
Leading large: emergent learning and adaptation in complex social networks

James K. Hazy

Adelphi University,
1 South Avenue, Garden City NY, USA
and
Mälardalen University,
Box 325, SE-631 05 Eskilstuna, Sweden
E-mail: hazy@adelphi.edu

Abstract: The mechanism by which adaptation in large organisations emerges from localised bottom-up processes remains largely unexplored. This paper describes the emergence of a learning algorithm in organisations which crosses levels of analysis. It posits that what is essentially a neural network arises naturally in organisations with individuals as nodes, interactions as edges, and influence relationships among them performing a function that is analogous to synaptic weights. This network structure enables organisations to adapt through a process of influence process structural learning that is analogous to back propagation learning in traditional neural networks. The model describes leadership within top management as expressing the organisation's response to environmental stimuli about which top managers have little direct knowledge. Leadership acts to change influence relationships among managers by altering their relative status and reputations. The theory implies that influence relationships exhibit a power-law distribution, a potential marker of emergent collective agency.

Keywords: leadership; complexity; organisational learning; adaptation; human interaction dynamics; agent modelling; organisation theory; social networks; power laws; emergent innovation.

Reference to this paper should be made as follows: Hazy, J.K. (2012) 'Leading large: emergent learning and adaptation in complex social networks', *Int. J. Complexity in Leadership and Management*, Vol. 2, Nos. 1/2, pp.52–73.

Biographical notes: James K. Hazy is a Full Professor at Adelphi University and a Visiting Professor at Mälardalen University in Sweden. He has published over 40 articles and book chapters in journals such as *The Leadership Quarterly*, *Emergence: Complexity and Organization*, and *Nonlinear Dynamics in Psychology and the Life Sciences*, and three books, most recently: *Complexity and the Nexus of Leadership*. He obtained his doctoral degree from the George Washington University and his MBA from Wharton School. Previously, he was an executive with AT&T and Ernst & Young and is a Founding Director at the Institute for Research in Complexity and Society (IRCS).

This paper is a revised and expanded version of a paper entitled 'Emergent agency in collectives: influence power laws and evolutionary selection in the garbage can' presented at the 14th Organization Science Winter Conference (OSWC-XIV), The Resort, Olympic Valley, Squaw Creek, California, USA, 6–10 February 2008.

1 Introduction

The question of whether leadership at the top of large organisations is qualitatively different than elsewhere remains controversial. Do members of top management teams (TMTs) do the same things as other managers but do them better, or do they do things that are qualitatively different? This paper uses an illustrative agent-based model (ABM) that explores leadership (Hazy, 2007) to identify what might indeed be a qualitatively distinct mechanism whereby top management is able to impact an organisation's adaptive success. It describes a mechanism called influence process structural learning (IPSL) that gathers and uses information from events (Lichtenstein et al., 2006) in the environment (stimuli) to change the internal flow of social influence (and thus resources) through the organisation's internal network improving its resource allocation choices (response) over time. Thus, in complex environments (Marion and Uhl-Bien, 2001) when timely action is needed as events unfold, greater emphasis is placed on determining who will have how much influence rather than focusing on building a centralised process for assessing the situation and optimising decision analysis and execution from the centre.

1.1 Status, reputation and cumulative advantage

This paper argues that within large organisations, IPSL occurs by elevating the relative status and thereby the influence of objectively successful project champions, and by diminishing the relative influence of less successful individuals. This has the effect of indirectly placing more resources at the disposal of those who are able to identify better opportunities (which can be thought of as environmental stimuli that require responses) through the effective use of the social networks in which they are embedded. When in addition, increased status or reputation has the effect of bringing more and better projects to an individual's attention, a positive feedback cycle embeds cumulative advantage into the influence process structure. As a result, it is posited that relative status and reputation reflect a scale-free frequency distribution in organisations as a means to efficiently process information (Baek et al., 2011). One would expect that IPSL would thus improve fitness so that successful structures would be selected in competitive environments as marginal advantage separates winners from losers.

Presumably this dynamic is also present at lower levels in the organisation. However, this particular mechanism becomes much more important at the top of the organisation because it also serves a unifying function (Hazy, 2012; Hazy and Uhl-Bien, in press). By expressing the organisation as a single entity in the environment, feedback about the efficacy of the collective's organisational capabilities (Helfat et al., 2006) is clarified with regards the success or failure *of the system*. Thus, when leadership implements an effective learning algorithm such as IPSL, the system internally selects more successful capabilities by virtue of the choices made through its signaling and influence networks. In doing so, the system exhibits the emergence of collective agency, at least with regards the expression of a collective response to environmental stimuli. The system does this by learning from the results of its actions. It actualises this by changing its internal influence structure to improve its sensing and its responding capabilities as a system. Note that learning occurs even though no single individual, or even any set of individuals, fully comprehends what the system has learned. What has changed is the process through

which various individuals participate in the social influence structure that enacts future decisions.

1.2 Emergent agency in the garbage can

The argument averred herein assumes that a decision making process similar to the garbage can metaphor (Cohen et al., 1972) is at work in the organisation. Agents find projects to sponsor through their networks and their relative reputations. As decisions are taken, the status of each agent who advocated a successful project is enhanced. Those that fail lose relative status. When this happens the influence process structure of the organisation is changed as well. This change represents a structural variation in the evolutionary sense, and this variation can be adaptive if it enables enhanced fitness. [Note that throughout this article fitness is defined as the probability that the organisation will survive into the next time step. The focus is on intra-generational fitness relative to competition rather than on reproductive fitness which is the traditional focus of selection in evolutionary biology (Okasha, 2006).] System-level learning does not rely on any individual or individuals. Rather, it occurs when information about what has succeeded in the past is embedded into the structure of the system. This is done through a leadership mechanism that changes the relative influence of individuals within the organisation's social network, enhancing the influence of those who have succeeded and diminishing the influence of those who have failed.

Complex systems researcher Gell-Mann (2002) describes complex organising by distinguishing between events or interactions that occur among individuals at the 'fine-grained' level, from the emergent regularities that can be modelled as 'coarse-grained' properties observable at the system level. As the organisation's coarse-grained capabilities succeed or fail in the environment, information created from the resulting events is back-propagated into the system to influence future fine-grained interactions inside the organising effort. These changes, in turn, impact the properties, capabilities and organising forms that are emerging. Due to preferential attachment mechanisms in social networks (Barabasi, 2002), the presence of a structural learning process like what is described herein is likely to have a power-law signature.

By enabling the above described learning process, top leaders of complex organisations do something in addition to their other commonly enumerated duties (see for example Yukl, 2000) or even those duties described in complex systems leadership theory (Uhl-Bien et al., 2007; Uhl-Bien and Marion, 2008; Hazy, 2008, 2011; Hazy et al., 2007; Hazy and Uhl-Bien, in press; Surie and Hazy, 2006). Namely, these individuals actively facilitate the construction of influence process structures within their organisational networks in an effort to place the right individuals with the appropriate knowledge and experience into the best positions to make decisions and to execute programmes of action.

An organisational meta-capability (Hazy, 2008) like this is necessary in large organisations because information about the fitness of an organisation's coarse-grained capabilities is distributed in the environment and different aspects of that fitness are observable by many different individuals. By changing how influence is distributed across the organisation's internal network of relationships, the information gathered and used to formulate the system's response to these stimuli can be modified and potentially improved. By changing how the influence of various individuals acts to enable and constrain the allowable fine-grained interactions among other individuals, the 'rules'

governing fine-grained interactions change. Thus, identifying an IPSL process within an organisation would provide evidence that complex systems leadership is occurring in that organisation (Hazy et al., 2007).

Changing fine-grained interactions changes the nature of the coarse-grained capabilities that emerge from them (Gell-Mann, 2002). This facet of complex systems leadership bridges from fine-grained interactions to coarse-grained capabilities and back again, changing the unfolding of each of these according to the information gathered from events in the environment. In particular, this mechanism leads the organisation's adaptation process by helping it learn from experience *as a system* in a manner analogous to back-propagation in neural networks. The organisation effectively acts as its own agent, advancing its own ends, and not those of any individual agent, even those with relatively high status. Collective agency emerges even as individual agents continue to act in their own self-interest.

In sum, this paper argues that a previously unidentified adaptive mechanism – called IPSL – can arise naturally in organisations. Further, when it does, a qualitatively different type of executive function emerges to unify the collective and exhibit a level of collective agency. The illustrative ABM described in later sections demonstrates that under a simple set of assumptions, this emergent structure can occur in simulations. The model is not intended to simulate a real world organisation; rather, the model attempts to balance the benefits of simplicity and transparency against the complications that result from attempts to achieve veridicality (Carley, 2002). The model depicts an organising process that is enacted by relatively simple-minded agents, each acting in its own self-interest. Nonetheless, the system exhibits organising-level learning through a back-propagation mechanism similar to what is apparent in neural networks. Like the brain, the organising system learns by changing its structure, and as a result, a level of supra-individual agency at the system level becomes a possibility. The learning mechanism is described in the next section.

2 A neural network model of organisations

Drawing on network theory, this analysis assumes that prior to organising together, individuals or agents are initially connected in a substrate network, perhaps with a small world topology – a condition where groups of agents are connected locally in a dense network of strong ties while some agents are connected to others over weak ties (Barabasi, 2002; Granovetter, 1973) or even with a more complex network such as one that might be represented with CONSTRUCT (Carley and Ren, 2001; Schreiber and Carley, 2006) or as described by Goldstein et al. (2010). For the present analysis, assume that a leadership influence process network, that will be described later, develops on top of such a substrate and that it forms through existing and additional connections that are analogous to those of a neural network as shown in Figure 1.

2.1 Individuals as 'neurons' in the network

Considering an organisation as an interconnected network of heterogeneous agents has some interesting implications. Each of the agents, individually and in groups, interacts with each other and the environment. In so doing, each encounters different opportunities

for the collective, but they do so with a perspective that is limited by their position. These 'worker agents' (Ws) are the nodes or the 'neurons' in the network that act as the input layer for the organisation. Because each is limited by its local position, it is difficult for it to know if the opportunity that is observed – the stimulus – is worth pursuing in the larger organisational context – that is, that the stimulus deserves a system response. At the same time, if one assumes that these local stimuli are the only inputs into the system – in other words, if one assumes that from a practical perspective the decision-maker agents in the TMTs are isolated from primary field data that arises in day-to-day business interactions – then for the system to survive, it must in some way determine to which input stimuli to respond at the system level and in what way. [Note that this assumes that there is a class of decisions that are taken by the organisation as a whole and that these are qualitatively different than day-to-day decisions such as pricing and purchasing decisions that are made at lower levels in the organisation. The decisions referred to herein involve the recognition of patterns in the environment that require a system-level response.]

The challenge for a system of individuals facing the environment as a system is to recognise opportunities and to act upon them even though information and knowledge is widely distributed and no one agent, including any of the nominal executive-leaders or TMTs, is in a position to have perfect knowledge of the possibilities or even to operate with a complete picture of the system's needs.


The simplest answer is to add an output layer as shown in Figure 1. The output layer is composed of decision making individuals, the TMTs who acts as, the 'neurons' who decide what the collective says or does; in other words, although these agents do not have direct information about events in the environment, they nonetheless are the ones who decide how the organisation responds to stimuli in the environment. Note that in this conceptual framework, because these TMTs do not have direct knowledge of the environment, they must rely on ideas brought to them by middle-manager agents (MMs) in the hidden-layer (see next paragraph). MMs in turn rely on the worker agents (Ws) who take stimuli directly from the environment albeit without being able to judge their relevance or potential value to the system.

Early research in neural network learning demonstrated that for complex learning to occur – what is called back propagation learning – more is needed than the functional distinction between input and output nodes. What is also needed is a 'hidden layer' that serves to aggregate inputs from the workers and to represent those inputs to the TMTs in the output layer (Bossomaier, 2000). As shown in Figure 1, in an organisational setting, agents in middle management roles, MMs, act as this 'hidden layer' within the organisation's network.

Recent modelling has shown that as internal interaction and strategic complexity increase, it becomes exceedingly difficult if not impossible to accurately predict the consequences of these individual complex interactions inside the organisation; they are so highly contingent (Siggelkow and Rivkin, 2005) that attempted prediction using non-linear models that are based on ill-defined and poorly measured variables is problematic. One possible reason for this is the problem of interaction or 'complexity catastrophe' (Kauffman, 1993) that must be managed in organisations. Complexity catastrophe occurs when too many interactions overload individuals, a situation that must be managed through work flow design which clarifies individual responsibilities

(Hazy, 2008). This is not to say that individuals can't succeed in these situations; they can, and they do. The point is that one's ability to develop analytical models is limited and that traditional models are unable to accurately represent the process that leads to success.

Figure 1 Representation of an organisation as a neural network



Notes: As described in the text, intra-generational learning occurs as structural changes are propagated into the network; this occurs as influence weights among agents change as errors in responses are experienced and recognised; to respond to stimuli received by input agents, an output layer disambiguates internal information that is processed through agents that are 'hidden'.

As it turns out, the twin problems of distributed knowledge and interdependent complexity may both be solved in organisations by a single organisation-level adaptation: the implementation of what amounts to a back propagation learning algorithm operating on a signaling network in the organisation (Hazy, 2007; Holland, 2001). This conceptual 'signaling network' is embodied within the communications infrastructure that supports the organisation. It refers to the process whereby individuals chose whether to share information they have gathered and with whom to share what information; heuristically, one can think of this as 'picking the right people with the right connections for the right jobs'. In total, the signaling network that results for this selection process serves the same function as the connections and synapses in a neural network.

2.2 *'Training' the network*

In the language of neural networks, learning occurs as the network is 'trained'. What this means is that input layer agents – *Ws* – gather data from the environment during their day-to-day activities. In other words, they receive 'stimuli' for the system, but they don't know what to make of it. By prior network connections, certain hidden layer agents – *MMs* in our framing – learn of the stimulus through input layer *Ws*. These hidden layer *MMs* are in turn connected to output layer agents, *TMTs*. Based upon these historically determined connections, the *TMTs* choose to output a 'response' into the environment. A 'correct' response causes the network to reinforce the connections that participated in organising the successful response by increasing the status or reputations of those who participated. An incorrect response, in turn, causes the relevant connections to decay and lose strength according to the level of perceived error. How this feedback is processed back into the system is called the organisation's 'learning algorithm'. Over many trials, the network is 'trained' to respond correctly.

2.3 *Agent reputation and synaptic weights*

In neural networks, when one neuron or node connects to another and causes it to respond, the second node is said to be 'activated' along with the first. In situations where one neuron is connected through a synapse to another, when the first is 'activated', if the strength of the impulse along a synaptic connection is above an activation threshold, the second will activate along with the first (Bossomaier, 2000).

By analogy, in human organisations, the likelihood that one individual will activate along with another relates to the strength of connection between the two. Individuals are more likely to offer information to or accept it from those who are perceived as having relatively better reputations than others. As such, this paper posits that the reputation of an agent determines whether the other agent will take action or become 'activated' once information about an opportunity reaches the first agent. In this context, to 'activate' is to enlighten by recognising the stimulus, i.e., the opportunity in the environment. This implies a proposition.

Proposition 1: In an organisational setting, when an agent receives new knowledge – a stimulus from the environment – the likelihood that a second agent will receive that information from the first is positively related to the level of the reputation of each agent in the mind of the other, and this is mediated by the existence, strength and quality of the relationship connection between them.

2.4 *Agent status and disambiguation*

When stimuli that originated in the environment – and have been propagated through several nodes of the network in this way – finally reach the 'output layer', they are most likely ambiguous. The nodes in the hidden layer that eventually activate the output layer, are likely to have themselves been activated by several input nodes each with uniquely interpreted inputs; these may conflict or they may complement one another. In any case, at the output layer, the facts themselves may be indeterminate.

By analogy, in human organisations, information flows from *Ws* through *MMs* to *TMTs*. Information that originated in the environment has been propagated through

several layers of the organisation before reaching the ‘output layer’. Thus, for TMTs, the relevant facts are often ambiguous or even indeterminate. Recent computational models have explored this process and have determined that even in very simple cases, a decision process such as this can be highly contingent on the specifics of the system (Siggelkow and Rivkin, 2005) and might therefore be impossible to predict with any confidence. What then enables success?

The present analysis describes a decision process involving much simpler dynamics. Each of the TMTs might have prior experience with the hidden layer MMs and might therefore have a perception of their relative abilities to deliver the results that they had promised in the past. Perhaps, for example, some of the MMs have developed the skills needed to overcome the complexity catastrophe (Hazy, 2008) as they implement complex projects. These MMs would have demonstrated that they have the skills to succeed in delivering results, and this might be known to the TMTs.

An organisation is assumed to include an output layer of TMTs – comprising the TMT – who make decisions to move forward in complex and uncertain environments based not only on the relative merits of the projects, but also on the relative status of the agents who are proposing them as well as on their individualised reputations in the eyes of the TMTs. This analysis supports the heuristic that picking the right people and holding them to account, that is, *putting the right people in the right seats on the bus* (Collins, 2001) is key to success at the top. These ideas imply a proposition.

Proposition 2: In complex environments, decisions are taken by TMTs based in large part on the relative status of those middle management agents (MMs) proposing the alternatives and on the reputations of those agents in the mind of the TMTs.

2.5 Changing status and back-propagation learning

For a system to have the property of IPSL, a learning algorithm must be identifiable. In other words, feedback from the environment with respect to the appropriateness of the system’s response to its stimuli must be processed back into the system as changes to the systems influence process structure. This must be done in ways that help the system embed its experience – with respect to successful opportunity recognition and programme execution – into its influence structure as useable information and knowledge.

Focusing on obtaining useful feedback from the environment with respect to its success at identifying opportunities and then adjusting influence relationships according to success or failure – in particular, adjusting the coarse-grained structure defining the relative status relationships among its agents as influencers of decisions – is distinct from traditional views of organisational learning. These highlighted systemic information processing elements (March and Olsen, 1975; Schwandt and Marquardt, 2000) rather than the relative influence of individual agents within the organisation’s structure. This new perspective implies a proposition.

Proposition 3: In an organisational setting, the extent to which

- a the evaluation of success or failure is based upon factual data direct from the environment (and is not filtered through the interpretations of the project’s sponsor)
- b this information is used to adjust the relative status of proposing and implementing agents in a manner that correlates to success

Are each positively related to the level of IPSL and thus fitness with respect to identifying and capitalising on viable opportunities.

The structure of status relationships in organisations represents one of its coarse-grained properties. It reflects the organisation's adaptive capacity, and relates to its potential to reconfigure itself as circumstances change. The fitness benefits emanating from this structure can even survive turnover of individual agents and therefore reflect emergent agency at the collective level. Furthermore, because these status relationships cross scale, they are likely to be scale-free and thus might be expected to exhibit a power law relationship. As a result, the presence of IPSL is related to preferential selection with regard to reputations among agents (Barabasi, 2002), and this implies additional propositions.

Proposition 4: Power-law relationships in the relative aggregate reputations of agents and in the structural status that organises agents as the organisation learns to respond to stimuli in the environment are necessary but not sufficient conditions for IPSL to occur.

Proposition 5: A positive relationship between agent intelligence at picking projects (as defined below) and both the status and the reputation relationships among agents is a necessary but not sufficient condition for IPSL to occur.

3 Computational experiment and model

To further explore the analytical implications of the relationships hypothesised in the propositions described herein, an ABM was designed to perform a computational experiment.

3.1 Hypothesis and model overview

The hypothesis of the computational experiment is that an organising system of agents can in the aggregate learn to recognise the highest value projects from a field of randomly generated projects drawn from a uniform distribution. These would be the projects that the organisation would choose to staff and execute even though no one single agent would have all of the information needed and the wherewithal to make a good decision. In particular:

- 1 Ws, who come across the opportunities in the environment have no way to assess their relative value.
- 2 MMs, interact with Ws but have no direct visibility into the environment itself and have a varying ability (called 'intelligence') to assess the value of the projects presented to them for sponsorship.
- 3 TMTs, must select from among projects presented by MMs with no direct knowledge of their value. Rather, the TMTs must decide based upon the relative reputation of the MMs with regards their ability to select projects, an attribute of MMs that is determined by the agent's prior record of success.

Thus, it is hypothesised that information about past performance can be embedded in the status structure of the MMs through a learning algorithm. The algorithm changes the

system's structure to ratchet remembered success as the complex social network within the system learns from experience and adapts to the environment.

To support the hypothesis, one must show that the statistical mean of the projects that are selected and executed, called 'fitness', when the learning algorithm is turned 'on' is greater than would otherwise be the case, and further that this difference in fitness is statistically significant.

To test the hypothesis, an ABM was designed such that an organisation's fitness can be tested in one of two modes: either projects are presented, sponsored or approved at random throughout the organisation, or the selection process occurs with the learning algorithm turned 'on'. Namely, when the IPSL algorithm is 'on':

- 1 information about opportunities is brought to MMs by Ws according to their relationships with those MMs (Proposition 1)
- 2 the TMTs use the status of sponsoring MMs to select projects (Proposition 2).

Finally, to distinguish incremental fitness that results directly from intelligent MMs from incremental fitness which results from the status structure alone, the model also has the capacity to look at the system in two hidden-layer modes: either the MMs have no intelligence at all, accepting projects from Ws on a first-come-first-served basis, or the MMs exhibit varying degrees of intelligence with regards project selection. This latter selection can be used to isolate the fitness benefits that accrue to the system due to the intelligence of agents alone, independent of their status.

When the model runs, projects are generated at random with values ranging from 1 to 100 drawn from a uniform distribution. Ws randomly encounter these and carry knowledge about them to MMs to whom they are connected. Each of the MMs who is not currently busy, and who has encountered Ws with new projects, selects a single project to sponsor and presents it to the TMT for approval. Every project has a life-span and once a project is sponsored, it keeps the MM busy for its entire life span whether it is approved or not. Only a limited number of projects are approved in a given time period, and the projects that are approved are determined by the relative status of the MMs sponsoring the projects. When approved projects are completed, their values are averaged. This average is tracked as the instantaneous fitness of the system for a given time period. For analysis purposes, the model tracks moving averages over 10 projects, 100 projects, 1,000 projects, and 10,000 projects. These moving averages are plotted. An example is shown in Figure 2.

3.2 *The model implementation*

The analytical consistency of the above theoretical discussion was tested using a preliminary ABM to illustrate its basic functions. The model was built with NetLogo¹ to illustrate the IPSL learning algorithm in an ABM format and is intended as an illustration of these ideas not as an analytical tool.

3.3 *The agents and the environment*

The model consists of four classes of agents interacting on a toroidal grid as their environment. The environment produces a user-defined number of value-creating *opportunities* (stimuli), called 'opportunity agents' or just 'opportunities', at regular

intervals and at random locations on the grid. Each opportunity agent only exists for a random number of time steps up to a user-defined maximum. If executed, each opportunity agent or project produces a random number of value units drawn from a uniform distribution in the interval [0,100]. Thus the size of the environment, the number of opportunity agents, how often new opportunities appear, and the maximum number of time steps each opportunity takes for completion are parameters in the model that determine the relative resource-richness of the environment.

As described above, opportunity agents come and go and also vary in value. The challenge for the system is to learn how to find and to capitalise on the highest value opportunities. Because the value of each opportunity agent varies randomly from 1 to 100, an average performance by the system would equate to approximately 50 units per project, and this would imply that neither intelligent search nor learning has occurred. There is no advantage to organising under this scenario; thus, it is the null hypothesis.

The number of Ws (input layer Worker agents) is a parameter determined by the user. Ws randomly move on the grid. When they encounter an opportunity, they become aware of it and carry it with them as an idea (it turns yellow in the model) as long as it is viable; no other agent can independently find the opportunity in the environment. Ws can collect a user determined maximum number of unexpired opportunities at a given time step. As such, Ws begin as homogeneous, but they become heterogeneous as they collect unique knowledge about opportunities in the environment. These Ws make no judgements about the value of the various opportunities and collect them on a first-come-first-served basis. To simulate interdependent tasks, Ws cannot execute or otherwise take advantage of opportunities without others. To execute the opportunity, the W must find a sponsoring MMs agent and eventually a TMT to approve the project.

The user also determines the number of MMs. These are the agents that act as the hidden layer. As such, they do not take inputs directly from the environment but only get opportunities from Ws. (In other words, they cannot find opportunity agents in the environment on their own.) To distinguish between intelligent interaction by MMs as they consider projects versus purely random exchanges between Ws and MMs, the user determines whether MMs, are ‘intelligent’ or not, and whether intelligence varies randomly among MMs from 0 to 1, or if all MMs exhibit the same, maximum level of intelligence, a value of 1. In this context, being ‘intelligent’ means that the MM agent has a level of visibility into the value of the opportunities that it sees and can therefore intelligently choose which ones to sponsor.


The Ws to MMs interactions work as follows: based upon a user selected interaction strategy, Ws seek out potential sponsoring MMs to whom they have a connection and positive experience (as described below) and present them with the opportunities that they have found. In other words, if Ws engage in ‘preferential selection’ with respect to MMs, then a scale-free structure might emerge over time (Barabasi, 2002). MMs evaluate the opportunities that they see based upon their intelligence as described above. For example, in the case where each MM has a randomly generated ‘intelligence’ from the interval [0, 1], the MM intelligence value determines how accurately the MM prioritises the opportunities that are presented to it with respect their actual fitness value.

Intelligence works as follows: if the intelligence value is set to 0.4, the MMs perception of the value of an opportunity is equal to 0.4 times the actual value of the opportunity plus 0.6 times a randomly generated number between 1 and 100. Thus the MM agent can either overestimate or underestimate the actual value of an opportunity. Perfect intelligence means perception matches reality. For simplicity, each MM is limited

and can only sponsor one project at a time, and once a project is selected, it is sponsored until it expires. As hidden-layer agents, however, MMs cannot execute the strategy. They must seek approval from the TMTs to execute and realise the project's value. Once a project is sponsored, however, the MM stays with that project until it is approved and executed (and creates a fitness value), or until it expires unapproved (and creates no fitness value). MMs that have no projects approved by TMTs create no fitness value for the organisation. They thus have no opportunity to improve their status.

The current model has only one TMT agent. At each time step, the TMT agent chooses a user-defined number of projects to approve. It does so based only on the relative status of the sponsoring MM. In other words, projects are chosen not by their merits, but by the status of the sponsoring agents. The status of the MMs as perceived by the TMT relates to the prior success of the MMs at identifying value-creating projects. Past success implies that future projects are more likely to be approved. For modelling purposes, when the model is initiated, each of the MMs is assigned a status from a random draw taken from a uniform distribution from [0,1].

Figure 2 Fitness as moving averages of projects selected (see only version for colours)



Notes: As described in the text, this graph shows the average fitness over various moving average ranges of recent projects implemented by the organisation. As expected, the least variance is in the solid white line representing the moving average of the last 10,000 projects versus the black line which is the moving average over 1,000 projects or the two shades of grey dots representing 100 projects (dense clustering) and 10 projects (loose clustering) respectively. As referenced, the average project has fitness of 50 (shown as the lower dotted white line). Any result greater than 50 demonstrates intelligent selection. When projects are selected randomly, the moving average hovers around 50 from time zero to approximately time 5,000. At time 5,000 middle manager intelligence is turned 'on' meaning MMs can assess the value of the projects that Ws bring them with varying degrees of accuracy, but the learning algorithm of the system remains 'off'; fitness improves to about 59. At approximately time 12,000 the learning algorithm is activated. It takes until about time step 18,000, but the system eventually (crosses the higher dotted white line at 60 and) resets its status and reputation levels reaching its fitness peak at about 62.

When projects are approved, they are executed, and the predetermined fitness value is realised by the organisation. At any given time step, one or many projects may be realised depending upon the number of projects that have been approved and executed. The resulting value for that time step is the simple sum of the completed projects since this simplified model assumes there are no interactions in the value creation that results from completing multiple projects together. The fitness of the organisation at a given time step is defined as the arithmetic mean of the projects completed at that time step. Because this value is very volatile, moving averages are also calculated and displayed. As can be seen in Figure 2, the moving average of the last 10,000 completed projects is displayed as the solid white line. This value is reasonably stable and can be used to compare fitness across model runs.

3.4 *The neural network embedded in the model*

Opportunities are informational ‘stimuli’ in the environment. These are recognised by the coarse-grained organisation and are interpreted as requiring a ‘response’. This stimulus-response dynamic is actualised within a human organising form through a process that is analogous to a feed forward neural network embedded in the organisation’s influence structure². In this framework, Ws represent the input layer of the network, sensing a stimulus when it is encountered, but they are also not in a position to recognise its significance nor to formulate a response. These agents *feed forward* information about the stimuli they recognise to MMs according to a strategy that incorporates the relative weighting in the connection between the input layer Ws and the MMs.

MMs only see opportunities presented to them by input layer agents, the Ws. They do not see the environment directly, and they cannot directly produce an output. Thus they are part of the *hidden layer*. There may be differentiation among MMs, but in any case, MMs are dependent upon their network of Ws for their ideas. Each MM presents the opportunity it has chosen to sponsor to the TMTs for approval.

All of the MMs approach the TMTs with their sponsored projects. This is another *feed forward* process whereby the various and sundry stimuli from the environment are disambiguated for a single informational response, a decision. When a project is approved, acting on the project becomes the network’s unambiguous output to the environment; it is the organisation’s *trial response* to stimuli in the environment. Thus the TMTs represent the *output layer* of the neural network, and its output is a decision to implement a trial response to the affecting stimuli received by the input layer of Ws.

The final step, IPSL, a process analogous to *back propagation learning*, occurs when a *learning algorithm* is implemented. The learning algorithm in this model has two elements. Firstly, the relative status levels of MMs impact which projects are selected. Thus, if the organisation is good at approving the highest value projects proposed by the MMs (recall that MMs ability to perceive the inherent value in a project can be made to vary through a parameter setting), then the organisation’s performance will be better. Thus *assignment of status* to MMs is a key element of the learning algorithm. Secondly, as they are hidden, MMs are dependent on their workers to bring them ideas. Workers bring ideas to MMs based upon that MMs reputation in the eyes of the W. Thus, *assignment of reputation* by workers to MMs is also a key element of the learning algorithm. Both of these elements – status and reputation – measure the strength of ‘synaptic connections’ between agents in the network.

4 Illustrative results and discussion

This section describes the outcomes of the computational experiment and interprets the results from three separate analyses. The first describes an example model run which illustrates the potential for the emergence of IPSL and demonstrates its benefits. The second involves hundreds of model runs and explores the statistical significance of the fitness benefits provided by IPSL. The third analysis looks at agent intelligence and status and explores relationships among these attributes. The section ends with a discussion of some practical implications of these findings.

4.1 Illustrative model run and interpretation

The left side of Figure 2 shows an example of the model's output when the system selects projects randomly rather than in learning mode. As can be seen, with random selection, the average fitness hovers at about 50, the mean fitness of all of the projects. (This is best seen as the solid white line which averages the last 10,000 projects; note that the broken white lines in the figure identify the fitness levels of 50 and 60 for reference purposes.) The fitness value for each project is a random draw from the interval [0,100] with a uniform distribution. This means that in the random case, the selection of projects occurs without any intelligent deliberation, and thus the results are random (i.e., the mean varies around 50). More precisely, there is neither intelligent selection by MMs – projects are selected by MMs at random – nor does the organisation itself select intelligently by choosing the projects that are championed by the MMs who had been most successful in the past (and thus had achieved higher status) presumably by being better at selecting projects. It is clear that for this case there is no evolutionary fitness benefit to organising; random selection would do just as well. The randomised setting of the model with a mean value of 50 thus represents the null hypothesis; it is equivalent to the toss of a fair coin with regards picking projects.

As can be seen in Figure 2, at about time step 5,000 the user has turned MM intelligence 'on'. Project selection is no longer random, and MMs are made to select projects using their intelligence. In this middle section of the display in the figure, the MMs relative status and reputations remain constant and are not altered based upon success. As can be seen by looking at the solid white line, the moving average over the 10,000 projects improves nine points to approximately 59, a statistically significant (to the 1% level using the Z-test) improvement when compared to the null hypothesis. This occurs because MMs begin to apply their own brain power internally in the system. They sift through projects that are brought to them by Ws and choose to sponsor the ones they believe have the highest value potential. The benefit of this collective deliberation is then aggregated to the system level as the simple sum of all of the MMs individual performance. Collective fitness is the arithmetic mean of the projects that are completed.

Clearly, smart MMs that use their intelligence to select projects provide better performance than would be the case if they selected their projects randomly. However, several questions remain: Can the broader collective do even better? Would the system benefit by organising its internal structure around MMs who are intelligent and use that intelligence to select projects? To test this idea, the system can be made to identify MMs who have identified the highest value projects in the past. Once identified, these agents are given higher status thus increasing the probability that their projects will be approved

during future time steps. By increasing the percent of projects from better performing MMs, the system achieves better average fitness overall. It learns to use intelligence differences among its MMs as a ratcheting process to remember successful champions and to provide fitness advantage.

The organisation does this endogenously. It identifies and selects the most intelligent MMs as they perform over time; those that succeed are given enhanced status, increasing their influence on decision making. As events unfold, the system continues to reinforce or diminish each MMs status selecting those agents with recent success. This ratcheting further increases the probability that better projects will be identified and selected as the organisation responds to changes in its environment.

This incremental ‘system intelligence’ over and above smart MMs is shown in Figure 2 beginning at time step 12,000. One can see that fitness gradually increases by an additional 2.5 to 3 points from 59 to between 61 and 62. This additional increase from IPSL alone adds to the total improvement over and above the random case. Both of these increments are statistically significant (at a 1% level using a Z-test). Accurate processing of feedback and effective adjustments to the relative status and reputations of those accountable for prior success affords the system incremental fitness advantage versus what would otherwise be the case.

Note that the model run shown in Figure 2 represents tens of thousands of projects selected and executed by the system in each of the three cases described above:

- 1 random selection of projects
- 2 smart MMs
- 3 smart MMs together with the presence in the system of IPSL.

The differences in mean values in each case are statistically significant to the 1% level using the Z-test. However, every individual run is unique and highly contingent. The projects, the distribution of agent intelligence, the initial network connections, and initial status levels are all selected randomly from a uniform distribution. Thus, every run has unique initial conditions. As a result, the model exhibits stochastic behaviour throughout its life; each individual case is unique and is therefore only illustrative of many possible cases. Next the results of hundreds of model runs are analysed.


4.2 Robustness of the fitness advantage of IPSL

In a separate analysis (shown in Figure 3) with intelligent MMs, the results of 400 different stochastic runs (200 for each of two cases) were compared: 200 runs with the learning algorithm turned on, and 200 runs with it turned off. The difference in the mean fitness across these 400 runs (between the two cases) was measured at each time step over a period of 5,000 time steps. The average difference between these two cases across these 400 model runs is shown in Figure 3 at each time step. Note that this analysis simulates 400 organisations sorting through one million projects for each case, each of these projects drawn at random from a uniform distribution. The main result of this study is that the organisations with the learning algorithm turned ‘on’ were found to select the projects of higher value than could be done by intelligent agents alone, and this difference in fitness was statistically significant.

As can be seen in Figure 3, the average difference between the means of the 200 model runs for each case when they are compared over each of 5,000 time steps is 2.64.

(Note that any difference greater than 0.6 is statistically significant to the 1% level using a Z-test. Also note that a difference that was greater than 0.6 was found for 70% of the 5,000 time steps.) However, one can also infer that this performance improvement is apparent from within a very noisy signal. Thus, it could be quite difficult to isolate the precise level of improvement at any particular point in time during any particular instance of the model running; the result is a systemic improvement in performance, but one that may be hard or even impossible for any individual with only local information at a single point in time to isolate absent a holistic perspective that follows overall system performance over time. This finding is consistent with computational modelling mentioned earlier where non-linear processes and contingencies make it very difficult for managers embedded in a complex social network to select good projects (Siggelkow and Rivkin, 2005). The point here is that it is easier, and perhaps more efficacious to select people that have been successful in the past, and then approve their projects.

Figure 3 Plot of difference in mean fitness with IPSL active or ‘on’ versus ‘off’ in model for 200 runs for each case and plotted over 5,000 time steps (see online version for colours)



Notes: This graph represents the difference in mean fitness at 5,000 time step for two cases described in the text: the case with the learning algorithm off is subtracted from and the case with learning turned on. Each case is the mean over 200 runs at each time step and the difference in means is plotted. When the difference in means are averaged over 5,000 time steps the average improvement is 2.64 points, but the volatility is quite large.

A significant result from the model, then, is that the strategies employed by organisations for acquiring and retaining talent and for assigning reputation and status among employees do matter. As was shown in the prior section, when a learning algorithm is not activated and when intelligence in MMs (its middle managers) is also not activated, the mean fitness value created by the organisation’s various projects varies around 50. By allowing the intelligence feature in MMs, the mean fitness of the system increases to about 59. To further improve performance, an organisation should enact an effective reputation and status assignment process to act as a learning algorithm for the

organisation, increasing its mean fitness beyond even the level that intelligent MMs enable. It also enables the organisation to continuously gather new information as a means to keep current with changing conditions in the environment. This provides support for Proposition 3 and is an example of IPSL.

4.3 The distribution of status and its relationship to agent intelligence

Figure 4 shows outputs related to agent status when the ISPL learning algorithm is operating. As can be seen in the top panel, there is a power-law relationship in the histogram of agents at each level of status when IPSL is active. If this relationship is verified empirically, this would provide tentative support for Proposition 4. The bottom panel in Figure 4 demonstrates a positive relationship between an agent's ability to correctly perceive value in projects, the trait called 'intelligence' in the model, and the status that an agent achieves in the collective. If this outcome were to be verified empirically it would support Proposition 5.

One thing that is enlightening is that many intelligent MMs who could contribute at high levels under other circumstances, do not achieve high status during these model runs. [In the above analysis, agents who fail to achieve any status are removed from the data set (because they have too few projects approved early in the process)]. This seems to imply that the prospects for personal success experienced by each individual agent, might be highly contingent depending upon the specific initial situation within which they are embedded. Since resources are assigned based upon prior success, new agents must achieve initial success if they are ever to improve, or even maintain, their status.

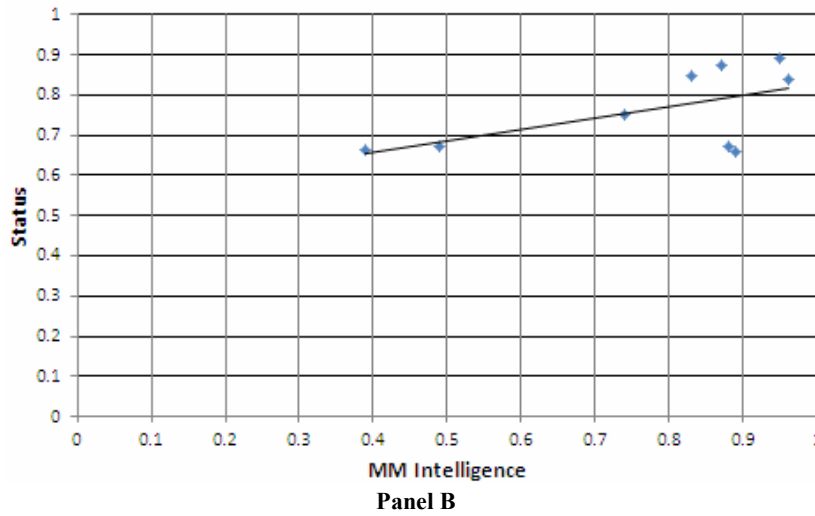
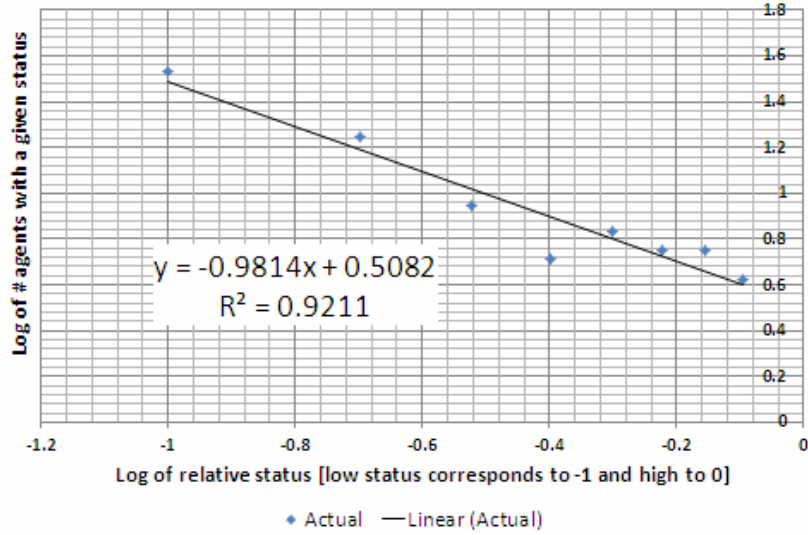
It is also significant that a learning process does not always emerge. Because the initial assignment of status is random, a bad structure can persist. If intelligent agents have no status initially, they never have their projects approved and thus have no way to improve their relative position. Although the scenario used in Figure 2 for illustration shows a case with a visible positive outcome, the reader should be aware that each run is a unique case and the results vary greatly due to the stochastic nature of the process that is modelled.

4.4 Practical implications for managers

This modelling result has important practical implications for managers. In competitive environments where every bit of incremental fitness is needed to gain relative advantage, organisations must thoughtfully manage reputations and assign status if this process is to act as a learning algorithm for the system. Businesses can easily err in assigning status to middle managers based upon apparent versus actual success, for example, or by letting reputations rest upon things other than success at improving the organisation's 'fitness'.

If status or reputation assignment is flawed, an organisation can easily get caught in a counterproductive status structure where agents with low intelligence (in the sense of picking projects) maintain a high level of status by virtue of their continued access to resources and the ability to 'spin' outcomes so as to appear successful. This pathology can cripple the adaptability of the system because the best projects do not get an airing, and the middle managers who identify opportunities with the best prospects do not get their projects approved. Because of this, managers with the best potential cannot improve their status and are thus unable to help the system learn and improve performance. Once a poorly performing structure develops, it can be self-reinforcing and difficult to change.

Figure 4 Agent status relationships from a series of simulation runs are shown as of time step 3,500 (see online version for colours)



Notes: (1) Panel A shows the log-log plot for the mean values of MM agent frequency for each status level calculated over five model runs of 200 MMs. A power-law relationship is demonstrated with an exponent of -0.98 and an r -squared of 0.92 . Panel B shows the status of MM agents versus their ability to accurately perceive the potential value in randomly generated projects, an attribute called intelligence. A positive relationship between intelligence and status can be observed in the chart once very low status individuals are removed from the data set.

(2) The plot in Panel A excludes the point of maximum status due to special modelling code for high status. When this point is included, the r -squared is 0.77 and the slope is -1.26 .

(3) In Panel B, the analytical model for the line is: $y = 0.29 * x + 0.54$ which has an r^2 of 0.35 .

Finally, there is a lesson for individuals seeking to advance into higher status management positions. This study emphasises the importance of achieving early, visible success within one's organisation. If a young manager – no matter what his or her potential – gets caught up in an organisation that does not offer opportunities to prove oneself and thus to gain the notice of those who proffer status, he or she should seek a different opportunity either by changing jobs or moving to another department. This study implies that it is easy for talented individuals to get lost in complex organisations if they do not have early success that brings with it access to the organisation's resources.

5 Concluding comments and future directions

This paper describes a mechanism called IPSL in organisations that may account for a previously unrecognised, qualitatively different type of organisational learning that occurs outside individual agents. IPSL uses scale-free complexity dynamics within complex social networks to enable fine-grained individual actions to have broad implications across scale. In so doing, the individual agent can impact coarse-grained performance by leveraging its network position and connections. A simple simulation was used to demonstrate the analytically adequacy of the model with respect to theory (McKelvey, 1999). It shows how IPSL might occur in organisations to enable adaptation, and it demonstrates that emergent agency at the system level is a possibility.

This study is intended to demonstrate the possibility of emergent IPSL in organisations; it does not attempt to explore its mechanisms in detail. Further research is needed to develop the simulation model so that it can be used to determine the exact conditions under which this type of emergence is most likely to occur. In such a model, different approaches could be tested to explore, for example, the relative impact of the number of projects approved per time step on the emergent structure over a given time period, or simulations could explore the relative impact of various parts of the learning algorithm on emergent IPSL. Empirical data could be gathered and used to develop a baseline scenario to validate the model, and the initial conditions could then be varied to test alternative scenarios that might improve performance. For example, a realistic substrate network that matches an actual organisation could be used in models to identify the specific topology that might characterise a situation that would enable IPSL. In this way, future simulations that closely match real world settings could be developed. More generally, these ideas should be taken to the field to evaluate their efficacy in real-world organisations.

As a tentative first step, however, the results of this simulation square with conventional wisdom, and indeed with intuition, that when managing within a complex organisation a first step is to 'hire the best and the brightest' at all levels of management, but particularly at the middle management level. The second step is to put the right people into the right positions. In this simple simulation, hiring intelligent middle managers has the effect of increasing mean fitness over random performance levels by about nine points, increasing from 50 when projects are drawn at random from the interval [0,100] to about 59.

Likewise, at the executive level, the benefit of supporting the projects proposed by high status individuals and then holding them accountable is also clearly demonstrated with an increase of three additional points, increasing fitness from 59 to 62. This three point advantage, a 33% increase in incremental fitness, represents additional fitness

gained by organising intelligent MMs into a complex influence network. Thus, for the case illustrated, an effective learning algorithm has 33% greater value than what is provided by intelligent MMs alone. This only matters, however, when IPSL is in place to ensure that the agents relative status is continually correlated with recent prior performance as a means to continuously adapt to changing markets. The insights described here thus have important implications for both research and practice.

In sum, this analysis contributes theoretical support and provides specification around what is conventional wisdom in practice: that executive leadership is about picking and mentoring the best people for the job and that leaders switch their organisations into overdrive and achieve even higher levels of fitness when the TMT is effective at picking strong performers and assigning them influence according to their continuing success. The analysis described herein implies that for the TMT, ‘effectiveness’ means instituting and maintaining an effective learning algorithm that enables IPSL.

Finally, this study helps to clarify the unique contribution of social capital to an organisation’s fitness. As such, it makes a general contribution to complexity science (Solow and Szmerekovsky, 2006) by describing a mechanism of emergent agency in complex adaptive systems of all types.

Acknowledgements

Ongoing analysis and revisions were supported in part through The Knowledge Foundation, Sweden.

References

- Baek, S.K., Bernhardsson, S. and Minnhagen, P. (2011) ‘Zipf’s law unzipped’, *New J. Phys.*, Vol. 13, 043004, doi:10.1088/1367-2630/13/4/043004.
- Barabasi, A.L. (2002) *Linked: The New Science Of Networks*, Perseus Publishing, Cambridge, MA.
- Bossomaier, T.R.J. (2000) ‘Complexity and neural networks’, T.R.J. Bossomaier and D.G. Green (Eds.): *Complex Systems*, pp.367–406, Cambridge University Press, Cambridge.
- Carley, K.M. (2002) ‘Simulating society: the tension between transparency and veridicality’, *Proceedings of Agents 2002*, Chicago IL.
- Carley, K.M. and Ren, Y. (2001) ‘Tradeoffs between performance and adaptability for C3I architectures’, paper presented at the *Command and Control Research Symposium*, Annapolis, MD.
- Cohen, M.D., March, J.G. and Olsen, J.P. (1972) ‘A garbage can theory of organizational choice’, *Administrative Science Quarterly*, Vol. 17, No. 1, pp.1–25.
- Collins, J. (2001) *Good to Great*, HarperCollins Publishers, New York.
- Gell-Mann, M. (2002) ‘What is complexity?’, A.Q. Curzio and M. Fortis (Eds.): *Complexity and Industrial Clusters: Dynamics and Models in Theory and Practice*, pp.13–24, Physica-Verlag, Berlin.
- Goldstein, J., Hazy, J.K. and Lichtenstein, B. (2010) *Complexity and the Nexus of Leadership: Leveraging Nonlinear Science to Create Ecologies of Innovation*, Palgrave Macmillan, Englewood Cliffs.
- Granovetter, M. (1973) ‘The strength of weak ties’, *American Journal of Sociology*, May, Vol. 78, No. 6, pp.1360–1380.

- Hazy, J.K. (2007) 'Computer models of leadership: foundation for a new discipline or meaningless diversion?', *The Leadership Quarterly*, Vol. 18, No. 4, pp.391–410.
- Hazy, J.K. (2008) 'Toward a theory of leadership in complex adaptive systems: computational modeling explorations', *Nonlinear Dynamics, Psychology and Life Sciences*, Vol. 12, No. 3, pp.281–310.
- Hazy, J.K. (2011) 'Parsing the influential increment the language of complexity: uncovering the systemic mechanisms of leadership influence', *International Journal of Complexity in Leadership and Management*, Vol. 1, No. 2, pp.164–191.
- Hazy, J.K. (2012) 'Unifying leadership: shaping identity, ethics and the rules of interaction', *International Journal of Society Systems and Science*, Vol. 4, No. 3, pp.222–241.
- Hazy, J.K., Goldstein, J.A. and Lichtenstein, B.B. (2007) *Complex Systems Leadership Theory: New Perspective from Complexity Science on Social and Organizational Effectiveness*, ISCE Publishing, Mansfield, MA.
- Hazy, J.K. and Uhl-Bien, M. (in press) 'Changing the rules: the implications of complexity science for leadership research and practice', David Day (Ed.): *Oxford Handbook of Leadership and Organizations*, Oxford University Press, Oxford.
- Helfat, C.E., Finkelstein, S., Mitchell, W., Peteraf, M.A., Singh, H., Teece, D. and Winter, S.G. (2006) *Dynamic Capabilities: Understanding Strategic Change in Organizations*, Blackwell Publishing, Malden, MA.
- Holland, J.H. (2001) 'Exploring the evolution of complexity in signaling networks', available at <http://www.santafe.edu/sfi/publications/wplist/2001> (accessed on 22 August 2002).
- Kauffman, S.A. (1993) *The Origins of Order*, Oxford University Press, New York.
- Lichtenstein, B.B., Uhl-Bien, M., Marion, R., Seers, A., Orton, J. D. and Schreiber, C. (2006) 'Complexity leadership theory: an interactive process on leading in complex adaptive systems', *Emergence: Complexity and Organization (E:CO)*, Vol. 8, No. 4, pp.2–12.
- March, J.G. and Olsen, J.P. (1975) 'The uncertainty of the past: organizational learning under ambiguity', *European Journal of Political Research*, Vol. 3, pp.147–171.
- Marion, R. and Uhl-Bien, M. (2001) 'Leadership in complex organizations', *The Leadership Quarterly*, Vol. 12, No. 4, pp.389–418.
- McKelvey, B. (1999) 'Complexity theory in organization science: seizing the promise or becoming a fad?', *Emergence*, Vol. 1, No. 1, pp.5–32.
- Okasha, S. (2006) *Evolution and the Levels of Selection*, Oxford University Press, Oxford.
- Schwandt, D.R. and Marquardt, M.J. (2000) *Organizational Learning: From World Class Theories to Global Best Practices*, St. Lucie Press, Boca Raton.
- Schreiber, C. and Carley, K.M. (2006) 'Leadership style as an enabler of organizational complex functioning', *Emergence: Complexity and Organization*, Vol. 8, No. 4, pp.61–76.
- Siggelkow, N. and Rivkin, J.W. (2005) 'Speed and search: designing organizations for turbulence and complexity', *Organization Science*, Vol. 16, No. 2, pp.101–122.
- Solow, D. and Szmerekovsky, J.G. (2006) 'The role of leadership: what management science can give back to the study of complex systems', *Emergence: Complexity and Organization*, Vol. 8, No. 4, pp.52–60.
- Surie, G. and Hazy, J.K. (2006) 'Generative leadership: nurturing innovation in complex systems', *Emergence: Complexity and Organization (E:CO)*, Vol. 8, No. 4, pp.13–26.
- Uhl-Bien, M. and Marion, R. (Eds.) (2008) *Complexity and Leadership, Volume I: Conceptual Foundations*, Information Publishing Associates, Charlotte, NC.
- Uhl-Bien, M., Marion, R. and McKelvey, B. (2007) 'Complexity leadership theory: shifting leadership from the industrial age to the knowledge era', *The Leadership Quarterly*, Vol. 18, No. 4, pp.298–318.
- Yukl, G. (2010) *Leadership in Organization*, Prentice Hall, Inc., New York.

Notes

- 1 NetLogo is available online at <http://ccl.northwestern.edu/netlogo/>.
- 2 This would almost certainly be true if interaction among opportunities (in the sense of fitness landscapes) was permitted in the simulation. For example, if one opportunity was in the works and a second was presented that offered synergy value above and beyond the simple sum of the two projects, a recurrent network could remember that the a project from a prior time step was in process and learn to recognize the synergy value inherent in doing both projects at once. Although interesting and potentially important, this realization doesn't directly further the simple existence argument for IPSL that I'm making herein. Thus, this interesting enhancement is left to future research.