long as K takes on a value of 1 or more, it is possible to identify an imperfect evaluation screening value that results in superior performance.

Figure 10 displays the performance level reached in the final period for different screening abilities (α values) in landscapes that vary in their degree of ruggedness (i.e., their K value).[8] Not surprisingly, shifting from a zero intelligence screener (α value of zero) to evaluators with some ability to discern superior from inferior alternatives (α values of one or more) leads to a marked improvement in performance. However, more surprisingly, as long as $K > 0$, we find that as the precision of the screeners is increased beyond a threshold level, the performance level that is achieved begins to decline. This threshold is realized at a lower value of α in more complex environments; it is in more complex, multi-peak landscapes that the enhanced tendency for imperfect evaluators to search is most valuable.

Insert Figure 10 here

This demarcation of the space defined by α, the precision of performance evaluation, and K, the complexity of the task environment, in terms of their implication

[8] The values of $K=3$ and $\alpha = 10$ and perfect evaluation are indicated in this plot as these two points correspond to the values used in our prior baseline analysis.

for performance is further illustrated by the level curves plotted in Figure 11. Each point of a given curve yields the same performance. We see that for all values of α of 3 or greater, the level curves are backward bending, in some cases markedly so, implying that more precise screeners can only obtain the same level of performance achieved by less able screeners if they are operating in a less complex (lower K) environment. Or viewed differently, for a given fixed value of K (indicated by a vertical line), the maximal level curve is obtained by a decidedly more imperfect evaluator.

Insert Figure 11 here

Organizations of Evaluators

Search is not merely carried out by individuals in isolation, but such individual evaluation is typically embedded in a larger organizational context. One actor may endorse an initiative and pass it along to another, perhaps hierarchical superior, for approval. Other actors may have sufficient authority to endorse or terminate an initiative on their own. We characterize an organization as consisting of a set of individuals who vary with respect to their authority to terminate proposals (terminating by their own evaluation or recommending termination to others) and endorsing proposals (authorizing the proposal on the basis of their own assessment or recommending acceptance to others).

As previously noted, we can characterize two extreme forms of organizational architectures: hierarchy and polyarchy. Hierarchy requires that for an alternative to be accepted it must pass through an approval process at each level of the organization. In this sense, hierarchy is very conservative and is unlikely to make Type II errors of accepting inferior alternatives. In contrast, polyarchy is defined as a structure in which approval by any agent within the organization is sufficient for acceptance of an alternative. As a result, the polyarchy structure tends to be very “pro” innovation and change and tends not to make errors of Type I of rejecting superior alternatives. Only one actor in the organization needs to see merit in the initiative for it to be adopted; however, this same property makes polyarchies prone to making Type II errors of adopting inferior alternatives.

We model organizations of a fixed number of actors, 6, that include the two pure forms of hierarchy and polyarchy, as well as the four forms intermediate to them (see Figure 2 and Figure 3). Obviously, organizational structure is an interesting property for a population of imperfect evaluators while, in contrast, with perfect evaluation the decision outcome would be invariant to structure. Thus, we take the same imperfect evaluation function previously examined in the individual actor analysis and examine organizations of 6 such actors arrayed according to the 6 alternative organizational forms.[9]

What is the effect of organizational architecture on search processes? We see that the hierarchical form has many of the properties of the perfect evaluator. Such organizations tend only to “walk” up-hill, all be it slowly, as they tend only to accept new alternatives that do in fact lead to an increase in performance. Thus, as with perfect evaluators, they tend to be “prisoners” of their starting positions, identifying local peaks

[9] Clearly, size is another facet of organizational form that could be varied. However, the set of possible hybrid forms grows exponentially with organizational size. Indeed, Christensen and Knudsen (2004) show that with a sufficient number of agents, an organizational form can be specified that approximates arbitrarily closely a perfect screening function --- i.e., a function that rejects inferior alternatives and accepts superior alternatives with certainty. However, note that per our analysis of perfect screening in the prior section, perfect screening need not be a desired property. To focus the attention on the role of changing organizational structure from more hierarchy ones to ones that are more polyarchy like, we keep the number of actors within the organization fixed.

but not exploring more broadly in the landscape. Further we find that our intermediate forms can offer an effective mix of exploration and exploitation (Holland, 1975; March, 1991). Elements of polyarchy enhance the breadth of search, but some degree of hierarchy facilitates the organization's ability to reliably sustain an attractive position in the landscape once identified. However, given that even a population of single evaluators is able to cluster rather closely to the most attractive peaks in the landscape, only a modest degree of hierarchy is needed to reliably sustain an attractive position in the performance landscape.

Reflecting these trade-offs between the search inducing polyarchy forms and the inertia generating hierarchical forms, we find an important complementarity between organizational form and screening ability of the agents who comprise the organization. Figure 12 examines the performance implications of the 6 organizational forms that we consider and contrasts this performance with the performance of a single agent, with the screening ability for both the single agent and the set of agents within the organization set at our baseline case of $\alpha = 10$. We find, in this setting, that the polyarchy yields the highest performance among the six organizational forms. However, it is also important to note that four out of the six organizational forms result in a lower performance than could be generate by an individual member of the organization (i.e., the imperfect agent shown with a dashed line).[10] Organizations have the potential to compensate for weaknesses of individual screeners (hierarchy potentially helping to reduce the extensiveness of search in the case of highly inaccurate screeners and polyarchy forms usefully enhancing the degree of search for more accurate screeners), but as indicated in

[10] The hierarchy and hybrids 1-3 reduces the performance of their members. Note also that the difference in performance between hybrid 1 and 2 is so slight that these two organizational forms cannot be visually distinguished in the comparison shown in Figure 12 (it seems that there are only 3 hybrids).

the results in Figure 12 the inappropriate organizational form may exacerbate the pathologies associated with an individual screener.

Insert Figure 12 here

Figure 13 examines this interrelationship between organizational form and screening ability across a range of screening levels. For very imperfect screening ability (α values of 1 and 2), we see that hierarchy yields a substantially higher level of performance relative to the polyarchy form. Hierarchy is a useful complement to very imperfect screeners. Individuals who evaluate alternatives with considerable noise naturally induce considerable breadth of search. Breadth of search has the virtuous quality of preventing organizations from locking in pre-maturely to inferior peaks. Sustained breadth of search, however, has the liability of generating a more dispersed “cloud”, or distribution, of organizations around the superior alternatives that come to be identified in the long-run. Thus, highly imperfect evaluators, even if placed in a hierarchical structure, are likely to generate broad ranging search, but the hierarchical form will enhance the ability of such organizations to retain the attractive solutions that are identified.

Insert Figure 13 here

Conversely, polyarchy is a desired complement to organizations composed of highly accurate screeners. Accurate screeners are likely to rapidly identify a local peak in the landscape. Polyarchy, a form that permits any individual within the organization to approve an alternative, only requires one of the 6 actors in the organization to view an alternative as favorable in order to result in its adoption. Thus, as long as the evaluation of the individual agents composing the organization has some possibility of error, polyarchy compounds the likelihood of accepting an alternative that results in an immediate decline in performance; at the same time, however, polyarchy offers the possibility of broadening search to new regions of the performance landscape. If actors are highly accurate in their individual screening, such organizations do not pay a significant price for their pro-acceptance bias of the polyarchy form in that, in the long-run, the distribution of organizations is still tightly packed around the superior peaks in the performance landscape.

Indeed, these results suggest that organizational forms must be designed to fit the contingencies of the available workforce (screening ability) as well as the task environment (level of uncertain evaluation and interdependencies among policy attributes). Thus, in the same task environment, the more able are the individual agents composing the organization, the more that organizational form should shift towards the permissiveness of the polyarchy form. Very able agents need a structure that accepts and empowers the divergent views of organizational members. Conversely, agents who are less able and therefore less discriminating require the repeated checks on behavior that hierarchy provides. Note that, as agents become near-perfect screeners, performance becomes insensitive to the specification of organizational form. In the limit, with perfect screeners, agents would simply replicate each others evaluation decision; thus, in the limit, performance is invariant to organizational form and the number of agents engaged in evaluation. Thus, a perfect evaluator would not benefit from being member of an organization.

Conclusion

Much of our analysis of search processes has been very one-sided. We recognize that choice sets are not presented to decision-makers but must be identified through search processes. However, in considering these important issues of discovery, we have tended to treat the problem of evaluation as to be trivial or self-evident. But the question of valuation is far from trivial and indeed forms the crux of resource allocation processes (Bower, 1970). If, as suggested by the "Carnegie School" view of the firm, organizations engage in problemistic search comparing small sets of alternatives to a status quo performance, then understanding the nature of that evaluation process is critical to understanding the dynamics of search processes and ultimately the pattern of firm adaptation that we observe.

While clearly a stylized and admittedly incomplete treatment of this question of evaluation, the work provides some useful redirection of the field's attention, as well as some initial results of interest. An unintended by-product of precise evaluation mechanism is the short-circuiting of the search process. Highly accurate evaluation systems will rapidly identify one of the local peaks in the neighborhood of the location at which the search process commenced. Evaluation processes under a rather wide range of imprecision will yield superior performance. It is perhaps not surprising that less precise evaluation would tend to result in search processes persisting for longer periods of time. However, the relatively low variability in the range of behavior and performance such populations experience is quite surprising. Indeed, in a cross-sectional sense, highly accurate evaluators generate greater variability in performance and behavior than do populations of less precise evaluators. Under moderate ranges of evaluation ability, we observe “clouds” of agents clustering rather tightly around the superior peaks in the performance landscape. In contrast, with highly precise evaluation, we observe distinct mass points of agents spread out evenly on a variety of local peaks that vary considerably in their performance value.

The precision of evaluation and the relative rates of Type I and Type II errors in the evaluation process are importantly affected by the structure of organizational decision processes. Hierarchical organizational forms, for a given screening ability of the agents who comprise it, tend to be cautious and are unlikely to mistakenly shift to less favorable alternatives. In contrast, flat forms that have a polyarchy quality tend to induce greater search as a result of the greater probability of making such errors in evaluation. Given that highly accurate screeners are likely to stop their search process prematurely, polyarchy as an organizational form is a useful complement to highly accurate screening ability. By the same token, elements of hierarchy can facilitate an organization of rather inaccurate screeners persisting on a favorable course of action, once identified.

This basic analytical structure that we have developed can be enriched and built upon in a number of ways. In our current analysis, the generation of alternatives is specified exogenously and is determined by the structure of the performance landscape. As Nelson (1982) argues, a better cognitive understanding of one’s task environment may allow for more intelligent identification of the alternatives to be sampled. We have treated the sampling process as defined by local search. While this is an important line of argument in the literature (from March and Simon (1958) onwards), it is important to consider the intelligent identification of non-local options (Gavetti and Levinthal, 2000).

A different form of endogeneity that would be interesting to consider is with respect to actor's screening ability. There is a vast literature on experiential learning (Argote, 1996) which suggests that skill at tasks increases with repeated trials. Therefore, it is reasonable to expect that screening ability may change with an actor's experience with a class of problems. Thus, actors may become quite skillful and accurate in evaluating one class of alternatives, but rather inaccurate in evaluating a different, and for them, novel set of alternatives. Experiential learning of this form should tend to acerbate the problem of competency traps previously identified in the literature (Levinthal and March, 1981; Levit and March, 1988). By becoming more expert evaluators in the domain of the organization's current activities, actors are less likely to engage in further search. Consistent with work on organizational demography and innovation and similar to March's (1991) model of exploration and exploitation, turnover in personnel and the introduction of more novice actors may be necessary to facilitate search processes.

While there are many such avenues of further inquiry, we wish to reiterate the basic call with which we started. Search is not merely about generating and discovering alternatives. It is equally about judging the value of those alternatives with which one is presented with. We hope to have, at a broad level, redirected the conversation in the organization's literature to a more balanced consideration of search processes, as well as provided a particular structure and set of results to facilitate such consideration.

Figure 1: Six levels of screening ability for an agent, ranging from completely random screening (α= 0) to perfect screening (α → ∞).

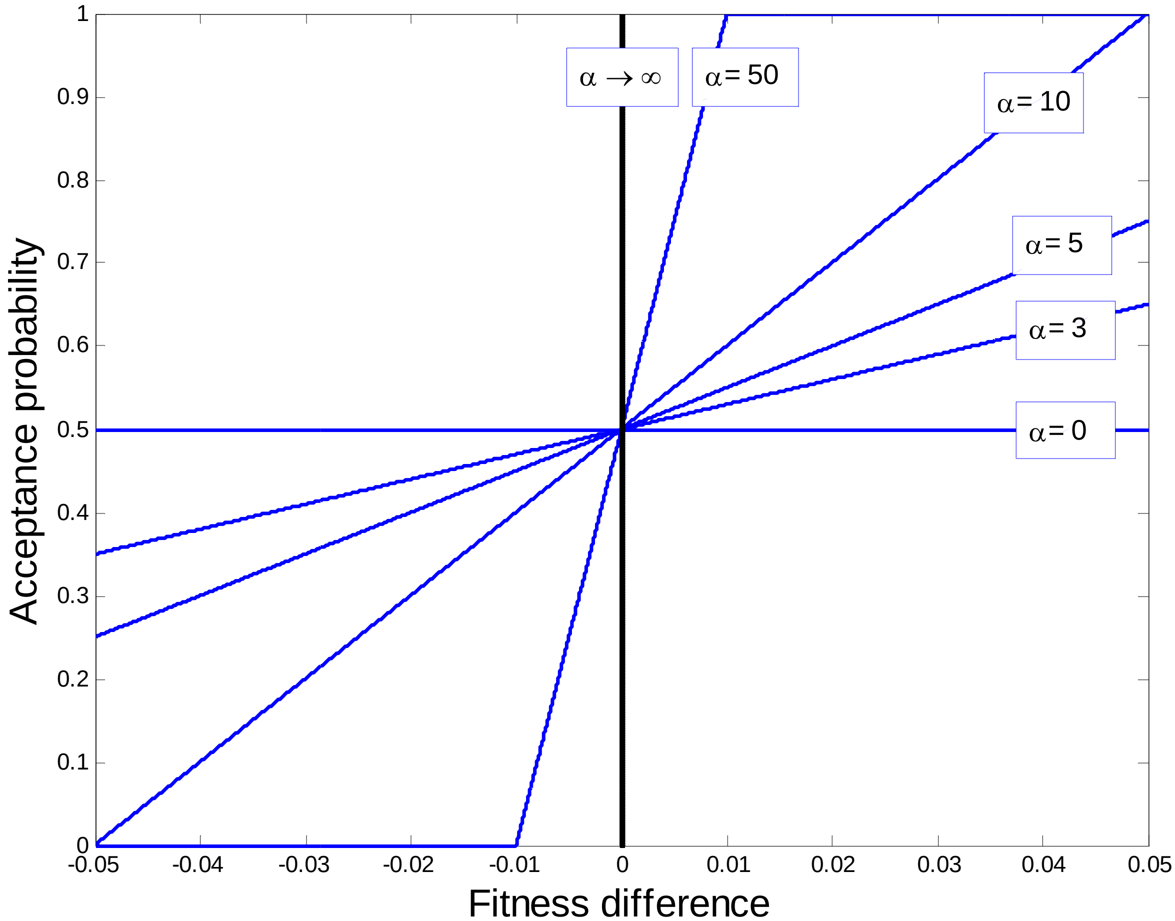

Figure 2: Flow of decisions in six organizational forms. In each of the six organizational forms, proposals enter with the actor in the lower left corner and then flow from the left towards the right. Dashed lines show rejection of proposals and solid lines show acceptance. Proposal that exit to the right are adopted by the organization (the number of organization members that may have a final say as to whether the organization adopts a proposal increases from one in the hierarchy to six in the the polyarchy).

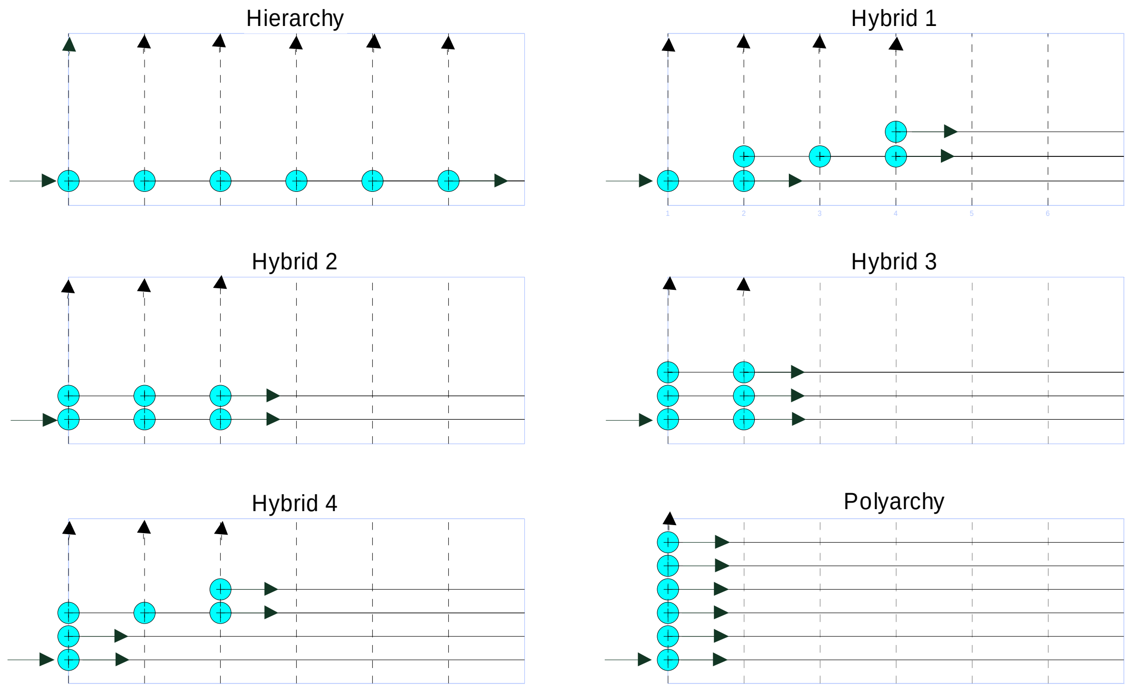

Figure 3: Ability of imperfect agent compared to six organizational forms, each built of six identical imperfect agents (α= 10). The ability of the organizational forms were derived according to the procedure shown in Appendix 1.

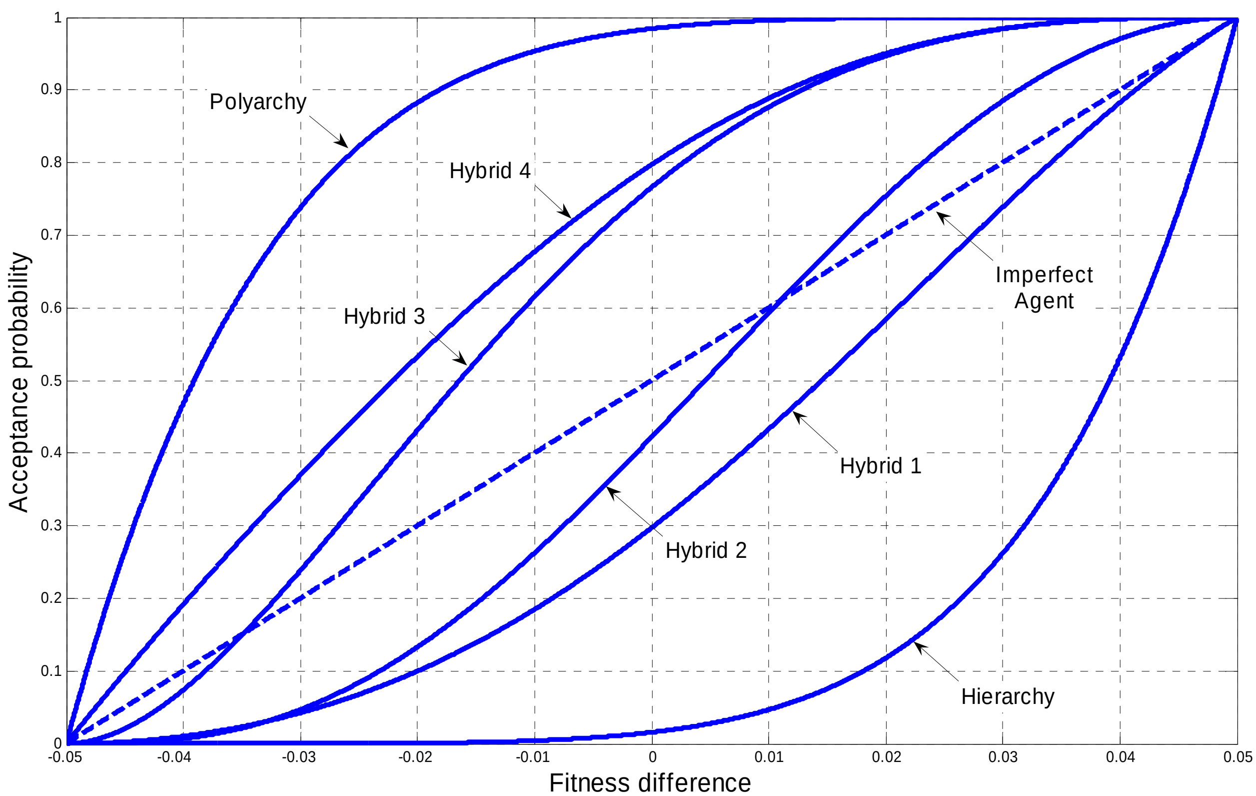

Figure 4: Fitness for perfect agent and imperfect agent (K=3, α= 10, 10,000 agents: 100 distinct landscapes with 100 agents on each).

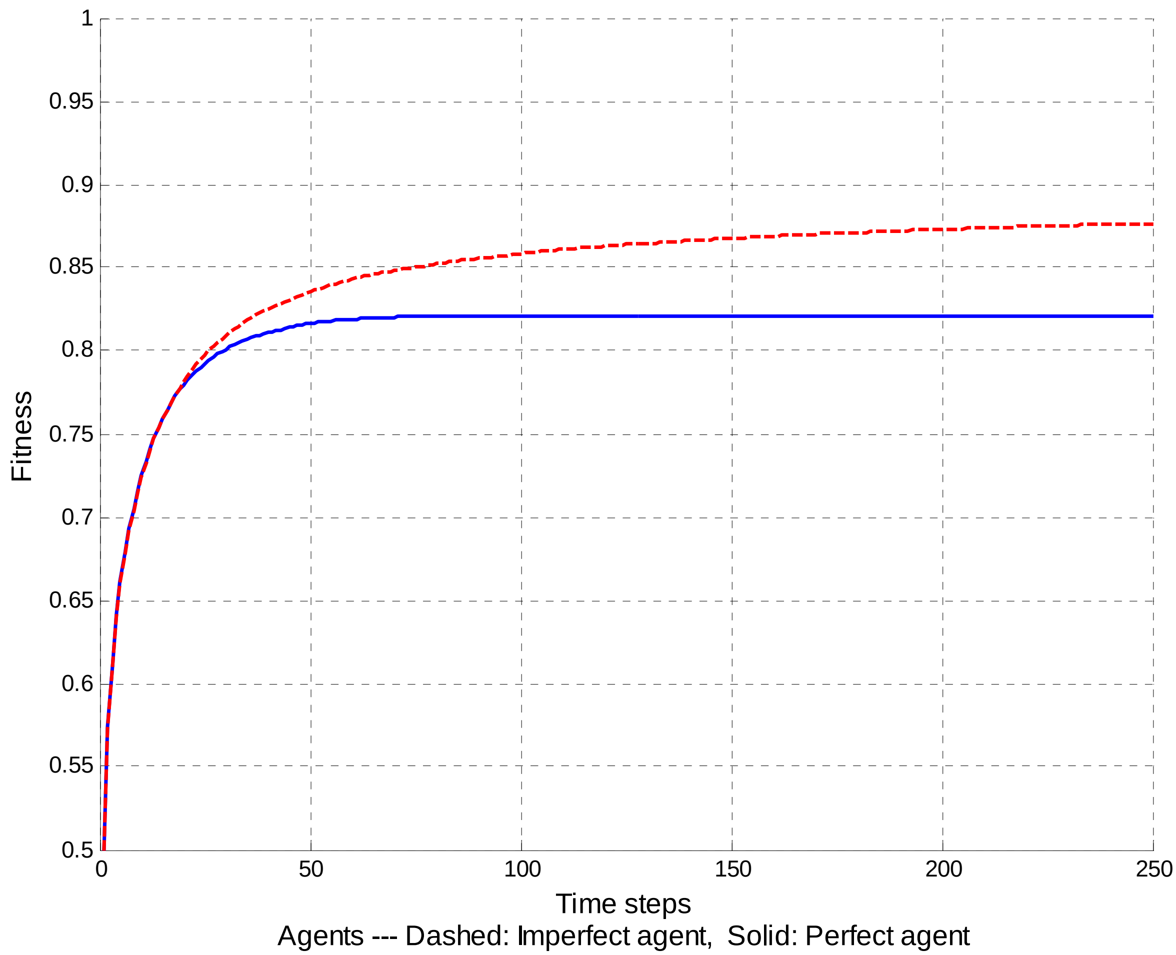

Figure 5: Distribution of imperfect agents at each peak during last 10% of run (K=3, 10,000 agents each on a single randomly selected landscape).

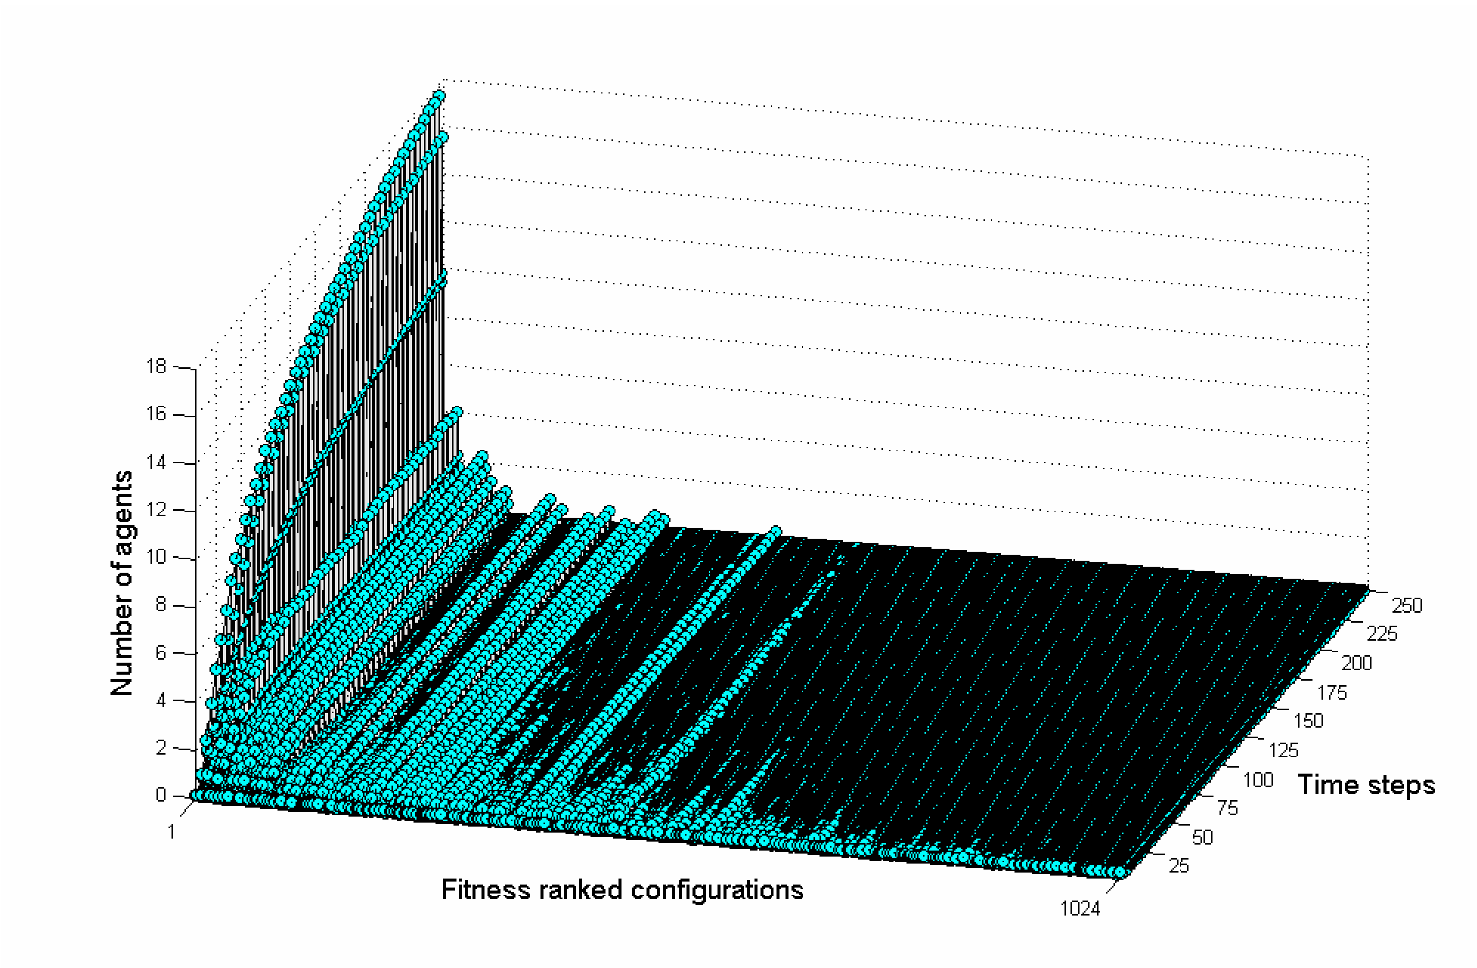

Figure 6: Distribution of perfect agents at each peak during last 10% of run (K=3, 10,000 agents each on a single randomly selected landscape).

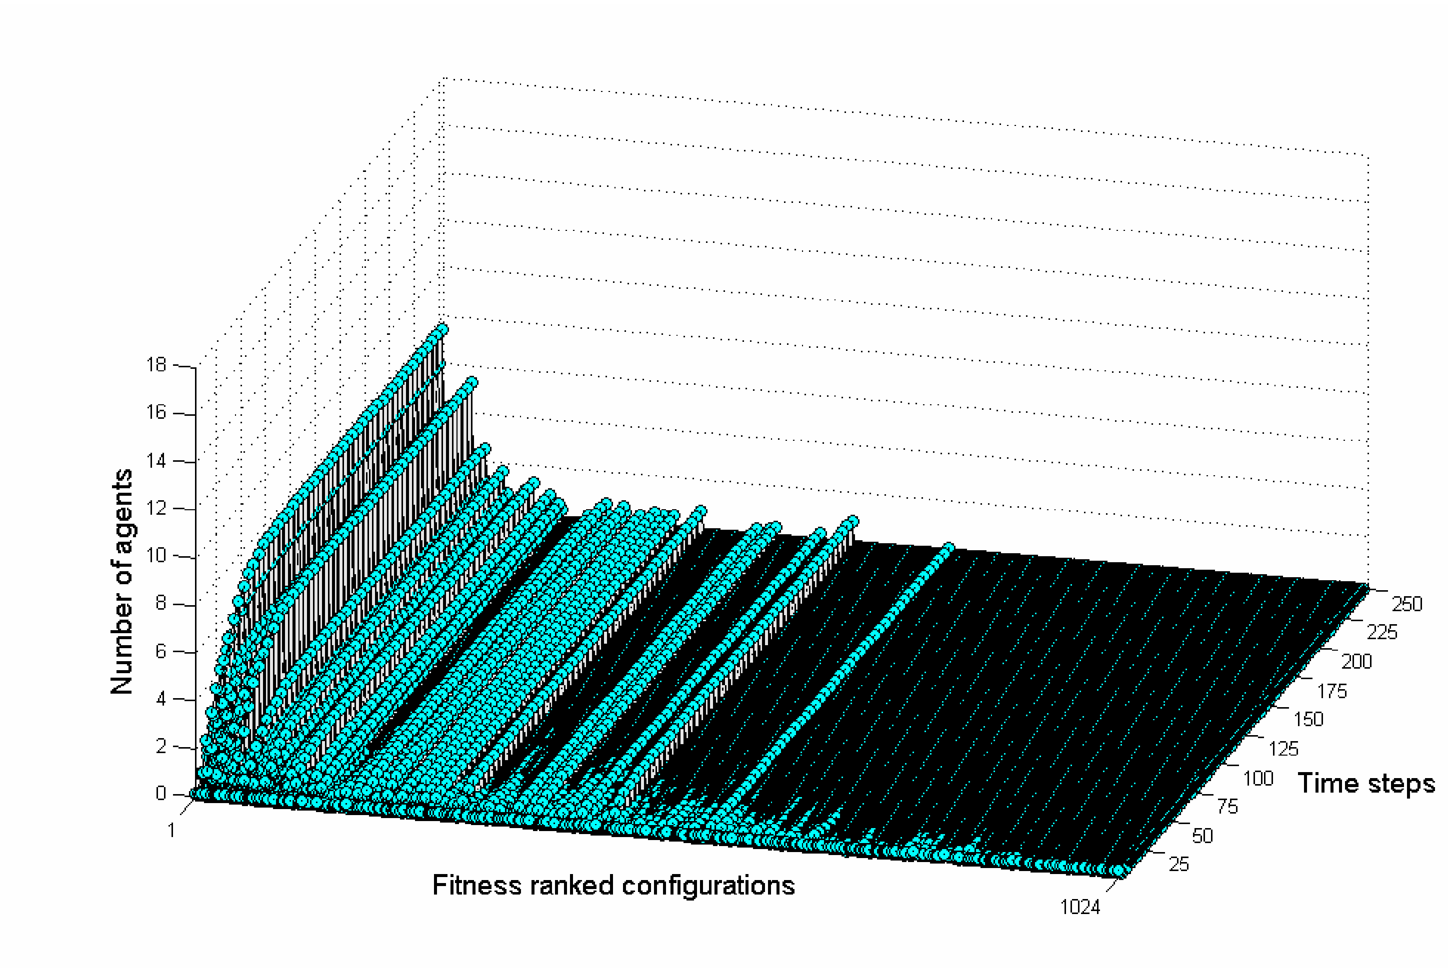

Figure 7: Fitness ranked average distribution of imperfect agents at each peak during last 10% of run (K=3, α= 10, 10,000 agents each on a single randomly selected landscape).

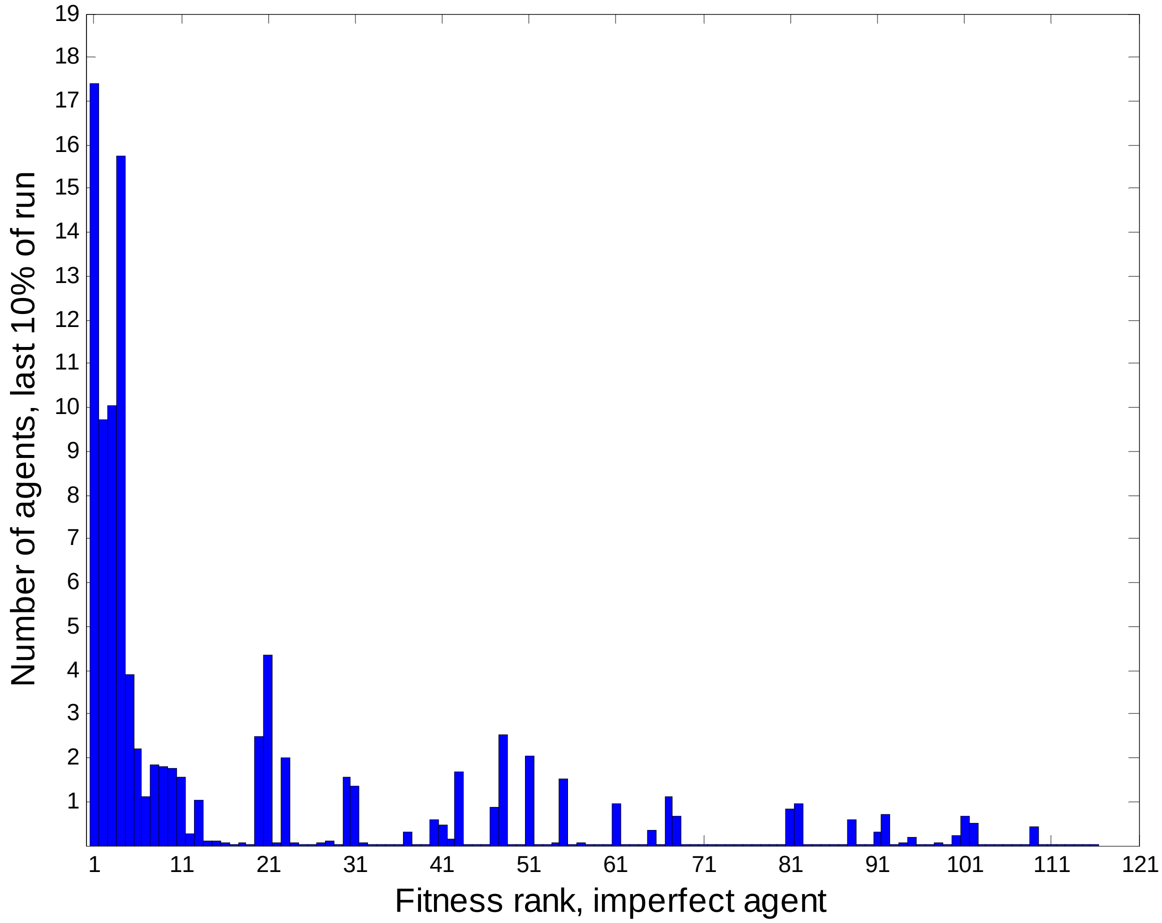

Figure 8: Fitness ranked average distribution of perfect agents at each peak during last 10% of run (K=3, 10,000 agents each on a single randomly selected landscape).

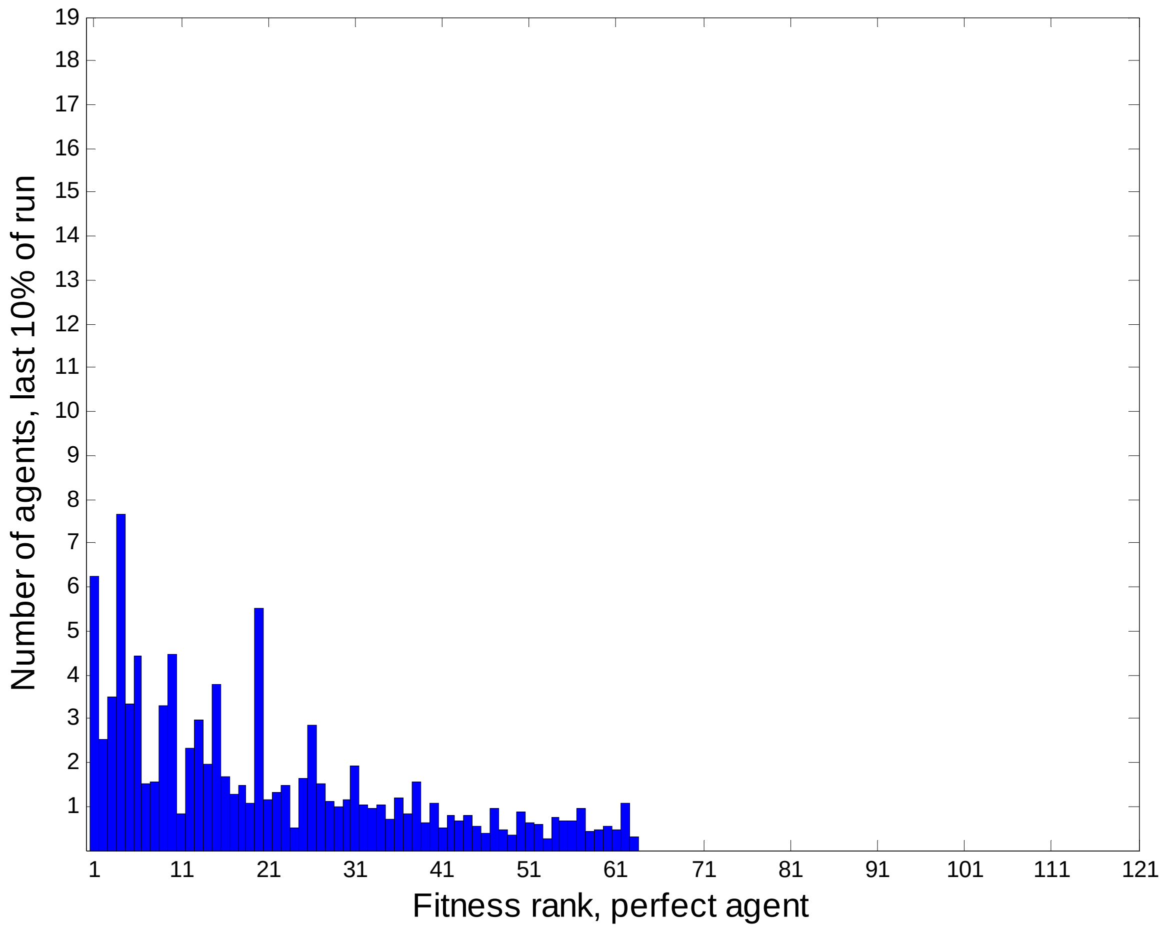

Figure 9: Fitness ranked distribution of agents last 10% of run (K=3, α= {1, 3, 5, 10, 50}, for each level of α: 10,000 agents: 100 distinct landscapes with 100 agents on each). Number of agents at a peak represented by size of circle (scaled by square root). Height of circle represents the fitness of the peak. Darker color of circle represents higher turbulence at the peak measured by the ratio of the number of distinct agents to the average number of agents at that peak. The last panel illustrates the results with a perfect evaluator.

Ability: α= 1

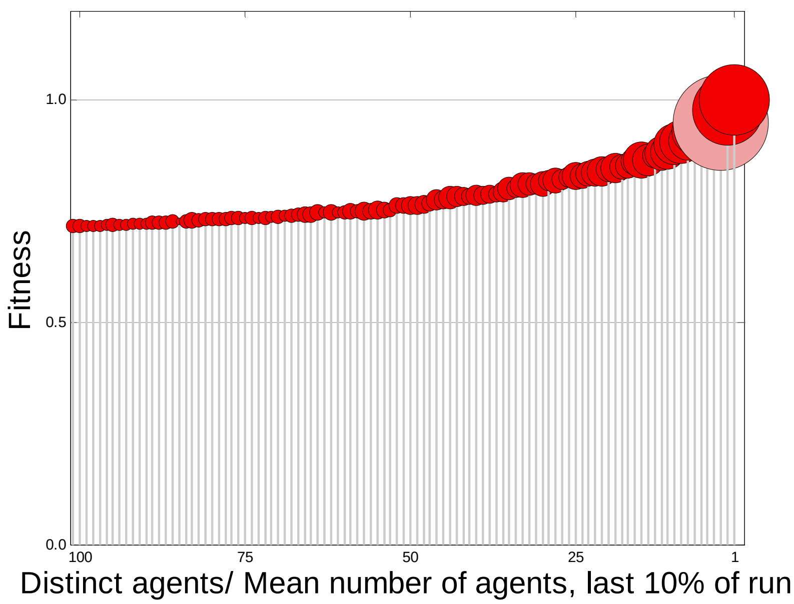

Ability: α= 3

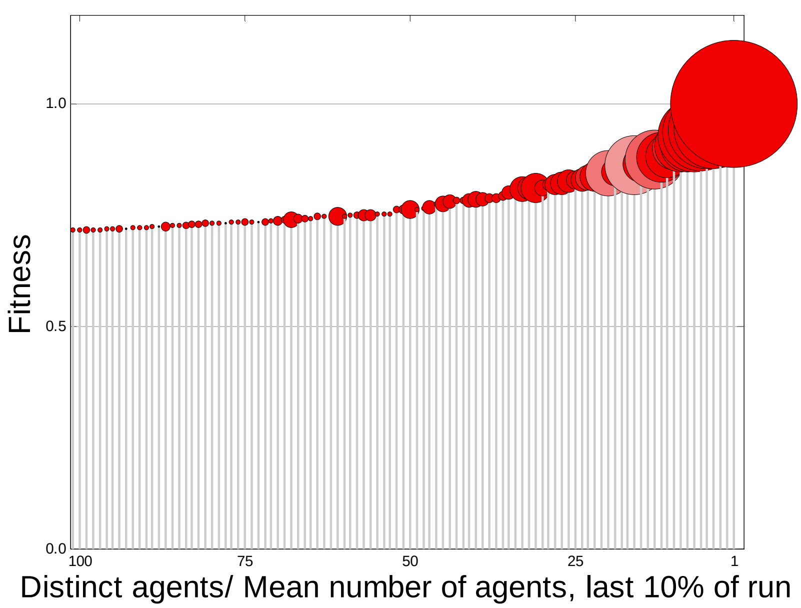

Ability: α= 5

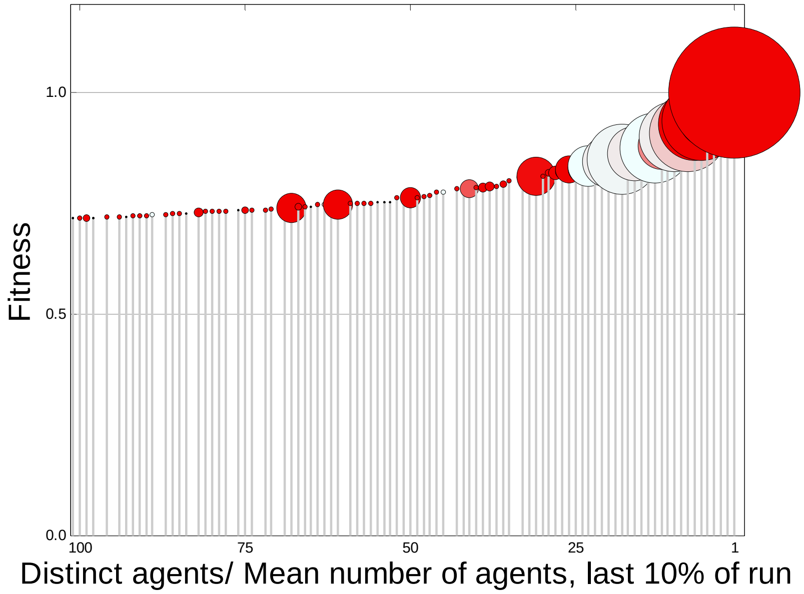

Ability: α= 10

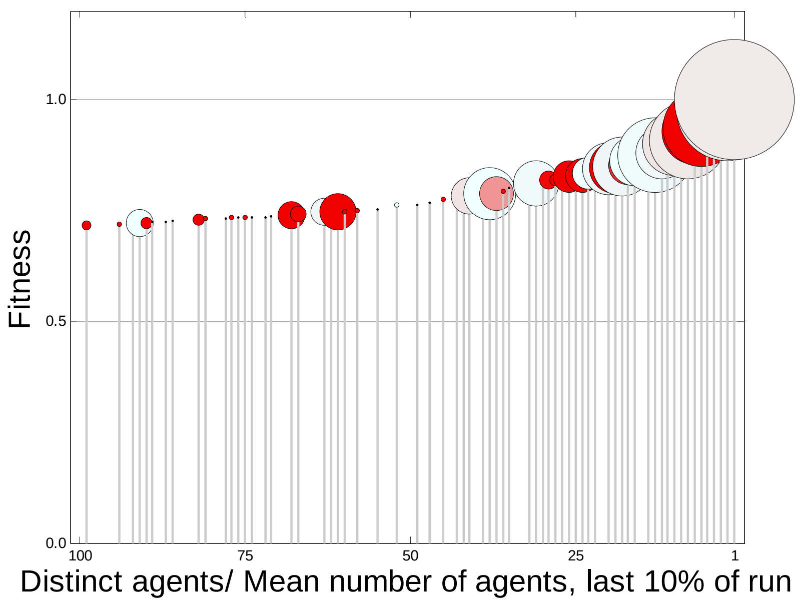

Ability: α= 50

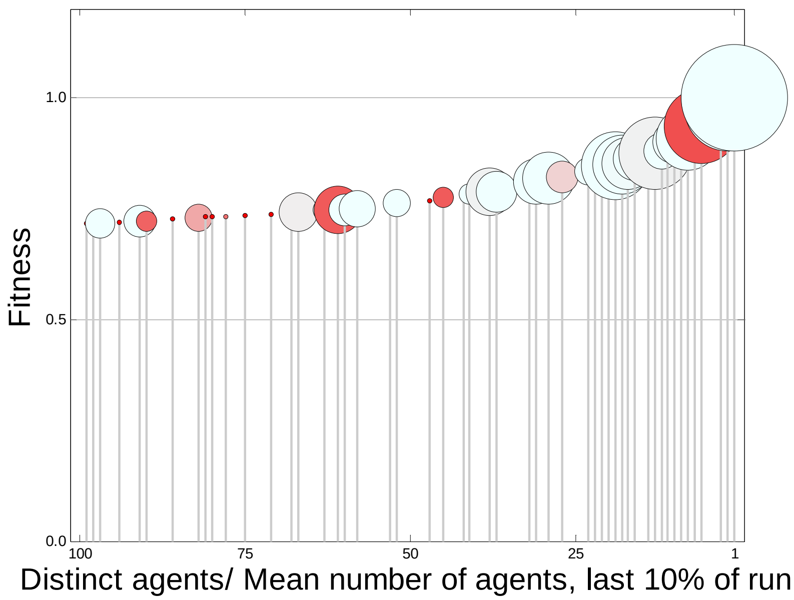

Ability: Perfect evaluator

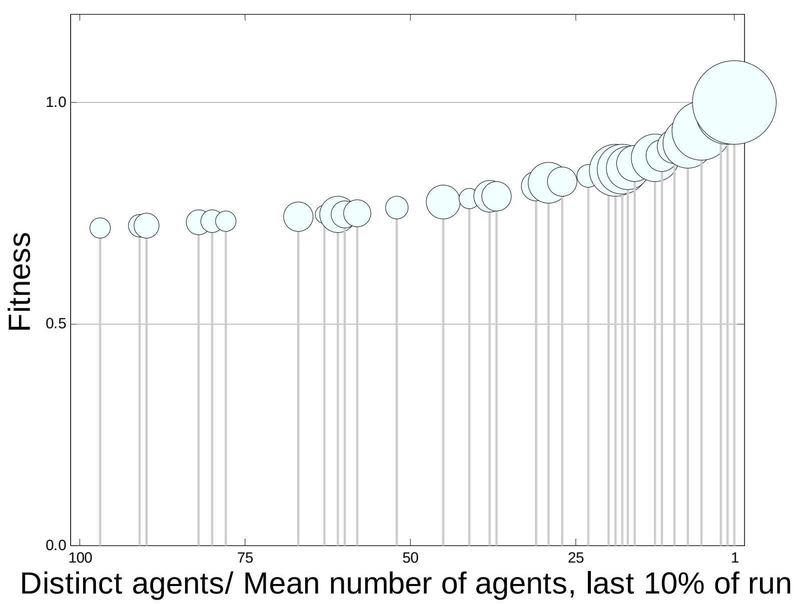

Figure 10: Fitnesses for perfect agent versus imperfect agent (for each level of α and K, 10,000 agents: 100 distinct landscapes with 100 agents on each). Fitness of perfect agent and imperfect agent (α= 10) marked at K=3.

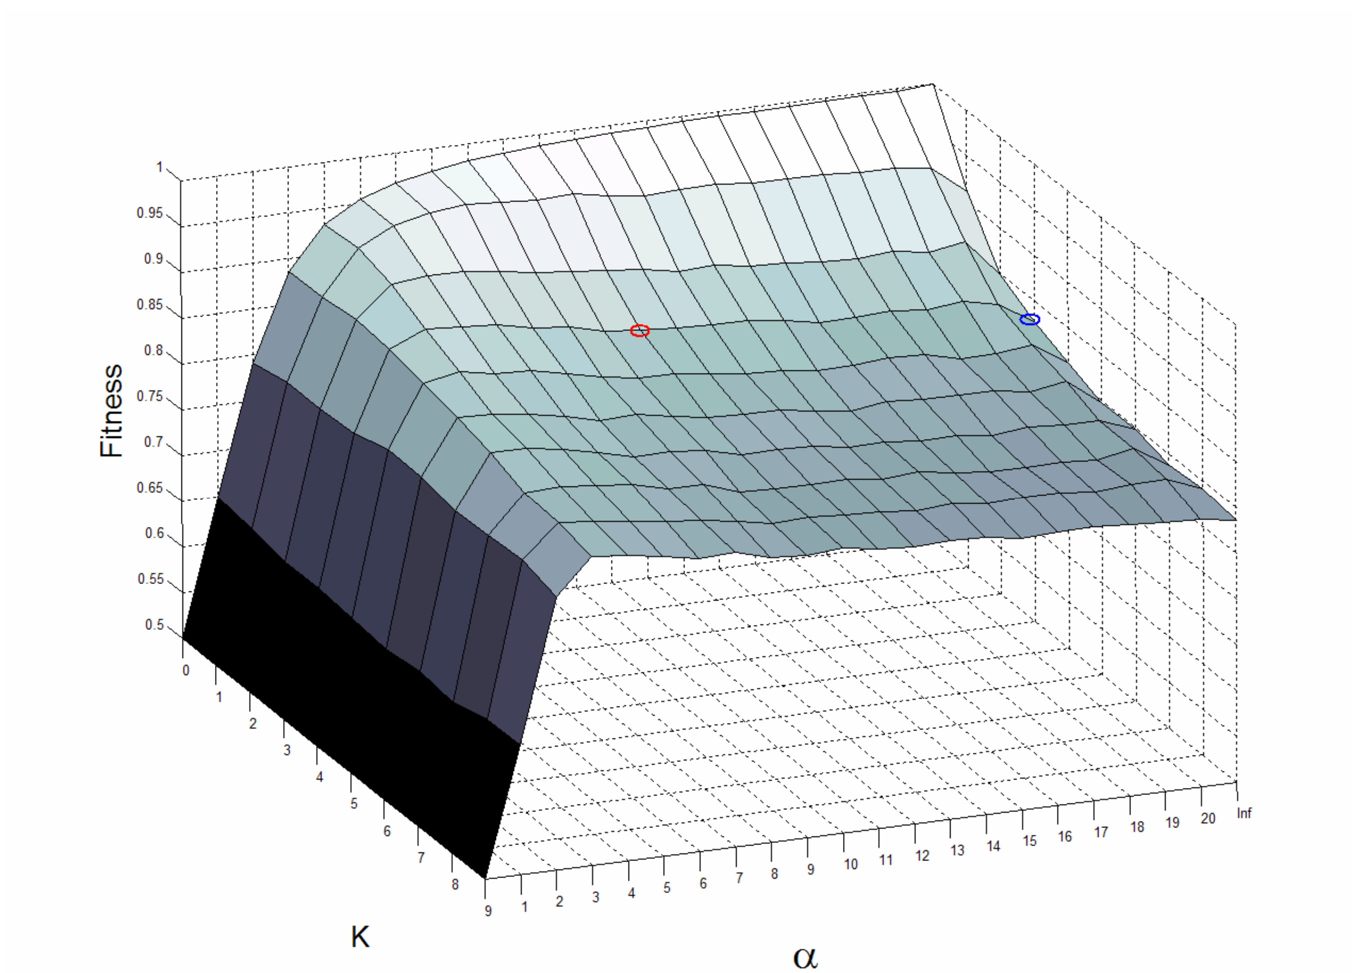

Figure 11: Level curves of fitnesses for perfect agent versus imperfect agent (for each level of α and K, 10,000 agents: 100 distinct landscapes with 100 agents on each). Fitness of perfect agent and imperfect agent (α= 10) marked at K=3.

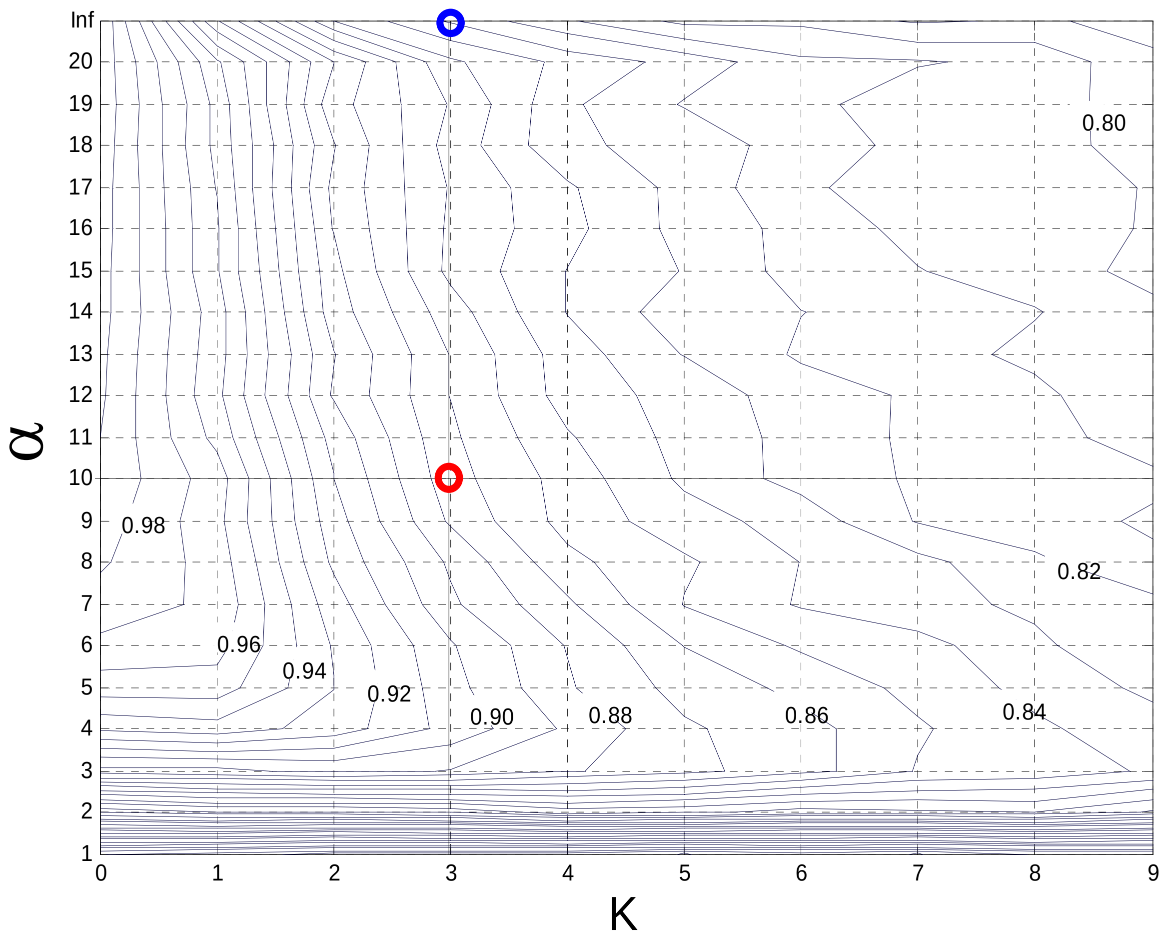

Figure 12: Fitness for perfect agent, imperfect agent and six organizational forms (K=3, α= 10, 10,000 agents: 100 distinct landscapes with 100 agents on each).

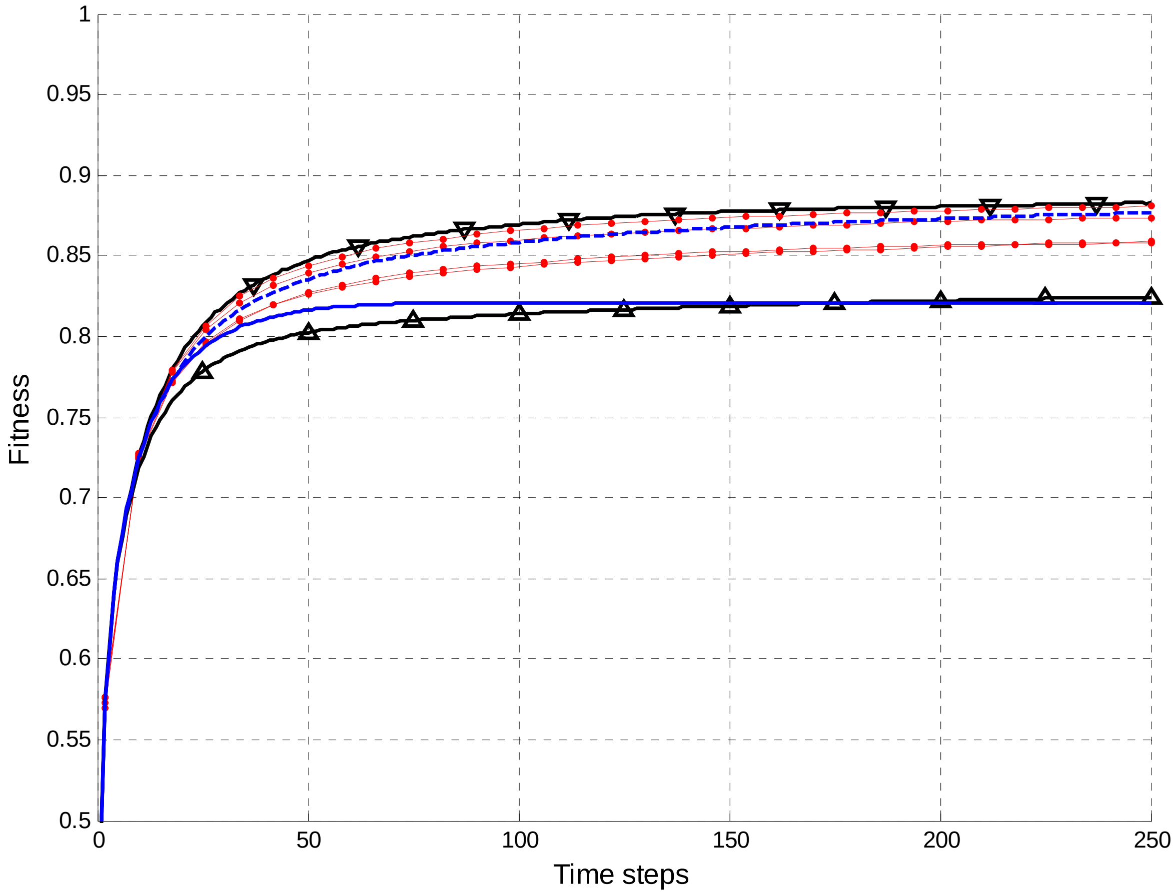

Six organizational forms --- Pyramid: hierarchy, Inverted pyramid: polyarchy, Light dash: hybrid,
Agents --- Solid: perfect agent, Dashed: imperfect agent

Figure 13: Ratio of fitness of hierarchical form to polyarchy for seven levels of α (K=3, for each level of α, 10,000 agents: 100 distinct landscapes with 100 agents on each).

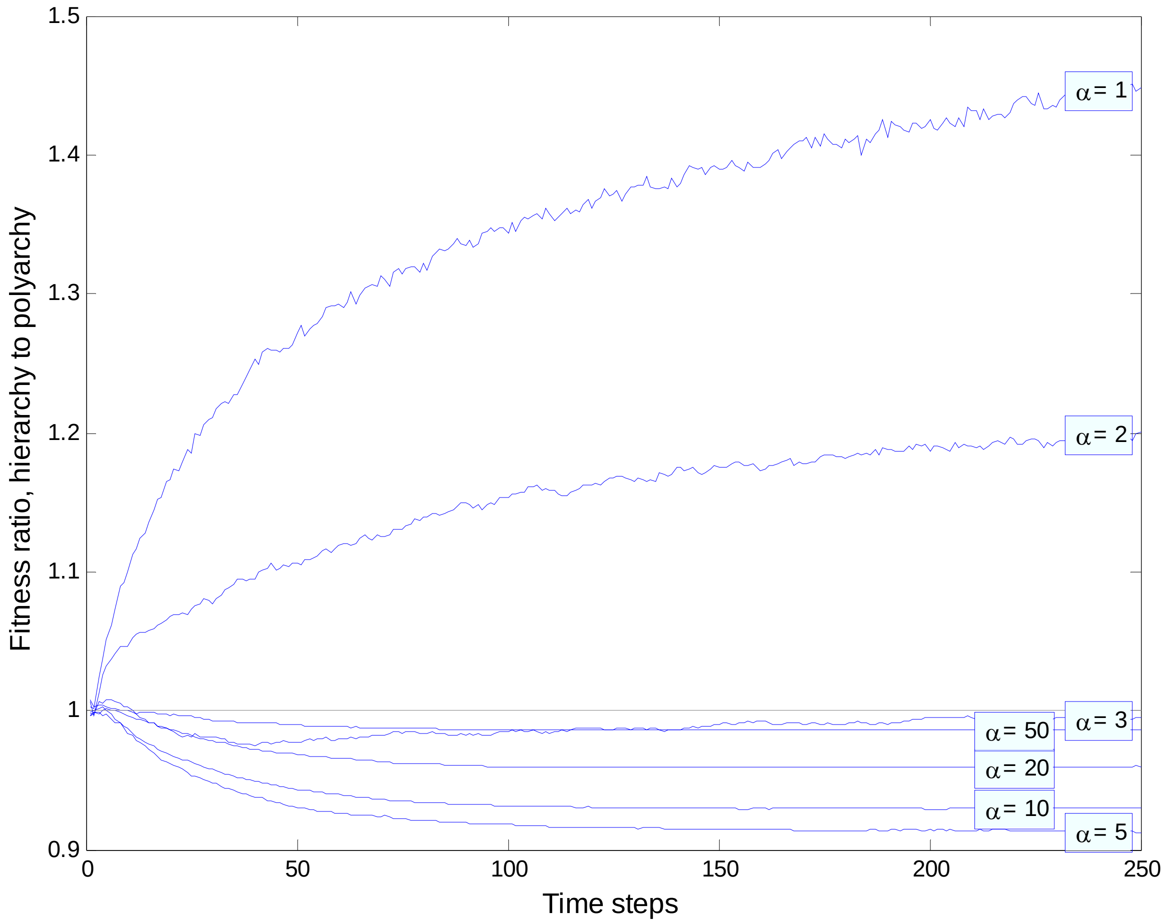

References

Argote, L., (1999). *Organizational Learning: Creating, Retaining and Transferring Knowledge,* Boston, Kluwer Academic Publishers.

Bower, Joseph. (1970). *Managing the Resource Allocation Process*. Boston, MA: Harvard University Press.

Christensen, Michael and Thorbjørn Knudsen. (2003). "The Architecture of Economic Organization: Toward a General Framework". LINK Working Paper.

Cyert, R. M., and March, J. G. (1963). *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice Hall.

Dosi, G., D. Levinthal, and L. Marengo (2003). "Bridging Contested Terrain: Linking Incentive-Based and Learning Perspectives on Organizational Evolution". *Industrial and Corporate Change,* 12: 413- 436.

Ethiraj, S. and D. Levinthal (2004). "Modularity and Innovation in Complex Systems". *Management Science*, 50: 159-174.

Gavetti, G. and D. Levinthal. (2000). "Looking Forward and Looking Backward: Cognitive and Experiential Search". *Administrative Science Quarterly*, 45: 113-137.

Greve, H. 2003. *Organizational learning from performance feedback*. Cambridge, UK: Cambridge University Press.

Holland, John H. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor, University of Michigan Press.

Kauffman, S. A. (1993) *The Origins of Order: Self-Organization and Selection in Evolution.* New York: Oxford University Press.

Lant, T. 1992. Aspiration level adaptation: An empirical exploration. *Management Sci.* **38** (5) 623-644.

Lant, T. K. and S. J. Mezias (1990). "Managing Discontinuous Change: A Simulation Study of Organizational Learning and Entrepreneurship," *Strategic Management Journal,* 11: 147-179.

Levinthal, D. A. (1997) 'Adaptation on rugged landscapes', *Management Science* 43: 934-950.

Levinthal, D. A., and March, J. G. (1981) 'A model of adaptive organizational search', *Journal of Economic Behavior and Organizations* 2: 307-333.

Levitt, Barbara and James G. March. (1988). “Organizational learning”. *Annual Review of Sociology*, 14: 319-340.

Lin, Z. and K. M. Carley. (1997). “Organizational Response: The Cost Performance Tradeoff,“*Management Science*, 43: 217-234.

March, J. G. (1991) ‘Exploration and exploitation in organizational learning’, *Organization Science* 2: 71-87.

March, James G. and Herbert A. Simon. (1958). *Organizations*. New York: John Wiley & Sons.

Mintzberg, H. (1979) *The Structuring of Organizations*. Englewood Cliffs, NJ: Prentice Hall.

Nelson, P. (1970). “Information and consumer behavior”. *Journal of Political Economy,* 78: 311-329.

Nelson, Richard. (1982). “The role of knowledge in r&d efficiency”. *Quarterly Journal of Economics,* 97: 453-470.

Nelson, Richard and Sidney Winter. (1982). *An Evolutionary Theory of Economic Change*. Harvard University Press.

Rivkin, J. W., and Siggelkow, N. (2003) ‘Balancing search and stability: Interdependencies among elements of organizational design’, *Management Science* 49: 290-311.

Rosenkopf, Lori and Atul Nerkar. (2001). “Beyond local search: Boundary-spanning, exploration, and impact in the optical disk industry”. *Strategic Management Journal,* 22: 287-306.

Sah, R. and J. Stiglitz. (1986). “The Architecture of Economic Systems: Hierarchies and Polyarchies”. *American Economic Review,* 76: 716-727.

Seshadri, Sridhar and Zur Shapira. (2003). “The flow of ideas and timing of evaluation as determinants of knowledge creation”. *Industrial and Corporate Change,* 12: 1099-1124.

Simon, Herbert A. (1955), “A behavioral model of rational choice”, *American Economic Review*, 69, 99-118.

Simon, Herbert A. (1962). "The architecture of complexity." *Proceedings of the American Philosophical Society*, 106: 467-482.

Sorenson, O. (2002) ‘Interorganizational complexity and computation’, in J. A. C. Baum (ed) *Companion to organizations*, pp. 664-685. Oxford: Blackwell.

Stigler, George (1961). “The economics of information”. *Journal of Political Economy,* 69: 213-225.

Stuart, Toby and Joel Podolny. (1996). “Local search and the evolution of technological capabilities”. *Strategic Management Journal,* 17: 21-39.

Wright, Sewall. (1931). "Evolution in Mendelian populations". *Genetics*, 16: 97-159.

Appendix 1

The organizational level screening function, *F*, is a polynomial in the individual level screening function, $f(x)$, under the assumption that all members of an organization have identical screening functions. Given the flow of decisions shown in Figure 2, we derived organizational level screening functions (shown in Figure 3) for the six organizational forms. They are as follows.

Hierarchy: $F = f(x)^6$

Hybrid 1: $F = f(x)^6 - 3f(x)^5 + 2f(x)^4 + f(x)^2$

Hybrid 2: $F = f(x)^6 - 6f(x)^5 + 15f(x)^4 - 18f(x)^3 + 9f(x)^2$

Hybrid 3: $F = -f(x)^6 + 6f(x)^5 - 12f(x)^4 + 8f(x)^3$

Hybrid 4: $F = -f(x)^6 + 4f(x)^5 - 5f(x)^4 + 2f(x)^3 - f(x)^2 + 2f(x)$

Polyarchy: $F = 1-(1-f(x))^6 = -f(x)^6 + 6f(x)^5 - 15f(x)^4 + 20f(x)^3 - 15f(x)^2 + 6f(x)$